UNIVERSITY OF
BATH

University of Bath

**Department of
Computer Science**

UNIVERSITY OF
BATH

Technical Report

Ph.D. Dissertation

A design framework for pervasive computing systems

Vassilis Kostakos

A DESIGN FRAMEWORK FOR PERVASIVE COMPUTING SYSTEMS

Vassilis Kostakos

A thesis submitted for the degree of Doctor of Philosophy

University of Bath

Department of Computer Science

November 2004

Aristotle and Plato in a detail of Raphael's fresco *The School of Athens* are depicted in a way that symbolizes their approach to knowledge and the constant struggle of science. Plato points his finger to the heavens, while Aristotle gestures toward the earth. Plato searched the heavens for the ideal; Aristotle looked to nature for pragmatic answers.

TABLE OF CONTENTS

Chapter 3 • Users or Citizens?   57

Chapter 4 • Information Spheres   75

Chapter 8 • Case Study: Designing for a Hospital Environment   159

Chapter 9 • Case Study: The City of Bath   195

LIST OF TABLES

LIST OF FIGURES

LIST OF INFORMATION BOXES

ACKNOWLEDGMENTS

First and foremost I would like to thank my family and especially my father Michael Kostakos for providing me with moral, psychological and economic support throughout my studies. Thanks for your help and guidance. I hope I was worth it.

I also wish to thank my personal supervisor, Eamonn O'Neill for all his support, comments, corrections and ideas. Thank you for spending so much time with me and being the best supervisor I could have asked for. Your constant challenging of my ideas has made me work even harder. I can only hope to achieve your perfectionism.

Furthermore, I wish to thank the members of the HCI Group at the University of Bath, Peter Johnson, Hilary Johnson, Manasawee (Jay) Kaenampornpan, Andy Warr, Peter Wild, Rachid Hourizi, Iya Solodilova, Anne Bruseberg and Stavros Garzonis. You have been my academic family and your help and support has been most appreciated. Special thanks are due to Dawn Woodgate for her ethnographic studies and all her suggestions on citizenship. Thank you for taking time to answer my questions.

I am grateful to the University of Bath and the Department of Computer Science for providing me with funding during my studies and making it possible for me to attend conferences. I also wish to thank the anonymous reviewers for their useful and insightful comments.

Finally, I wish to thank Peggy Kalantzis, Maria Gabriella Gonzales and Eftychia Chatzopoulou for helping me get through the difficult times and making my life colourful.

ABSTRACT

The question this thesis addresses is: How can we design pervasive systems? Designing a system goes much beyond giving building instructions. To design a system, we would like to be able to relate our system to other existing systems. We would like to be able to have grounds on which we can make design choices amongst various possibilities. We would like to be able to learn from existing systems, and thus improve future systems. Finally, we would like to be able to prescribe the required technology, thus pushing the development of technology along the line of satisfying actual needs.

To answer our question, we build on the established HCI tenet that there are three dimensions to situations where humans and computers, people and technology, come in contact: the user, the task and the domain. In this work, we explain why the notions of user, task, and domain are not sufficient to help design the grand vision of pervasive or ubiquitous computing. The concepts we propose instead are citizen, sphere, and space respectively. These three elements form our design framework, based on which we have developed a design tool and method. Our design tool can be used to model and represent pervasive systems, evaluate the potential of privacy breaching, indicate situations where physical interaction with the system is not possible, and inform the designer of situations where cognitive overload could happen. Coupled with our method for inspecting such problems, we show how design choices can be explored, design alternatives evaluated and compared.

We illustrate the applicability of our ideas on several levels. First we apply our ideas to the implementation of a gestural interaction technique. Then we draw on the results of an ethnographic study of the A&E department of a London hospital to propose design solutions. Finally we look at the city of Bath, where we apply our ideas and framework to generate design recommendations.

CHAPTER 1

INTRODUCTION

Pervasive computing is a family of technologies that aims to become part of our everyday life. Such technology will be available to us everywhere, and for any purpose. Being "online" everywhere and anytime is what pervasive technologies are about.

The Internet was termed the "global digital village" by popular media such as television and newspapers. However, the most common way of accessing the Internet, computers, did a pretty good job of keeping it out of our way. Computers were the end of the line: switch the computer off, and the Internet disappears. The Internet, therefore, is not as global as it was hyped to be. Pervasive technologies take a step further by having literally thousands of devices connected to the global network. All sorts of everyday objects will be part of the network: coffee mugs, refrigerators,

sofas, you name it. All these devices and objects simply cannot be switched off. Pervasive technologies will eventually pervade our lives.

Economic and technological developments are encouraging the movement towards a more interconnected worldwide community. But on the other hand, recent protests illustrate a distrust on the part of citizens towards some of these developments. Concerns are raised over issues such as privacy and rights such as freedom of speech and democratic participation. Yet, there are huge benefits, realised and potential, that come with the increasing ease with which people may communicate and collaborate across borders and oceans. Why is it then that such new technology is not universally welcomed, and how does this affect the design of pervasive systems?

An important contributing factor is that the design, implementation and deployment of technologies are being driven almost universally by commercial and governmental (primarily military) needs, rather than civic needs. Hence, the technology available to citizens worldwide fails adequately to serve the needs of society in general; instead the focus is on the needs of commercial or governmental organisations. Our initial motivation in the work presented here is that civic needs have to be addressed by any pervasive system aiming to be successful on a large scale.

The discipline of Human-Computer Interaction (HCI) has potentially much to contribute in moving us closer to the vision of widespread, if not universal, access to computing resources. HCI may contribute from informing the design of relatively simple interface features, so that individuals can easily access systems at the point of use, to ensuring that the wider social implications of such systems inform their design and development. While much of HCI has focused on interaction design concerns, there is also a tradition of considering the broader social concerns.

In our work we have extended previous HCI theory and practice in order to account for the gap between technology on one hand and social and civic needs on the other. Enormous potential benefits are available for delivering computing resources throughout society but with very high associated risks including, for example, loss of privacy and great potential for abuse of power. These risks are mainly the result of a lack of understanding.

1.1 Pervasive technologies are here, the understanding is not

In the 1970s and early 1980s, organisational intranets were a rarity and the Internet did not exist in its current global form. Back then, designing a system that supported an entire workplace and its diverse activities was large scale.

Now, for the first time, we do have the technological potential to enable global infrastructures for computing support. What we do not yet have, however, is the theory, tools or practices to design systems that can exploit this technological potential to deliver truly pervasive systems to serve the needs of society in general; that is, to deliver usable, accessible computing resources on a large public scale.

The lack of theory, tools and established practices has left the field of pervasive technologies rather fragmented. A good endeixis of this fragmentation is the list[1] of various names or expressions that have been coined to describe the general research domain of pervasive technologies:

- Ubiquitous Computing (Mark Weiser, Xerox PARC 1988)

- Calm Computing (John Brown, Xerox PARC 1996)

- Universal Computing (James Landay, Berkeley 1998)

- Invisible Computing (G. Barriello, UoWashington 1999)

- Tangible Computing (Ishii, 1997)

---

1. Many thanks to Alois Ferscha for compiling most of this list.

- Pervasive Computing (Academia, IBM 1999, SAP 2000)

- Context Based Computing (Berkeley/IBM 1999)

- Hidden Computing (Toshiba 1999)

- Post PC Computing (Popular media)

- Ambient Intelligence (European Commission, FP5)

- Everyday Computing (Georgia Tech, 2000)

- Sentient Computing (AT&T, 2002)

- Autonomous Computing (IBM, 2002)

- Amorphous Computing (DARPA, 2002)

- Spray Computing (Zambonelli, 2003)

Each of the above terms has a slightly different focus, but the number and variety of terms indicate that this is still a very young and very active research area. These research areas are closely related to what we call mobile computing. Projects in these areas utilise existing mobile technologies and devices. For the purposes of this thesis we will be referring to those technologies that deliver pervasive access to information as pervasive systems.

We have seen considerable advances in technical experience in implementing and configuring pervasive technologies and environments. Unfortunately, theoretical development has not kept pace with technical development, a problem of very long standing in HCI [Barnard, 1991]. While our experience with the technology becomes greater, the technology itself advances and our techniques are fine-tuned, yet our understanding of pervasive systems has not substantially improved. We still have little idea what it means to have a truly pervasive system, with wide -- ideally universal -- physical coverage, access and usability, or how we can achieve that.

To date, the designs produced within pervasive computing (and related fields), as well as the tools, ideas and frameworks to support these designs, mostly address

commercial and government/military interests. Such interests and demand drivers are very helpful in the development of technologies. However, they may inspire concerns about the range of uses to which they are put, especially where those uses impact directly and personally on members of the public. For example, Benetton, Gap and Wal-Mart have successfully used RFID tagging on warehouse pallets for some time but all three companies abandoned prototype trials of RFID tagging individual goods in the face of customer concerns about privacy [McGinity, 2004].

Mobile and pervasive computing systems are becoming increasingly significant in our lives and it is, therefore, vitally important that we understand these systems and know how to design them to serve us best. Flawed design of such systems will have serious consequences for individuals, for groups and for wider societies. If we are to allay the fears of citizens about the potential dangers and abuses of pervasive computing, there needs to be a balance between commercial interests, government interests, and the interests of citizens themselves. Striking a balance amongst these competing interests can result in designs, tools and technologies more appropriate for civic interests, yet able to absorb and reflect commercial interests, as well as government interests, such as national security. In order to achieve such a balance, we must base our designs and design tools on sound fundamental understanding of pervasive systems that take into account social interests and their competition with commercial and governmental interests.

1.2 Our vision: public pervasive systems

Our vision for pervasive technologies is to ultimately provide truly pervasive public systems as a public service or public good to society. By pervasive systems we refer to pervasive access to information (as opposed to systems that use actuators to change the environment - for example systems that open doors or turn on the lights). By truly pervasive, we mean systems that pervade the *physical*, *social* and

Mark Weiser

"The most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it". This is perhaps the most popular quote within pervasive computing research. It is due to Mark Weiser [1991]. Mark Weiser was best known for his advocacy of "ubiquitous computing", a concept he first proposed in 1988. The idea of ubiquitous computing built on Mark's earlier research on human-computer interaction, and was further influenced by Xerox PARC's work in networking, the ethnography of computing and workplaces (and its critique of traditional computer design), and graphical user interface research. Building on "a new way of thinking about computers in the world, one that takes into account the natural human environment", Mark hoped to create a world in which people interacted with and used computers without thinking about them. He proposed wall, pad and tab sized computer devices for various everyday uses. Ultimately, computers would "vanish into the background", weaving "themselves into the fabric of everyday life until they are indistinguishable from it". On April 27, 1999, Dr Mark Weiser, Chief Technologist at the Xerox Palo Alto Research Center, passed away.

cognitive environments. Although necessary, it is not sufficient for pervasive systems simply to be available everywhere. The two additional dimensions we have mentioned, social and cognitive, play a very important role in the success of pervasive systems. For these systems to become part of everyday life, they need to address and pervade the way we think (psychological environment) as well as the way we behave and communicate with others (social environment).

In terms of these systems pervading the social and psychological environments, we distinguish between *domestic* and *public* pervasive systems. This distinction reflects the difference between, on one hand, the currently dominant implementation of pervasive systems in tightly constrained domains such as the home and, on the other hand, the truly pervasive systems that could potentially be made publicly available for general use. We envision a public pervasive system as a system that anyone may use, without regard for the physical location or identity of the user.

In this definition, domestic pervasive systems typically are owned by private individuals or companies, much as current domestic appliances or ISP arrangements. Public pervasive systems may follow an open source model and have no single owner or may be owned by government or communities for the public good, simi-

lar to current council housing and Housing Association arrangements. Domestic pervasive systems are small-scale. They are the smart buildings and smart cars of current pervasive computing implementations. Public pervasive systems are very wide area, providing coverage to entire communities and societies. Domestic pervasive systems are optimised for particular functionality to support specific purposes. In the main, these will be defined by the owners in terms of the services they offer, with some user customisation, much as current desktop software applications. Public pervasive systems need to be much more flexible, in order to offer useful, usable computing resources to the indefinitely wide range of potential users, individuals and groups, performing an indefinitely wide range of activities. The main characteristics of both public and domestic pervasive systems are summarised in Table 1.1 on page 7.

Physical limitations play a central role in limiting the potential success of a pervasive system. Some of these limitations may be overcome by providing a system that offers very wide area coverage. Our envisioned public pervasive systems may offer

---

TABLE 1.1: Characteristics of public and domestic pervasive systems.

Public pervasive systems are open, flexible, public systems. On the other hand, domestic pervasive systems are private, closed and designed for specific environments.

	Public Pervasive Systems	Domestic Pervasive Systems
Ownership	Owned by the community, government etc. Can be used by anyone who is a member of the community.	Private or corporate ownership. For use by the owners such as members of the family, company, organisation, etc.
Coverage	Large-scale. Public areas such as squares and parks, social units such as towns, cities and countries.	Small-scale. Specific locations such as a house, company headquarters, building complex.
Functionality	Flexible	Optimised for specific purposes

coverage for an entire city or even a whole country. In line with our definition, this wide-area coverage is a minimum requirement for truly pervasive systems.

With wide coverage, however, comes a complex set of requirements. Such a system will offer coverage to a wide range of people, who are in a wide range of locations and situations, and who will probably wish to perform a very wide range of tasks. In such a public setting, social requirements and constraints must be taken into account. This implies, for example, that a pervasive system should be compatible with (or at least not compete with) other pre-existing social and non-social systems in the environment. Pervasive systems should be introduced taking account of existing social models and norms, so as to avoid failure due to their lack of touch with social reality. Many similarly ambitious projects, technologies and proposals have failed in the past because they were out of touch with reality and their contemporary social milieu and situations [Schuler, 2001].

Furthermore, such wide-ranging systems with ambitious goals of being used in many aspects of everyday life are much more than simply software; they have been termed social software. According to Shirky [2003], designers of social software are, in spirit, closer to political scientists and economists than to compiler writers. This comment reflects the importance to society of such systems, and highlights some of the non-technical areas that must be considered in the design of such systems.

1.3   A disjoint research area

How can we move towards our vision of designing public pervasive systems? Currently, pervasive systems are in their infancy, and thus our attempts at designing such systems can at best be described as spasmodic. A wide range of projects classify themselves under the umbrella of pervasive computing. These projects have varying foci, including interface design, multimodal interaction, input and output

technologies, hardware of all sorts, middleware, context modelling, communication and networking protocols, database and storage models.

Yet, for all their contribution to technology advances, these systems bear little resemblance to each other, nor do they provide insight and understanding for the improvement of other systems. The absence of a common frame of reference, a common language, a single benchmarking effort, forces most of these projects to be designed, deployed and evaluated in relative isolation. Most projects take a vertical approach, by implementing from scratch all the required elements from interface and interaction to back-end. On the other hand, the research that is indeed focused on specific issues is mostly defined by technological capabilities rather than pervasive computing needs.

1.4 Learning from the success of GUIs

The apparent fragmentation within pervasive computing is caused by, and contributes to, the lack of common concepts, underlying ideas, tools, practices and, ultimately, theory. Unless a set of underlying and unifying principles is developed, we fear that pervasive computing will never enjoy the success of conventional systems with their GUIs, but instead will remain a kind of expensive art, the kind that would be found in rich homes and luxurious buildings. We see great potential in this type of artwork (such as by Ishii and Rekimoto), and do appreciate this approach. Although it offers artistic pleasure to, for example, control sounds by opening and closing physical bottles, it is still much more functional to use a GUI to do so.

Despite the criticism that GUIs receive, they brought a great advantage in the form of an underlying concept: the window metaphor. The popularity of GUIs is far from accidental. Almost any task can be supported using a window. The easy manipulation of digital artefacts, along with the powerful underlying concepts of

GUIs epitomise their great advantages. Another motivation for our work is that for pervasive systems to enjoy the same success as GUIs there has to be a similar set of underlying concepts and ideas that unify this currently fragmented domain.

There have been attempts to provide a unifying framework for dealing with all the pervasive technologies. Rodden & Benford [2003] provide a framework based on the six levels of buildings' lives. Each level is associated with a timescale (e.g. the exterior surface of a building may change every 20 years), as well as a group of stake holders (people who are involved in the organization and execution of changes to a building). Current technologies are classified into the framework according to which of the six levels they address. Rodden & Benford conclude that technologies are currently focusing only on the interior of buildings, with the danger of ignoring the broader settings.

We believe that few pervasive systems currently being implemented address the true pervasive nature and potential of such systems - their ability to integrate into a very wide range of aspects of our everyday lives. For instance, a smart car or a smart

Artistic approaches

A number of research projects have explored artistic approaches and have produced designs based on aesthetic values. For example the Pinwheel project takes fields of pinwheels and explores what arrangements create interfaces that are intuitive and informative for the user while being an ambient source of information. Its current application domains include stock market activity monitoring, web site hits tracking, natural wind movement, and server packet monitoring.

The musicBottles project introduced a tangible interface that deploys bottles as containers and controls for digital information. The system consists of a specially designed table and three corked bottles that "contain" the sounds of the violin, the cello and the piano in Edouard Lalo's Piano Trio in C Minor, Op. 7. Custom-designed electromagnetic tags embedded in the bottles enable each one to be wirelessly identified. The opening and closing of a bottle is also detected. When a bottle is placed onto the stage area of the table and the cork is removed, the corresponding instrument becomes audible. A pattern of coloured light is rear-projected onto the table's translucent surface to reflect changes in pitch and volume. The interface allows users to structure the experience of the musical composition by physically manipulating the different sound tracks.

(Source: http://tangible.media.mit.edu/projects/musicBottles/musicBottles.htm)

house may very well provide a pervasive environment in its own right, but the moment one steps outside its physical borders, the pervasive environment stops being of any use.

Such small-scale pervasive technologies and systems do, however, provide proto-types and testbeds for potentially larger systems. Further research is needed to eval-uate what we have learned from these systems and how we may most effectively move towards our vision of truly pervasive computing. Currently, the focus remains technological, resulting in relatively simple, small-scale situations that are amenable to building heavily constrained systems within limited physical areas or locations.

1.5 Need to focus on design, not technology

An approach for unifying pervasive computing needs to focus on design require-ments, not implementation. Currently, the closest we have achieved is work done on middleware and toolkits to support the creation of pervasive technologies. But could a toolkit really provide an underlying conceptual framework? It depends. Such a toolkit would need to do much more than offer class wrappers and aggrega-tors, object proxies, network layers, etc. We need new powerful metaphors, just as we had the mouse pointer, the menu, the button, the window, the click, the desk-top, the file, the folder. We believe that such metaphors cannot be defined by a simple toolkit; rather a holistic approach is required, even a complete system, like the Xerox Star.

By considering such an initial pervasive system we soon identify an interesting par-adox. Such an initial system would probably offer limited functionality. Therefore, it's pervasiveness would be compromised by this limitation. But a pervasive system has to be pervasive; anything less would render it a simple application, much like today's projects. Yet, a prototype or initial system cannot reach its full envisioned

potential; we just can't get it right straight from the start. How are we then to design pervasive systems?

1.6   Research Question

Given the claims and visions of pervasive systems, the impossibility of studying a pervasive system without having built one, and the wide range of issues that are raised by the design and presence of pervasive systems, the question this thesis addresses is: *How can we design for pervasive access to information*? More specifically,

- We would like to be able to relate our system to other existing systems.

- We would like to be able to have grounds on which we can make design choices amongst various possibilities.

- We would like to be able to learn from every system, and thus improve future systems.

- We would like to be able to prescribe the required technology, thus pushing the development of technology along the line of satisfying actual needs.

All the above currently impossible to do within the domain of pervasive computing systems.

1.7   An overview of this thesis

In the development of the work presented in this thesis, we have had a strategic research goal of developing an applied science of HCI. This involves developing a sound theoretical footing for HCI and deriving design principles for the development of human-computer systems that are theoretically well-founded, empirically tested and operationalised for designers' use. Hence, our research is focused on developing the necessary underlying theory before going on to develop design tools and recommendations for practice. In the work reported in this thesis, we have

extended previous HCI theory and practice in order to account for the gap between technology on one hand and social and civic needs on the other.

Rather than moving towards large-scale pervasive systems in a bottom-up way from today's small-scale systems, associated challenges and ad hoc solutions, we argue for following a top-down approach, drawing on the human-computer interaction lessons that have already been learned in the development of more traditional computer systems.

Our starting point

Our starting point has been the established HCI position that there are 3 dimensions to a situation where humans and computers, people and technology, come in contact: the *user*, the *task* and the *domain*. Notice that from the start there is no mention of the enabling technology, yet as we have already noted much current research is segmented according to technology. From this initial approach, we have augmented these 3 concepts to take into account specific issues.

We have considered the claims and goals of putting the human at the centre of attention, of making computing part of everyday life ("much like electricity" is the common analogy), and of focusing on human needs. In Chapter 2 we explain why the notions of user, task and domain are not sufficient to be helpful with the vision of pervasive computing. The concepts we propose instead are *citizen*, *sphere*, and *space* respectively [Kostakos & O'Neill, 2004b].

An initial framework

The three dimensions we propose - citizens, spheres and space - are a general framework for designing and evaluating pervasive computing systems. By framework we refer to a general set of ideas [Eysenck & Keane, 2000] which are useful in addressing our research question and issues. The three main concepts of our framework, citizens (Chapter 3), spheres (Chapter 4) and spaces (Chapter 5), which we develop

separately over three chapters, help us in developing and designing pervasive systems for information access.

The purpose of our framework is to provide concepts which can be used a priori to raise and predict issues in the design of public pervasive systems. In Chapter 9 we do this in a case study. Furthermore, our framework has given rise to a design tool and method whose purpose is to be used both a priori and post hoc to represent and analyse pervasive systems. In Chapter 8 we use our design tool post hoc to describe existing problems in a London hospital. We also use our design tool and method a priori to design a pervasive system and predict issues with its operation and use within the hospital. Finally, our framework and design tool identify a number of issues which are important in the design of pervasive systems. We address some of these issues in Chapter 7 in respect to user interaction. In this chapter we develop an interaction technique that takes into account the ideas and issues that our framework and design tool raise respectively.

A design tool   The main result of our framework takes the form of a design tool (Chapter 6). This tool, which can be used a priori and post hoc, offers a representation of a pervasive system in terms of the three dimensions we have described. Furthermore, we offer a method for manipulating these representations and exploring alternative ones. Starting with the abstract descriptions of citizen, sphere and space, we perform an analysis of the possible relationships that can exist between the three dimensions, and their impact when and should they appear in a real-world system.

Based on these analyses, for any pervasive system which is described using our tool we can predict or explain a number of issues. For instance, we can spot potential privacy problems in the design of the system. Besides spotting problems, we can get an understanding of the dynamics of the system we are proposing. We can tell whether physical interaction will be possible or not, as well as whether there will be

a cognitive overload on the part of the users. Using our design tool, we can propose alterations to the system we are examining, and re-evaluate the proposed system both for privacy problems or any other reason that caused us to make the changes.

Over the following three chapters (Chapters 3, 4 and 5) we use our framework in three distinct ways which are complementary and range from informing technical solutions (interaction technique described in Chapter 7) to post hoc use of our design tool (A&E case study in Chapter 8) and to a priori use of our framework to generate high-level design recommendations (city of Bath case study in Chapter 9).

An interaction technique

Our framework has been applied in the development of an interaction technique [Kostakos & O'Neill, 2003] which takes into account our ideas about interaction spaces, and how they relate to privacy issues. By interacting with a system, a user gives away information about the interaction itself and the information being accessed. Here we develop an interaction technique that allows users to control how much information is withheld or not about the interaction itself. This interaction technique is gesture-based, and provides a flexible way of providing input to a system in ways that do not compromise privacy requirements. In Chapter 7 we discuss this interaction technique, which uses strokes on a virtual compass as its main metaphor. In this chapter we show how our framework can have direct practical implications, and how some of the abstract ideas we describe can be instantiated and addressed firmly at a technical level.

Ethnographic study

In Chapter 8 we present the results of an ethnographic study of the Accident and Emergency (A&E) department in a large hospital in London. In this chapter we use our design tool to evaluate design solutions and possibilities. We use our design tool to describe the information that exists in this department, as well as the means by which it is disseminated. We then explain in terms of our design tool some problems that the hospital faced, and also identify further issues that our design

tool has raised. We then go on to propose design solutions based on the recommendations of our ethnographic work. In the process of doing so we discuss how our systems level applications and solutions may be brought in and used.

Large-scale case study

To complete our work, we show how we can generate design recommendations using our framework. In Chapter 9 we describe a case study of a large-scale location, the city of Bath. In this case study we analyse various locations within the city of Bath, and then derive design recommendations and predictions. In the process of doing so, we also discuss various issues that are relevant at such a high level, such as how to ensure a balance between commercial and civic needs, and how to assess the success of our system. The recommendations we provide in this chapter may be followed through using the process we presented in the Hospital case study. Thus, we have shown how our ideas may be applied and used to generate overall design recommendations and objectives, to derive design alternatives which can then be evaluated and explored, and to implement system-level solutions.

1.8   Looking forward

In an attempt to answer our research question, "*How can we design for pervasive access to information?",* we now turn to existing work, and see how existing design approaches can be of help in the design of pervasive systems. In Chapter 2 we showcase many design approaches, tools and even full-scale commercial and non-commercial projects that are related to the design of pervasive systems. We relate all these to traditional HCI design practice, and explain why we must take into account the fact that pervasive systems are not the same as traditional systems. As a result, by the end of Chapter 2 we show that traditional HCI practices need to be extended, upgraded if you wish, so that they can remain relevant for the domain of pervasive computer systems.

HOW DO WE DESIGN SYSTEMS TODAY?

2.1   Following the lifecycle of information

In this chapter we show what the field of HCI and related areas have to offer in terms of designing computer systems. To begin with, we state one of the fundamental tenets of Human Computer Interaction [Preece et al., 2002]: *design for a specific user, performing a specific task, in a specific domain*. As we go describe various projects, technologies, methods and approaches, we keep in mind these three dimensions: user, task, domain.

We proceed with our survey by following through a simplified lifecycle of information: first, information must be *generated*, then *distributed*, and then *presented* for interaction with the user. For each of these stages, we show work that is related to

the design and implementation of computer systems and in particular pervasive computer systems. Since the latter is a new area within Computer Science, the full arsenal of systems design has not yet been tried and applied to pervasive systems. Thus, it can be worthwhile to study design techniques that have been used elsewhere, even though they are not of direct relation to the design of pervasive systems.

Having surveyed numerous design approaches, frameworks and implementations, we then turn back to HCI. We show that, in light of our survey of design issues and pervasive computing, the traditional HCI approach of user - task - domain is inadequate. Over the following three chapters (3 - 5) we describe our approach, in the form of a framework for designing pervasive systems.

2.2 Generating the Information

Our goal for generating information should be clear: To generate the right information at the right place and at the right time. Note that right does not necessarily mean correct, but it should be interpreted as appropriate. If we manage to generate appropriate information, we immediately increase the efficiency of our information network, and reduce the chance of providing information that is unneeded, even unwanted.

However, it can be difficult to decide which piece of information is really relevant, and therefore is worth generating. For instance, keeping track of a pillow's temperature probably sounds like a waste of resources, but perhaps it could be useful in certain situations. There can be a virtually endless list of such situations, and this has caused the pervasive systems community many problems. In an attempt to battle this overabundance of information, involving the user in deciding which information to generate is important. A lot of research has been directed towards trying

to address the problem of information overabundance, and generally falls under the research area known as context or context awareness.

Context

There is no standard definition of context, and people seem to define it according to their needs. As defined by [Dey et al., 2001], context is "any information that can be used to characterise the situation of entities (i.e., whether a person, place, or object) that are considered relevant to the interaction between a user and an application, including the user and the application themselves".

Because of its broad range of definitions, the notion of context is being explored by a number of research communities. Lots of different applications have been developed for implementing context awareness. Example applications include active badge call-forwarding [Want et al., 1992] and GroupWear nametags [Borovoy et al., 1998], to even the simplistic example of the light that turns on when the refrigerator door opens.

Most context-aware applications try to change their way of operation according to the situation, or try to retrieve and provide the most relevant information for the user. This often relies on sensors in the environment. However, as noted in [Dey et al., 2001], a large-scale implementation of context-aware applications will require a large number of sensors and services. Therefore, such a bottom-up approach has questionable usefulness. Furthermore, as we argue in Chapter 5, context and domain have characteristics that cannot be physically sensed.

Jonathan Grudin [2001a] explains that technology can enhance efficiency gains, by providing additional context digitally. Despite the efficiency gains however, new technology can disrupt existing practices. According to Grudin, such disruptions may be resolved in three different ways. First, the technology may be completely rejected; secondly, in the long run people may develop new practices and conven-

tions for the use of technology; and thirdly, the technology can be enhanced, often by providing even more contextual information. Thus it is likely that we will try to improve technology by providing ever more contextual information, which will probably never be enough.

This can lead to problems if we accept Grudin's proposition that the process of making contextual information digital can introduce distortions or even fundamental transformations. Because the world provides us with very dense and multidimensional contextual information, we must always decide on the level of context that we wish to acquire, essentially the "cut-off points". Therefore, contextual information that is transformed to digital form is never fully accurate.

Involving the user

One way of compensating for the inherent deficiencies in the generation of digital information is to involve users in the process of generating information. There are situations where, ultimately, the only entity that can decide which information is relevant to you in a certain situation is you. Because humans tend to improvise [e.g. Suchman, 1987], only the fundamentally non-human aspects of information may be efficiently supported by devices on their own.

As Bellotti and Edwards point out [2001], human intervention is required in situations where human intent and interpersonal relations are present. They propose that systems must incorporate two key features: accountability and intelligibility, i.e. that systems "must be able to represent to their users what they know, how they know it, and what they are doing about it". In their framework, Bellotti and Edwards include four design principles to realise their two key features.

- Users must be informed of the system's capabilities and understandings.
- The system must provide feedback, including feedforward and confirmation.
- Systems must enforce identity and action disclosure
- Systems must provide the user with control over the system.

Another domain where user intervention is used extensively is privacy. Different people will have different views on privacy, and therefore different needs and requirements. Furthermore, different environments result in varying levels of privacy being required. Additionally, as noted by Ackerman [2000], the requirement for attention on behalf of devices, when followed by a prompt for consent, can be very disturbing to users. This means that users must shift their focus from their activity [Ackerman et al., 2001]. The extent of pervasive systems, and the number of devices in such an environment increase exponentially this problem. Because of such unresolved issues, the general HCI requirements for privacy need to be better understood, especially in light of pervasive technologies' extension into what hitherto have been public and private areas of life.

2.3   Distributing the information

Having generated the information that we deem appropriate for a user, the next step is somehow to distribute this information. Currently, the information networks available make no distinction between different types of information. They transmit data between computers, devices and other entities connected to the network. The demand for bandwidth seems to keep growing, and the technology seems to respond adequately, but without ever considering what sort of information is being transmitted over the network. We address and discuss this issue by introducing the notion of information spheres in Chapter 4.

A number of research efforts have been directed towards generating and implementing a framework for coping with the distribution and "arrangement" of large amounts of data between electronic devices. Although these systems have the potential to be pervasive, we believe they have not reached the level of true pervasiveness.[2] This section presents some of these efforts at creating a "complete" pervasive system.

The grid

The idea behind computational grids can be considered as a model for pervasive systems' infrastructure. The term Grid was introduced in the late 1990s [Foster & Kesselman, 1999] and refers to distributed networks with shared resources. More precisely, the issue addressed by grids is "coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organizations" [Foster et al., 2001]. Virtual organizations in this definition are the set of individuals and institutions defined by their sharing. In very simple terms then, a computational grid is a peer-to-peer network like Gnutella and KaZaa, where peers' available resources are files, storage, bandwidth and computation. The analogy commonly used compares computational grids to electrical power grids: getting access to computation and data should be as easy and standard as plugging an appliance in a power socket. This analogy focuses on the technical details of making this possible, the back-end.

The same analogy has been also used in pervasive computing. Electricity is such a profound technology that it has effectively disappeared. We do not care about the details and intrinsic workings electricity, we just "use it". The view focuses on how the users perceive and use the technology, i.e. the front-end.

A grid infrastructure provides two essential characteristics for pervasive computing: computational abundance and pervasive presence. Tapping into the computational abundance of the grid could be possible regardless of location. This, in turn, allows us to build mobile devices focusing on interaction, not computation, which in turn results in more mobile and wearable devices simply acting as portals to the infrastructure's data and computational resources.

Grid researchers have built prototypes of the infrastructure required to deploy services on a massive scale, and have even developed economic models for doing so

---

2. In Chapter 1 we argue that pervasive systems should pervade the physical, social and cognitive environments. The projects described here mostly address the physical dimension.

[Buyya et al., 2002; Vazhkudai & Von Laszewski, 2001], but struggle when it comes to providing scenarios and ideas about how to provide people with something useful for everyday life. Currently, the grid does not live up to the expectations we have sketched, mainly due to technological shortcomings. Issues such as network latency, distributed data coherence and efficient process migration are issues that are currently being addressed. Historically, however, similar problems have been overcome, and we can expect the same of the grid in due course.

The PIMA project

The PIMA project at the IBM T.J. Watson Research Center [Banavar et al., 2000] offers a view of pervasive computing that challenges the notions of devices, applications and environment. Essentially, Banavar et al. have decided to adopt the following view:

• Every device offers a portal into applications and data, not simply the user's software collection.

• An application enables users to perform tasks, and is not just a piece of software written to use the device's capabilities.

• The computing environment is the user's physical surroundings enhanced with information, not just a virtual space in which to execute programs.

In order to realise this vision, a new application model has been defined. The model consists of three main sections: the design-time, load-time, and run-time.

Design-time. During design-time, the application is defined in terms of its requirements. This allows applications to be device-independent. The designer must also define an abstract user interface, along with abstract services. This means that the services required by the application are not explicitly stated and that any unanticipated services that exist at run-time, unknown to the designer, may also be used consumed. This is achieved by abstractly specifying optional services.

Load-time. Ideally, devices should be able dynamically to discover the applications that are available in the environment, and applications should be able to adapt to a device's capabilities. This can be achieved if applications are defined in terms of their requirements, and devices are defined in terms of their capabilities. During load-time, some sort of mechanism tries to match these constraints.

One of these matching mechanisms is called dynamic discovery. Because the user tasks may be bound, or depend, on the physical surroundings, such tasks have to be enabled by contextual services. Therefore, the system must discover and consume the services available in the environment, in order to perform the desired tasks.

Furthermore, devices negotiate about requirements and capabilities. This enables a device to discover which applications and services can by hosted by its resources. Another concern is apportioning. By incorporating an efficient algorithm, and using resource information and application demands, there can be a split of the execution burden between the device and available servers.

Finally, during load-time, the user interface is appropriately adapted and composed, according to the resources available. The main goal is to be able to offer all the functions of an application on any device, by appropriately adapting the interface to the device's capabilities.

Run-time. At run-time, the currently active portal (i.e. device), or portal set, has to be constantly monitored. Any change of the resources must initiate an adaptation of the application to the new resources. Furthermore, changes initiated by the user must also be monitored.

Furthermore, users must be allowed to start new tasks without being interrupted. Therefore, any changes in the environment, handoffs between environments, or even network disconnections have to be catered for. In order to support a task-ori-

ented application, the user must have continuous access to the services requited to carry out the task.

Finally, unexpected failures, such as low batteries and service crashes, have to be handled appropriately during run-time. The current mechanisms used for failure detection and recovery have to be re-examined for their applicability to pervasive systems.

The Portolano Project

The Portolano project at the University of Washington [Esler et al., 1999] provides a view of the future where computers have been replaced by specialised devices. The analogy they present is that today's computers are essentially electronic Swiss Army knives: good for a camping trip, but impractical for activities requiring efficiency and quality. Computers are too complex because they try to be "all things to all people" (p. 256).

Esler et al. envision a future where all devices are task-specific, and ubiquitous. That means that their interfaces will not be visible to users. Furthermore, users will not have to deal with technical matters such as file formats, configurations or connectivity. Users will expect to accomplish their tasks easily and worry-free by using a market of information services to which they have subscribed.

Three key research areas have been identified in order to overcome the obstacles to realising this vision. These are user interfaces, distributed services, and infrastructure.

User Interfaces. The vision of the Portolano project requires a way of handling access to information services from more than one point. It is therefore essential that users can access the same information from different devices, and even provide a smooth transition when users switch devices. It is of great importance that users

can carry out their tasks with the same ease on any device. The interfaces should be usable on any device, and yet clearly represent the same information service.

A number of potential solutions to this problem are discussed. For example, the Interface Description Languages (IDLs), which describe an abstract user interface using a hierarchy of types, could be used [Hodes et al., 1997]. Additionally, a scheme built on top of XML which has superseded IDLs may also be used [Hodes & Katz, 1998]. Finally, Motorola's VoxML markup language is another interesting effort that could resolve a number of issues by integrating speech interfaces through simulated dialogues for interacting with web content.[3] Although Esler et al. seem to agree that markup languages are indeed an enabling technology, they point out that they will never be enough for describing the content and semantics of a document. The experience of numerous plug-ins and extensions to HTML are quite recent, and they suggest the need for a robust UI architecture to balance the needs of users and designers.

Infrastructure. A key feature of the supporting infrastructure is the discovery of resources. A number of research efforts have been directed towards this goal, such as RDP [Perkins & Harjono, 1996], SLP[4], JavaSpaces and Jini,[5] T-Spaces,[6] Universal Plug and Play,[7] and the consumer electronics consortium HAVi.[8]

A further key issue is the requirement for data-centric networks, proposed by Esler et al. This means that bundles of data should be able to use (and pay for) any resources needed, until they reach their destination. The network should be able to inform the devices about what network they are using, as well as the provision of

---

3. For more on the Motorola VoxML see http://www.motorola.com
4. For more information about SLP, see http://playground.sun.com/srvloc/
5. For details about JavaSpaces and Jini, see http://java.sun.com/
6. For more information on T-Spaces, see http://www.almaden.ibm.com/cs/TSpaces/
7. For more information on UPNP, see http://www.upnp.org/
8. For details about the HAVi specification, see http://www.havi.org/

admission control. Furthermore, the network infrastructure should offer quality of service (QoS) guarantees, for anyone willing to pay more. Related work in this area is the RSVP protocol [Zhang et al., 1993], and the QEX protocol [Davies et al., 1996]. Additionally, data-centric network should allow for the persistent storage of ubiquitous information. Therefore file systems like DFS[9], Coda,[10] or some combination is required. The ideal would be ubiquitous storage that is available to a number of distributed clients on any ad hoc network. This is in direct relation to our notion of spheres, described earlier, and how they can be maintained in a pervasive environment.

Finally, support for distributed computation is of great importance. Devices should be able to perform functions on their own, and these functions do not necessarily have to be provided by the device's designer. These functions should be able to be "downloaded" from the network, and executed whenever appropriate.

Cooltown

Cooltown (recently rebranded as HPBazaar)[11] is a project undertaken by Hewlett - Packard. It tries to utilise existing technology in order to create an infrastructure for supporting nomadic users. The goal is to bridge the physical world and the virtual world, by means of web technology.

Web Presence. The systematic integration of web services has lead to the notion of web presence. Web presence is defined as the representation of people, places, and things on the web. This set of categories is the same that was used in Taligent [Potel & Cotter, 1995].

Things obtain a web presence by installing small web servers on each device, and connecting the devices to the web. Places, which act as placeholders for things and

---

9. See http://www.opengroup.org/
10. See http://www.coda.cs.cmu.edu/
11. See http://www.hpbazaar.com/

people, provide a service called a "PlaceManager", which organises web things into collections. Peoples' web presence is represented by a global home page, which contains "WebLink" services to assist communication among individuals.

Infrastructure. The creators of CoolTown argue that existing web technology is enough for enabling people, places, and things to become web present. The web has widespread accessibility, and provides access for mobile users. Also, resources can be accessed from any device that supports the standard protocol (HTTP), and resources outside the current physical environment can be accessed transparently. Furthermore, the diversity of devices argues against a software or application based solution. Finally, the structure of the web does not require a complete operation network, but instead minimises the amount of infrastructure that needs to be operational in order to achieve basic communication.

CoolTown's infrastructure is divided into three layers, called bottom, middle and top. The bottom layer is responsible for obtaining points of web presence of people, places, and things. This essentially happens by sensing URLs. Three methods are used for achieving this: Discovery (e.g. a public "directory" on a network), direct sensing of a URL from a beacon, or indirect sensing, such as translating GPS coordinates into zip codes, and then into a place's web presence URL.

The middle layer provides services for exchanging content between the web presences of entities. Such interaction requires client software, which in this case is a web browser, and web servers, which are embedded into things. Using the standard HTTP protocol, these interactions can take place. It is assumed that devices will know the service name for the type of service they require, and it is hoped that standard names such as "printing" and "camera" will emerge from efforts such as the Universal Plug and Play Forum. Furthermore, a service provides a list of entry points, which are the functional units of the service. Another point to note is that

the servers which are embedded in things also support device discovery. This enables users simply to walk up to a public device and use it.

The top layer of CoolTown's infrastructure provides services related to places and nomadic users. A place is really a context for service provision, based on an underlying physical domain. Devices should be easy spontaneously to connect to networks, with little or no human intervention. This requires automatic service discovery. There is a problem, however, because real places don't always match the topology of a network. Furthermore, it is also a matter of policy what is included in a virtual place. All these issues are resolved with the use of a PlaceManager. The PlaceManager is a component responsible for providing secure views of the subset of resources present in the place, and the services that are provided. This content that is provided to a user by the PlaceManager is policy-driven: it depends on the user's security clearance, and the device's functional capabilities

Higher-level services for mobile work, such as remote access to things and people, or location-aware services, have been implemented. For secure remote access, CoolTown uses software called SecureWebTunnel, which employs secure web tunnelling for connecting remote devices. For remote access to people, WebLink has been used. A link from a person's home page to a globally accessible WebLink redirector service provides URLs of suitable communication services as the owner of the home page moves around. A similar effort is the mobile people architecture [Maniatis et al., 1999]. Finally, for location-aware applications, the WebBus has been used. It is an embedded web server with a GPS transponder and a wireless link to the Internet. According to its location, the WebBus server provides different services to its clients.

2.4   Presenting the information and interacting with it

Once the required infrastructure is in place and operational, information is able to travel through the network and reach the user. Information may be presented to humans in different ways, and each way has its own advantages and disadvantages. For instance, a graph may be better suited than a table of numbers in certain situations.

Furthermore, the interaction between the user and the information is of importance. There exists a relationship between the presentation of the information and the interaction between the user and the information. A certain presentation supports only certain types of interaction, and vice versa. In Chapter 7 we discuss our approach to interaction for pervasive systems and present an interaction technique we have developed. We believe that presentation and interaction ought to be seen as two different aspects of a system. In Chapter 7 we propose a way of separating the interaction from the physical form of the system.

The area of User Interfaces has seen recent research efforts resulting in two more categories of UIs, in addition to the traditional software UIs. These new categories can be seen as two sides of the same coin. On one hand, we have the Invisible UIs, which try completely to hide the interface, and make the user unaware of the interface's existence. On the other hand we have Tangible UIs, which try to present the user with a set of tangible objects that provide the interaction methods to the user.

This distinction of UIs is rather crude. In many cases certain types of interfaces use interaction techniques or presentation methods that can also be found in other types of interfaces as well. Yet, this distinction allows us to begin a general discussion, and consider various types of interfaces.

**Software
Interfaces**

The first types of interfaces to be used were the software interfaces. These interfaces were designed by the programmer and generated by the software, and usually displayed on a terminal screen. The early versions of such interfaces were Command Line Interfaces, such as Unix and DOS. The user was able to give input through a keyboard, while output was usually simple text restricted to the screen and the printer.

Later versions became much more graphical and so they were called GUIs, Graphical User Interfaces. Such interfaces involved more graphics in the presentation of information and status of the system, as well as the use of additional input devices like the mouse or a digital pen. These interfaces were based on the notion of direct manipulation [Shneiderman, 1983], metaphors and affordances.

Despite advances in technology since then, resulting in very graphical and interactive interfaces, the functionality of today's GUIs has not changed much compared to their early ancestors. The main drawback is that they are too static and present information in more or less the same way on every occasion, according to the way the programmer designed the interface.

In trying to address this limitation, a number of research efforts have focused on generating interfaces that are dynamically determined. Their main goal is to use contextual information, as well as some high-level guidelines provided by the designer, in order to present the user with an interface that is tailored and best suited to the situation. The area of context awareness, which we have already discussed in this chapter is of direct relation to this.

A number of discrete research areas have addressed the issue of dynamic interfaces. These include User Interface Management Systems [e.g. Kasik, 1982], systems based on state transition diagrams [e.g. Newman, 1968; Wasserman & Shewmake,

1982] and context-free grammars [Olsen & Dempsey, 1983]. Furthermore, a number of systems have implemented constraints, such as Sketchpad [Sutherland, 1963], ThingLab [Borning, 1981], and Garnet [Myers et al., 1990]. Finally automatic techniques have been utilised to generate dynamic interfaces. These include model-based tools such as Cousin [Hayes et al., 1985] and HP/Apollo's Open-Dialogue [Schulert et al., 1985], UIDE [Sukaviriya et al., 1993], Mike [Olsen, 1986], Humanoid [Szekely et al., 1993], and ITS [Wiecha et al., 1990]. These systems use heuristic rules to select components, layouts, and other details of the interface.

Invisible Interfaces

Invisible interfaces are a vision of the future for many researchers. Users will be able to operate and use devices without being aware of an interface. A simple example that illustrates this is the ability to switch on and off the lights in a room, by simply clapping our hands. In a more advanced scenario, we can consider the situation where a room is aware of the people in it, and responds to people's commands, in whatever form they may be: speech, gestures, writing, clapping, moving around, etc.

Effectively, the system provides a completely "natural" interface to the user. There is no need for keyboards, mice, etc. (although they could be useful in certain situations). Users will be able to communicate their wishes to the system, and the system will be able to interpret and perform the users' commands and wishes. By being so natural, such systems aim to become efficient, effective, and easy to learn.

There are two main steps that need to be taken in order to reach a stage where invisible interfaces could become a part of everyday life. The first step is to embed computers throughout the environment. With such an infrastructure devices can communicate with each other in order to co-ordinate their activities and produce the required results. This area of research is called ubiquitous computing, and more about various efforts in this area is described in the next part of this section.

Having embedded devices in the majority of everyday objects and devices, the next step is to allow users to communicate with the embedded devices in a natural and effortless manner, without having to be disturbed from their activities. To achieve this, new interaction techniques are needed, such as gesture recognition (which we describe in Chapter 7).

Pervasive and Ubiquitous Computing. The key idea underlying ubiquitous computing is that all kinds of devices will have computing power embedded in them [Weiser, 1991]. It has been said that the result of Moore's Law will not be supercomputers performing super-hard calculations very quickly, but that very small devices will be capable of performing everyday tasks with adequate speed [Miller, 2001].

Currently, the cost of embedding wireless computational power in an everyday object, such as food packaging is more than the price of the object. When this balance shifts we will see a huge increase in the number of ubiquitous devices available for households and everyday use. The problem at that point will be that of co-ordinating all these devices in order to generate something useful.

Users will most likely distribute their computational needs over a spectrum of devices, ranging from small, light, mobile devices to large, wall-sized screens. Therefore, it becomes vital that each of these devices is not treated as a single computation and storage entity, but rather as part of a larger, integrated, multi-user environment. This problem is largely addressed by the efforts at distributing information, described in the previous section.

Furthermore, users will need to communicate with each other, using this broad spectrum of diverse devices. Current systems, like Microsoft NetMeeting or GroupKit [Roseman & Greenberg, 1996], are too simplistic. Instead, we have to

understand how people with multiple different devices, in different places, and with different network connections, may communicate effectively and useably.

A number of efforts are trying to address such problems, by providing a research testbed for ubiquitous devices. By designing and building a so-called "Smart House", these projects study the developing user interface paradigms and communication techniques for large numbers of "intelligent" devices and objects. Usually, such projects provide an actual house, which has a number of intelligent devices, ranging from television sets to refrigerators. Furthermore, they also provide the technology for allowing communication and co-ordination between such devices. Examples include the Living Room Project [Vanhala, 2001], the inHaus project [Miller, 2001], and various other similar projects in USA, Japan, Holland, Sweden and Switzerland.

Interaction techniques. Providing the enabling technology and the required infrastructure is the first step in realising the vision of a ubiquitous computing future. What we are currently lacking is a way for humans to interact with all these smart, embedded devices. The current interaction techniques have been optimised over the last 20 years on systems that typically provide a screen, keyboard and mouse. Ubiquitous devices require new ways of interaction.

The diversity of devices, and the potential problems caused by it, may be seen even today. A typical example is that of the display screens. Today's devices employ screens that range from 5cm in diagonal, such as mobile phones, PDAs, etc., to screen sizes of 4.5 meters in diagonal, such the DynaWall in the i-Land project [Streitz et al., 1999], plasma screens, video walls, etc. This diversity implies that different devices need different interaction techniques, and that today's techniques are not adequate. Pier and Landay [1992] describe an example of potential problems. When using standard pull-down menus with large screens, they found that

some users were simply too short to reach the controls on the screen. The potential interaction problems on newly emerging displays such as 3D volumetric displays[12] will probably pose even more problems.

A number of different branches of interaction techniques have been developed. Speech recognition, which was one of the first alternative interaction techniques to be explored, has made a lot of progress. Currently, there are end-user software packages available (IBM ViaVoice, MS Office XP), which allow users to dictate text, or issue simple verbal commands to the interface. Although their capabilities are limited, they still offer a promising future.

Gestural interfaces and interaction techniques are also beginning to emerge. Such interfaces are based mainly on hand gestures made by the user, although other parts of the human body can be used as well. Users issue commands simply by performing a gesture. Many ways have been developed for recognising gestures, mainly optical recognition using cameras [Rui et al., 2001], and electronic or mechanical actuators attached to the human body [Rekimoto, 2001]. Although no real breakthroughs have been made yet, this area is quite interesting, and could offer very good results, especially when combined with speech recognition.

Various other interaction techniques are currently being researched, like eye tracking [Yamato et al., 2000], and biofeedback.[13] However, most of these efforts are based on existing interfaces, and they try to improve the interaction with existing interfaces. Unless the focus shifts from existing interfaces, not much progress will be made in. The main reason is that current systems are based on the event-based model, which assumes that interaction occurs at distinct points in time (events). This cannot hold for recognition-based systems, since they need to process continu-

---

12. See http://www.3dtl.com/
13. See http://www.biocontrol.com/

ous input. From this fundamental difference, a lot of differences in the interaction techniques are bound to follow.

**Tangible
Interfaces**

Tangible interfaces can be defined as those interfaces that involve physical objects as part of the interface. Most electric devices, such as ovens and cd players, provide physical, tangible interfaces to their users. However, the main research focus is not the design and layout of buttons for devices, but to provide a "real" interface to the digital world.

Throughout our lives, humans develop skills for manipulating objects in our world - the real world. These skills that we develop are not fully utilised by conventional interfaces. In other words, our knowledge and experience of real world objects and their behaviour, is in large part wasted.

To design a physical interface, one could follow design guidelines that have been developed especially for such interfaces. Alternatively, a good source of inspiration for tangible interfaces is metaphors.

Design Guidelines. It has been suggested that guidelines that are derived from object-oriented design of computer software could be applicable to the design of tangible interfaces. More specifically, Pendersen et al. [2000] have proposed that by following the OO paradigm, we could derive a parallel between conventional OO and what they call "tangible object-oriented abstraction".

They proceed by providing a set of four design guidelines for the design of tangible interfaces:

• "The physical instantiation given to central digital objects in the user's task domain should be objects that are appropriate and meaningful in the use situation".

- "All methods of relevance to the user of the digital object should be available as manipulations of the tangible representation".

- "The manipulation of the physical object should be simple, appropriate to that particular object and to the use situation".

- "A user-relevant method which is inherited through an object hierarchy should be reflected in similar physical triggers and manipulations".

The main goal of these guidelines is to follow the same process as conventional object-oriented design. This implies that the three characteristic dimensions of abstraction, encapsulation and inheritance have to be adopted for tangible interfaces. A useful adaptation, for instance, would be that the appearance of a physical object might provide information about relevant attributes.

General Themes. Instead of pursuing design guidelines, other research efforts have tried to produce general themes that may be applied to many situations of designing tangible interfaces.

For instance, the Tangible Bits project [Ishii & Ullmer, 1997] adopts the convention that there should be a seamless interface between people, bits, and atoms. They have demonstrated this by designing the metaDesk, transBOARD, and ambientROOM prototypes. Their key characteristics are interactive surfaces, the coupling of bits and atoms (i.e. information bound to physical objects), and the use of ambient media. Their central approach is the use of ambient media, such as background sounds (falling rain), and background light (ambient light) for representing "bits", i.e. information.

Other projects have used metaphors from everyday life. For instance, the Pick-and-Drop direct manipulation technique [Rekimoto, 1997] was inspired by the use of chopsticks to eat food in Eastern countries. The DataTiles project [Rekimoto et al.,

2001] created tagged transparent modular tiles as representations for physical and virtual objects, locations, and information.

Another characteristic example of the use of metaphors is Durrel Bishop's Marble Telephone Answering Machine [Crampton, 1995]. In this design, a small marble that pops out of the machine represents an incoming voice message. The users can manipulate messages by manipulating these marbles. For instance, if the user places a marble on the speaker, then the corresponding message is played. A message is deleted simply by dropping the corresponding marble into a "black hole" in the answering machine, so that the marble can be reused.

However, tangible user interfaces are still in their infancy both in terms of technology and of our understanding of their properties. Svanaes and Verplank [2000] propose three steps for exploring the design space of this new technology:

- "Finding the dimensions and elements of the technology. This can be done as a formal exercise, but should be grounded in empirical studies of how users structure their experiences with the technology."
- "Building simple demonstrators illustrating the dimensions and elements."
- "Finding metaphors that fit the new class of applications. This search can be formal, inspired by usability tests, or inspired by other media and cultural phenomena."

There is still a long way to go until tangible interfaces mature enough and are widely used for everyday tasks. However, the results that have been obtained so far are quite promising.

2.5   Interdisciplinary approaches

A number of interdisciplinary research efforts are indirectly related to the field of pervasive computing. These efforts can be roughly grouped in two categories:

efforts to understand and describe the use and adoption of technology in various settings, and efforts to create virtual environments (virtual reality, augmented reality) based on interdisciplinary research approaches, ideas and frameworks.

Use and acceptance of technology

Understanding the use and acceptance of technology in a public setting is a useful step towards successfully designing the large-scale pervasive systems that we envision. Several factors are known to influence the acceptance of technology [e.g. Davis, 1993], whole schools of though have been created to explain how and why technology becomes accepted or not [Green, 2002]. More specifically, environmental factors such as level of privacy and social density are known to influence behaviour of people using technology in public [Kaya & Erkip, 1999; Gifford, 1987]. Different types of territory exercise different levels of control [Ruback et al., 1989]. So for instance, people generally have permanent control of their home, in contrast to public territory in which many people can gain access and temporary control. Furthermore, it has been found that users of technology in public spaces (such as ATMs) may change their behavioural intention because of physical intrusion into their space [Little, 2003]. In this study, users stated that their initial intention was to perform two or more ATM transactions but because people who were waiting in the queue came too close or walked up behind them, thus causing intimidation and pressure, the users actually performed only one transaction.

Building on issues of territory ownership, Silverston & Hirsch [1992] proposed that media in the home posed a whole set of control problems, such as regulation and control of space. In a similar study [Baillie, 2002] it was shown that feelings of control (or lack of it) over spaces in the home were an important indicator of the participants' feelings towards the technologies that those spaces contained.

A more refined approach to understanding the use of technology has taken the form of a framework [Venkatesh, 1996]. This framework addresses technology

assimilation in the home space, and proposes two categories of space: the social space which constitutes the social structure and activities performed in the household, and the technological space which represents the nature of technology in the household. The typical example given is the use of the computer, which was viewed as a work tool in the 1980s, but became assimilated into the social space with the advent of email and the Internet in the 1990s. This framework was used in a recent study [Baillie, 2002] in order to understand spaces in the home. The study found that the framework was limited and did not provide a rich enough picture for all the types of spaces encountered in the home.

Although the understanding of the use and acceptance of technology in both public and private spaces has been explored, very few results have been presented in relation to how we can design computer systems, and specifically pervasive systems. This does not necessarily mean that the understanding we have gained is not important. Rather, we believe that we need to operationalise this understanding if we are to produce useful results in relation to design.

We now turn to virtual environments where, in contrast to what we have discussed so far, a lot of ideas and approaches have been implemented and embodied in systems that have been built.

Virtual environments and models of space

A lot of work from architecture and urban design has been applied to the design of virtual environments. Rather than duplicate here our survey of architectural and urban design ideas, we have placed it in our study of space in Chapter 5, where it is more appropriate. However, here we mention some work that has built on the same body of knowledge.

The PERSONA project [McCall, 2003] explored the concept of navigation within information spaces. Part of this project involved developing concepts which

designers or evaluators of electronic virtual environments could use to explore navigational issues. A recent study [McCall, 2003] explored whether people found design cues from the built environment useful within virtual environments. It was found that paths, signs and districts received strong positive responses. It was also found that landmarks, like statues, need to contain some semantic relevance in order to receive positive responses.

Research focusing on building virtual environments consists of an analysis of the concept of place as proposed within geographical theory, and its relevance for the design of spatially oriented technologies (such as tourist guides) [Brown & Perry, 2002]. Further research on the design of interactive physical environments has adopted a geographical perspective on place focused on its "experiential quality". Additionally, the design of virtual models of real places has been informed by "phenomenological" studies of place [Turner & Turner, 2003].

Finally, a lot of effort has focused on understanding the way in which multiple physical and virtual spaces co-exist. For instance, Dix [2003] suggests the existence of three types of space: real space (actual objects in actual physical space), measured space (the representation of real space in the computer) and virtual space (electronic spaces created to be portrayed to users, but not representing explicitly the real world). Of direct relevance to such work are models and taxonomies of spatial context [Dix et al., 2000], models for mixed reality boundaries [Koleva et al., 1999] and the capturing of human spatial understanding which can be collected from various sources [Dix, 2000]. Of relevance also is work from the Artificial Intelligence and Geographic Information Systems communities focusing on informal reasoning about space [Grigni et al., 1995; Papadias & Egenhofer 1997].

Many projects have developed an understanding of space and have applied this understanding to the design of virtual environments. However, this does not advance the design of pervasive systems for a number of reasons.

Design approaches and models

A number of approaches have been developed for designing traditional computer systems. For instance, USTM/CUSTOM [MacCaulay et al., 1990] uses diagrams of task models along with textual descriptions to establish stakeholder requirements. Another approach called ETHICS [Mumford, 1993] addresses the technical and social requirements of a system. Soft systems methodology [Checkland, 1983] assists designers in understanding the situation for which they are designing. This is achieved by developing a detailed description of the problem situation, generating "root definitions" for the system and defining how they are fulfilled, identifying discrepancies between the actual system and the conceptual model of the system and deciding which changes are necessary to the system as a whole. Finally, a number of cognitive models have been developed for understanding the users' mental processing. Such work includes GOMS [Card et al., 1983] and CCT [Kieras & Polson, 1985], ICS [Barnad, 1985] and TKS [Johnson & Johnson, 1991].

2.6   So, why can't we design pervasive systems?

As we mentioned in the previous chapter, there is an absence of theory within the pervasive systems community. In this chapter, we saw some design approaches that have been used elsewhere within HCI and Computer Science, with some success. Paradigms from object-oriented software design have been applied to the design of pervasive systems (or at least their interfaces). However, such an application of OO methodology has not been justified by Pendersen et al. [2000], and is questionable at best. Some researchers[14] have adopted a less rigorous approach by calling upon

---

14. The researchers best known for this approach are Hiroshi Ishii and Hun Rekimoto.

artistic value and aesthetics. Such projects attempt to design pervasive systems based on aesthetics and artistic values related to the purpose of the design. The aesthetic marrying of the real world and the digital world is usually seen as their ultimate goal. This approach can produce very pleasing results, but the main drawback is that such designs cannot be reproduced for different tasks or needs, cannot be easily extended, do not follow a specified set of rules, and do not address the issues of having a disjoint research area that we have discussed.

We have also seen a number of research efforts that deal with user interaction. This is an essential part of pervasive computing, since invisible interfaces, tangible interfaces, and to some extent context awareness are all related to how the user interacts with the system. However, such research addresses only a small part of the design of pervasive computing, and unfortunately most interaction methods cannot be completely independent of the system with which they are used. Therefore, deciding to use a new interaction technique with a pervasive system would impose constraints on the system as well as the interaction. Looking just at the interaction is not sufficient, unless it can be made completely independent of the underlying system.

We briefly mentioned in the previous section a number of models for understanding space and designing virtual environments. These are not sufficient for building pervasive systems for a number of reasons. First, these models were not developed with the aim of advancing our understanding of pervasive systems design. We saw an attempt at understanding the relationship between real and virtual spaces, but this has not yet produced the kind of design guidelines helpful for the design of pervasive systems. Furthermore, we saw that there exists an understanding of physical space, based on which virtual spaces have been designed. The nature of pervasive systems, however, is different to that of virtual environments and virtual reality; pervasive systems will be part of everyday life and objects, while VR systems offer

an alternative world. We do not yet have a thorough understanding of how pervasive systems impact the build environment, and vice versa. Furthermore, such models offer no clear understanding of how this new kind of computing system should be interweaved with our physical environment.

Lastly, the design methods we presented in the previous section have been developed with traditional systems in mind. Therefore, they consider systems that support predefined tasks and in many cases assume a job or office environment, i.e. issues that are different with pervasive systems. Furthermore, generic approaches like the soft systems methodology rely on the knowledge of the designer in order to recognise potential problems and mostly offer a generic approach to design. This can be useful at an initial stage, but for the design of pervasive systems we need an approach that is better suited for the issues involved. Also, such methods tend to focus on details of systems, something which is possible with traditional, static systems, but which cannot always happen for the dynamic and rich environments supported by pervasive systems. Additionally, as we describe in Chapter 3, the cognitive models we described earlier lose their usefulness when we design on a large scale and for the unanticipated situations of everyday life.

Back to square one?

At the beginning of the chapter we stated one of the fundamental tenets of HCI: *Design for a specific user, performing a specific task, in a specific domain*. Perhaps this can be of help to us in the design of pervasive systems. However, in light of existing practice, reflected in the literature review we have just seen, the terms user, task and domain have so far been interpreted in a way that we believe is inappropriate for pervasive systems.

The reason behind this is simple: the pervasive systems community has had to borrow heavily from conventional systems design. In such conventional designs, the *user* is studied at a cognitive level, and sometimes at a group cognition level. In

our discussion of user involvement, we saw guidelines as to how the user can be made a useful part of the system by providing input to the system at the right time. In general, traditional systems do not require a more refined "model" of a user, a model where the user is not seen as an individual or group member, but as a person who has an everyday life.

Similarly, traditional systems had only to worry about the *task* which they are supposed to support. Many techniques have been developed to manipulate the interface according to the task currently supported by the interface. However, no need existed to be aware of other desktop computers within, say, a room. Therefore, problems arise when, for instance, a user's task is complicated in terms of computer devices, such as "share my photographs". In this case, the user has to use various devices and applications, all of which live in their own little world where the user's task is "copy x y z files from memory stick to hard drive", "compress x y z files", and "email file x to address y". The user's point of view cannot be seen by any of the systems or devices that are being used.

Finally, in traditional systems (but also in most current pervasive systems) *domain* has been poorly substituted by "location", which greatly hampers the potential for analysis and evaluation of a system. Physical locations are much more than GPS coordinates. They have values and rules attached to them. They have histories. All of this valuable information is wasted if we only utilise the GPS coordinates of a location.

2.7   Arriving at our framework

The established interpretations of the terms user, task and domain in HCI are inadequate for pervasive computing. Here we explain why these notions do not adequately address the requirements for designing pervasive systems. As a way to

motivate a shift in focus on the issues we believe to be of importance in the design of pervasive systems, we introduce the terms citizen, sphere and space.

From users to citizens

In this chapter we have provided a review of research that is related to the design of pervasive systems. We now draw on our review to identify a number of characteristics which we can attribute to the existing focus of user.

Users tend to improvise [Suchman, 1987] and this needs to be taken into account both during the design and use of the system. Users also have intentions which can partly be modelled using human intervention and direct input from the users [Bellotti & Edwards, 2001]. Another characteristic of users is attention, which is limited. It can be disturbing for users if the devices and tools they use are competing with the activity itself for user attention [Ackerman, 2000; Ackerman et al., 2001].

An important characteristic of users is that they have knowledge of the physical world, and users apply this knowledge when using computer systems. Tangible interfaces build on users' ability to manipulate physical objects [Ishii & Ulmer, 1997] and understand or use metaphors from everyday life [Rekimoto, 1997]. Traditional computer systems have tapped into users' abilities directly to manipulate representations of objects (direct manipulation) [Shneiderman, 1983] and the affordances that they recognise in these represented objects [Dix et al., 1998].

Several theoretical approaches model users' cognitive processing, memory, senses and their associated limitations. Such work includes GOMS [Card et al., 1983] and CCT [Kieras & Polson, 1985], ICS [Barnad, 1985] and TKS [Johnson & Hyde, 2003].

The above work focusing on user characteristics addresses the design of traditional systems. However, large-scale pervasive systems are quite different from conventional computer systems. Such large-scale pervasive systems offer ubiquitous access

to information used in everyday life, in contrast to the tightly controlled and specified use of traditional computer systems. Inevitably, the extensive reach of such systems makes them interact with the norms and regulations that govern everyday life. Because of this, we argue that the design of pervasive systems needs to be informed of such issues.

Effectively to design for pervasive access to information, the first issue we need to consider is that by their nature, large-scale public pervasive systems are used by people that the designers could not study. In other words, we can't know the individual users of public pervasive systems. Indeed, in certain situations privacy requirements and restrictions may prohibit us from knowing anything about the users and their everyday lives.

To design for truly pervasive access to information we also need to consider that people are members of communities [Held, 1989] and study people's membership at the levels of group, organisation, community and society [Vicente, 2004]. Furthermore, we need to acknowledge the presence of norms and regulations in everyday life [Benderson et al., 2003] in order effectively to support everyday tasks. Many technological projects have failed because they were out of touch with their contemporary social settings and norms [Schuler, 2001]. In addition to these norms, people live their everyday lives according to rights and obligations that exist in our society [Marshall, 1950; Lister, 1990; Turner, 1994], and these need to be taken into account.

Furthermore, we need to consider that there already exist large-scale systems used in everyday life: public services. In providing large-scale pervasive systems as a way of improving everyday life, we need to consider how people regard existing large-scale public services [Krajewski, 2001], what people expect of public services [Stoker & Williams, 2001], and what rules and regulations are imposed on public services

[Harrison & Woods, 2001]. Additionally, functional characteristics of public serv-ices [Patterson, 1999] can be observed and studied for their applicability in public pervasive systems.

A comparison between the issues being studied under the focus of *users* and the issues related to pervasive systems reveals a significant mismatch. From our analysis we deduce that the user characteristics being studied do not match adequately the requirements for designing for pervasive access to information. To address these requirements we introduce a new term, *citizen*, in an attempt to motivate a refocus-ing on those aspects of people and society which are becoming more relevant with the advent of pervasive computing.

The concept of a *citizen* can be more meaningful in the realm of pervasive systems than the concept of a computer user. We may know little (besides the physical and psychological metrics) about the particular *users* of a publicly available, large-scale pervasive system. But there are a number of things we can know about *citizens*. Such information may include citizenship rights, how citizens view public systems, (e.g. broadcast television, public transport etc), and what types of access to public systems citizens prefer or require. All these have an impact in the design of perva-sive systems, as we show in Chapter 3.

From tasks to spheres

The established focus in HCI on task has resulted in models and representations of the tasks people carry out and methods to manipulate these representations. There exists a number of approaches to tasks analysis, each with an emphasis on a slightly different area. Task analysis tends to describe the observable behaviour of users, but is also used to build conceptual models of how the user views the system and task.

Specifically, hierarchical task analysis (HTA) [Annett & Duncan, 1967; Shepherd, 1989] performs a decomposition of the tasks and subtasks involved in a specific activity. The hierarchy of tasks and plans generated by HTA describe the order in which tasks are performed and the conditions under which this happens. Knowledge-based task analysis offers an understanding of the knowledge needed to perform a task by listing all the objects and actions involved in performing a specific task [Diaper, 1989]. Additionally, task analysis techniques such as task knowledge structures (TKS) [Johnson & Johnson, 1991] have been used to model collaborative tasks [Johnson & Hyde, 2003; Zachary et al., 2000], in an attempt to offer insights at modelling task and collaboration.

The models and representations provided by task analysis can be used in the design stage as well as during the system's operation. For example, USTM/CUSTOM [MacCaulay et al., 1990] uses diagrams of task models to establish requirements of a system. Additionally, a number of architectures and methods are used to manipulate and adapt a system's interface to suit the task being carried out. In systems such as Cousin [Hayes et al., 1985], Open Dialogue [Schulert et al., 1985] and UIDE [Sukaviriya et al., 1993] the characteristics of the specific task being carried out influence dynamically the appearance of the interface.

A number of tangible user interfaces have been developed to match the precise requirements and needs of a specific task. For instance, the marble answering machine [Crampton, 1995] provides a tangible interface which is suited specifically for the task of listening to messages on the answering machine. In general, tangible interfaces offer the advantage of tangible manipulable artefacts well suited for the task they support at the cost of not being able to dynamically change the interface they provide.

The traditional approach of focusing on specific tasks has helped improve the systems we design and deliver by addressing specific situations and tasks. However, fully to address the requirements of designing for pervasive information access, there are a number of additional issues we need to consider.

The specific tasks which are carried out within a truly pervasive cannot necessarily be known in advanced. Just as we cannot study the specific users of a large-scale pervasive system, we cannot know in advance how people will use the information we deliver to them. For example, one's address book may be used to send invitation cards to friends, for the delivery of a product to a specific address, to call or send an email to a contact, and so on. The multitude of devices and locations embraced by pervasive systems simply renders inadequate the study of specific tasks.

To design for pervasive access to information, we need to abstract away from specific tasks and consider information itself. We need to provide ways of conceptualising the information we access [Kindberg et al., 2002]. Considering the research visions of invisible interfaces, we also need to address and conceptualise ownership and control over information [Hong & Landay, 2004]. Additionally, our systems need to be able to describe, to some extent, what happens to the information we input or "give" to the pervasive system [Adams & Sasse, 1999b].

Another set of important issues we need to address is the interaction between the physical location, the technology used to deliver information, and the information itself. In providing pervasive access to information, the same information gets delivered to various locations by means of various technologies. Therefore, we need to be able to reason about information in relation to the devices and locations where it gets delivered.

We also need to consider that our wide-ranging systems are intended to support people in their everyday lives. Therefore, in designing for pervasive access to information we need an understanding of the use and exchange of information on a public scale [Malina, 1999], along with the social issues this raises [Habermas, 1962]. People's deliberation about common affairs, as well as the generation, circulation and reconstruction of information in this process [Fraser, 1995] are now becoming relevant in the design for pervasive access to information.

The tools and models developed for understanding *task* do not adequately match the issues involved in designing for truly pervasive systems. The traditional focus on analysing specific tasks in specific environments does not adequately support our need for describing the presence and ownership of information in a pervasive system, nor how different locations or technologies have an impact on information. Therefore, we introduce the term of *spheres* as a way to motivate a refocusing on those characteristics of task and information that are most appropriate for pervasive access to information.

Using spheres, which can be thought of as pools of information, in Chapter 4 we describe the dynamics between three distinct types of information: public, social and private. These constructs offer us a convenient top-down abstraction away from task, and a way of conceptualising information and ownership in a pervasive environment. Spheres assist us in examining the delivery of different types of information in varying locations by means of varying technologies. At the same time, these notions allow us to address the social issues that emerge as a result of the pervasiveness of our systems.

From domain to space

In HCI, the meaning attributed to domain has traditionally been that of an application domain. A domain describes an area of expertise and knowledge in some real-world activity [Dix et al., 1998]. Additionally, a domain consists of certain

important concepts which the users manipulate in order to carry out tasks. More generally, however, the domain has been seen as the situation and environment in which a system will be deployed. This view has been adopted by systems methodology, such as soft systems methodology [Checkland, 1983].

In more recent years, the general environment in which a system has been deployed has been termed context. One of the many approaches to context suggests that any information that can be used to characterise the situation forms part of the context [Dey et al., 2001]. Numerous application have been developed to make use of context, ranging from automated call-forwarding [Want et al., 1992] to conference assistants [Dey et al., 2001] and complete systems [e.g. Banavar et al., 2000]. A recent survey of numerous approaches to context [Kaenampornpan & O'Neill, 2004] reveals that physical location (GPS coordinates), physical parameters and technical parameters are the most used aspects of context. Work has also attempted to extend such parameters over time [e.g. Schmidt et al., 1999].

To design for pervasive access to information there are a number of additional issues that need to be considered. Recall that pervasive access implies that people may be located in varying situations and locations when accessing information. Therefore examining specific locations now becomes inadequate.

Although every specific location is unique, there exists work which abstracts away certain characteristics of locations, and groups all possible locations into specific categories [e.g. Hall, 1969; Green, 2001]. These categories help in reasoning abstractly about locations, thus overcoming the problems of having to study infinite specific locations. In addition, we need to recall that pervasive systems will be used within people's everyday life, in open or public locations. We therefore need to consider the issues surrounding information access and technology use in the public realm [Biddulph, 1993; Buchecker, 2003].

Furthermore, we can now make use of abstract concepts that help us reason about the technology used to deliver information [O'Neill et al., 1999]. Building on an abstract description and classification of spaces, we need to consider how the delivering technology is affected by space itself.

In the previous section, we described public services as existing large-scale systems (from the viewpoint of citizens). Adopting a spatial approach, we claim that Architecture stands as a large-scale system, pervading our lives in urban environments. A large-scale pervasive system will inevitably interact with architecture, as it does with social norms and regulations.

Understanding architecture, therefore, plays a key role in designing for pervasive access to information. We need to consider how architecture's manipulation of space, aiming to minimise obstacles [Bentley, 1985], interacts with delivery of information in spaces. Similarly, because different spaces interact with each other [Logie, 1954] it would be appropriate to consider how these spaces interact with our delivery mechanisms. Ultimately, architecture can be seen as a large-scale system, and as such, we can learn from its presence and design. Numerous guidelines have been suggested for designing public spaces, such as the project for public spaces [2000] and the UK Commission for Architecture and the Built Environment (CABE). The issues raised by such efforts can be considered in our designs for pervasive access to information.

The traditional characteristics attributed to *domain* support the design for specific situations and environments. Comparing this with the requirements for pervasive systems, we observe a substantial mismatch. We introduce the notion of *spaces* as a way to motivate a refocusing on those aspects of location and environment which are most relevant and important to the design of pervasive systems.

In Chapter 5 we describe our notion of spaces. We differentiate between architectural spaces, which are physical locations, and interaction spaces which are conceptual volumes of spaces within which an artefact is usable. In delivering information to users, the designers need to consider which means are appropriate. In our work we focus on providing the appropriate means to deliver information. These means are what we call interaction spaces, which are instantiated in physical space.

Over the next three chapters we develop and discuss our concepts of citizens, spheres and spaces. Then, in Chapter 6 we discuss how these three concepts of our framework can be used together to form a design tool to help in design for pervasive access to information. Then we show how our ideas may be used to generate overall design recommendations and objectives (Chapter 9), to derive design alternatives which can then be evaluated and explored (Chapter 8), and to support the design of appropriate interactions with the system (Chapter 7). We do so by applying our framework and design tool at different levels of resolution (design recommendations, design exploration, interaction design) as well as for the purposes of design (Chapter 9) and evaluation (Chapter 8).

More specifically, in Chapter 7 we explain how our framework can affect systems design at an interaction level by instantiating interaction spaces that are appropriate to the interaction taking place. Our framework has been applied in the development of an interaction technique which takes into account our ideas about interaction spaces, and how they relate to privacy issues. By interacting with a system, a user gives away information about the interaction itself and the information being accessed. Here we develop an interaction technique that allows users to control how much information is withheld or not about the interaction itself by allowing for the generation of appropriate interaction spaces.

In Chapter 8 we use our design tool to evaluate design solutions and possibilities. Here we use our design tool to describe the information that exists in a hospital A&E department, as well as the means by which it is disseminated. We then explain in terms of our design tool some problems that the hospital faced, and also identify further issues that our design tool has raised. We then go on to propose design solutions based on the recommendations of our ethnographic work. In the process of doing so we discuss how our systems level applications and solutions may be brought in and used.

To complete our work, we show how we can generate design recommendations using our framework. In Chapter 9 we make different use of our framework, and describe a case study of a large-scale location, the city of Bath. In this case study we analyse various locations within the city of Bath, and then derive design recommendations and predictions. The recommendations we provide in this chapter may be followed through using the process we present in Chapter 8.

USERS OR CITIZENS?

3.1 The unanticipated effects of technology

On Friday 18 July 2003, British Airways ticket and baggage handlers staged a walkout[15] in Terminals One and Four of Heathrow airport in London. The walkout resulted in more than 500 flights being cancelled. About 10,500 travellers were stranded, while more than 100,000 people faced disruption as they turned up for grounded and rescheduled flights. Hundreds of passengers, unable to find hotel accommodation, spent Friday night in the terminals, with many complaining that the airline had not done enough to keep them informed. Massive queues built up Saturday in and around Terminals One and Four as passengers showed up for their flights only to learn they had been cancelled. The walkout ended on Saturday, but as a result thousands of passengers endured days of chaos while the backlog was

---

15. http://news.bbc.co.uk/1/hi/business/3108853.stm

cleared. Shares in British Airways fell as much as 6.75 pence, or 3.9 percent, and the loss to profits was estimated to be 8 million pounds.

Why did the BA staff stage this walkout? They were protesting against a new swipe card entry system, called Automated Time Recording (ATR), which allows managers to monitor the employees' working hours. It is understood that staff were worried that the system, set to be introduced the following Tuesday, could have lead to staff being sent home during quiet periods. BA denied this, and said swipe cards had already been in use in some parts of its Heathrow operations for three years and were widespread across British industry. But one union official said BA managers had been warned that staff were deeply unhappy about the system. He said "These are not militant workers - they have just had enough". A BA spokeswoman said the airline was looking at moving all of its operations from paper-based to electronic systems over the next five years, and that ATR had been identified as a "suitable" way of managing staff.

Technology is not always the problem

Swipe card systems have been in use for a number of years, yet they have not been the cause of such problems elsewhere. Why did the designers of ATR fail to foresee such a response from the users of a simple swipe card system? Why did they fail to foresee that such a system could cause so much controversy, problems and damage? It could be argued that such events are rare, spontaneous, and no one can really study users and design a system that takes into account extreme cases (such as strikes). However, we believe that this line of thought is fundamentally flawed. The problem is that the notion of a *user* tends to view those who use the system as isolated individuals living in an idealised world (as far as the system is concerned). This view causes us to focus on the psychological aspects of people. As such, we could have never predicted that the British Airways staff would stage a walkout.

People have many aspects and dimensions to them, and psychological aspects are just one of such dimensions. Traditional systems design and approaches, such as the ones we described in Chapter 2, study the psychological aspects of people in relations to specific tasks. Furthermore, traditional systems supported tightly defined environments and therefore the design approaches considered those social aspects that were of direct relation. Now, additional dimensions, such as social life and creativity, are starting to become relevant, as technology becomes more widespread and embedded in our lives.

The changing scope of technology

The use of interactive technology is no longer limited to static and highly constrained situations, places and people. Pervasive computing is at the forefront of a new wave of technologies aiming to be used in constantly changing situations and for constantly changing purposes. The models and understandings of users and tasks which have been used to design software and systems for static users and for specific tasks, are now being stretched and in many cases misapplied. Technology is now being used by groups of people, in dynamic situations and for a wide variety of tasks. Looking at the level of a single person, something that psychology can do so well, is no longer sufficient [Vicente, 2004].

Recent research has been directed at understanding the social aspects surrounding technology, and how our social structures, needs, assumptions and restrictions affect and are affected by technology. Fully to understand the impact of pervasive technology we need to look beyond the individual at the group, organisational, and ultimately political levels. In this context, the notion of a user is limited in usefulness as well as applicability. Our understanding of the cognitive processes within the brain of a user is useful for designing interaction at the interface level. In other words, designing for direct interaction between one user and "the system". In addition to this, however, we need to inform our designs with the design requirements

we have produced by studying social needs at varying levels. People do not operate technology in isolation, but rather with others and in the presence of others. People do not exist in a constrained world where their purpose is to satisfy the needs and requirements of the computer system. Rather, people exist in the real world. They have real needs. They have real problems. They are not inclined to carry out inconvenient or impractical tasks, they cannot be expected to behave as cold machines. And they live in a world which provides them with a balance of laws, obligations and fundamental rights. People have rights, but when reduced to users, they lose those rights. How can this understanding improve our designs?

If not users, then who?

The intended users of information systems in public settings may be viewed as "the public" of the societies in which these systems will be embedded. Designing a highly usable system is difficult, even when we can specify the users, tasks and domains that the system will support. The problem becomes even more difficult if the specific users are not known in advance, such as when we are designing a system for the public. Designing a system without studying its users is anathema to many in the HCI community, yet many systems are used by the public, and thus were designed without knowing in advance the specific users. Such systems include the subway, trains, buses, electricity, highway management (signs, lights etc.), telephone and television, and any system that supports or imposes social rules and regulations, such as road signs and voting [Bederson et al., 2003]. These systems were designed (admittedly sometimes badly) without having specific users in mind, but by targeting citizens. Yet, it is widely accepted that such public provisions do improve the quality of life. So if we are to follow established HCI design wisdom in designing pervasive systems for the public, then we need to extend the golden rule of designing usable systems, "Know thy user", to "*Know thy citizen*".

3.2   Know thy citizen

The concept of a citizen can be more meaningful in the social realm than the concept of a computer user. We may know little (besides the physical and psychological metrics) about the particular *users* of a publicly available, large-scale pervasive system. Indeed, concerns about privacy and similar issues suggest that many users of such a system would prefer that little or nothing was known about them. But there are a number of things we can know about *citizens*. Such information may include citizenship rights, how citizens view public systems, (e.g. broadcast television, public transport etc), and what types of access to public systems citizens prefer or require.

The concept of citizenship has a long and chequered history, although throughout that history it has retained certain common attributes. For example, citizenship is often conceptualised as comprising certain rights against and obligations towards a community, and implies membership of that community [Held, 1989]. However, it has often been used to exclude, as well as to include, and there have been numerous attempts to restrict citizenship to certain categories of people, thereby excluding other categories. Such restrictions have tended to favour more powerful groups within a society, and to disadvantage the less powerful. In the UK, for instance, citizenship rights have been withheld at various times on the basis of factors such as gender, race, age and social class.

Marshall's view   Marshall's classic [1950] account viewed citizenship as comprising three distinct components, which were achieved in the UK in the eighteenth, nineteenth and twentieth centuries respectively.

The first component was a civil element, which encapsulated rights of the freedom of the individual, including such aspects as liberty of the person, freedom of speech,

thought and faith, the right to own property and to conclude valid contracts, and the right to justice.

The second component was a political element, which Marshall characterised as the right to participate in the exercise of power, whether as a representative, or as a voter.

The third component comprised a social element, which was itself composed of a wide range of rights, from the right to a degree of welfare and security, to the right to play a full part in social life, living the life of a civilised being according to the standards currently prevailing in that society. Existing social rights were expanded during the mid-twentieth century through the post-war settlement that led to the huge expansion of state welfare provision. However, it is not always a straightforward matter to separate out these different components of citizenship.

The interdependence of rights

Lister's [1990] study emphasised the interdependent character of citizenship rights. For example, changes and restrictions in entitlement to the Legal Aid scheme may have acted as a deterrent to some people in exercising their right of recourse to civil law. In similar fashion, the Poll Tax of the 1980s dissuaded some from registering for inclusion on the Electoral Roll, which excluded them from the right to vote in parliamentary and local elections. These examples show how people's ability to exercise their civil and political rights may be affected directly or indirectly by their social rights.

Marshall's [1950] ideas have been criticised from a number of standpoints. For example, his liberal perspective has been attacked from the political left, for failing to tackle the problematic relationship that exists between the idea of individual freedoms and the persistence of social inequality that inevitably renders some groups and individuals less free than others in a capitalist society [Laski, 1962].

Giddens [1982] criticised the seemingly evolutionary character of Marshall's account of the historical and social processes involved in the development of citizenship. According to Giddens, citizenship rights are not necessarily as unified and homogeneous in character as Marshall's account seems to indicate. On the contrary, the three different components were achieved not only at different times, but also as a result of very different types of political struggles that occurred between different factions within society, and in response to very diverse circumstances.

Furthermore, once achieved, citizenship rights are not necessarily irreversible, as was shown by the restructuring of the UK welfare state by successive governments post 1979.

Giddens [1982] and also Mann [1987] took issue with the "anglocentricity" of Marshall's approach, which seems unquestioningly to assume that the UK is typical of the capitalist west as a whole. Others have argued that Marshall's particular preoccupation with social class has ignored, or at least obscured, issues of other important social divisions such as those based around gender and ethnicity, even within the framework of exclusively British society [Pascall, 1986].

Marshall's conceptualisation of citizenship nevertheless remains an influential one, that deserves to be taken into consideration in any meaningful discussion of the concept. It is also one that has been built upon by more recent theorists. For example, Mann's [1987] criticism formed a basis for the development of a comparative model which specifies five different strategies by which citizenship developed in response to the problems and contradictions of industrialisation in different societies, with the necessity to incorporate and accommodate first the bourgeoisie, and later the urban working class during periods of rapid industrial development.

Turner's view   Another very influential critic, of both Marshall and Mann, is Turner [1990; 1994].

Turner's chief criticism of Mann's thesis is that, in its focus upon Marxist concerns of class, capitalism and nation states, it fails to take on board that citizenship is not always a matter of rights conceded from those in authority in return for the cooperation of subordinate groups, but can also come about as a result of active pressure from below, particularly in connection with new social movements. For example, it does not address feminism, or Green movements.

Turner [1994 p. 2] defines citizenship as a "set of practices" that define a person as "a competent member of society and as a consequence shape the flow of resources to persons and social groups". In characterising it as a set of practices, rather than as a simple bundle of rights and obligations, Turner's intention is to emphasise that citizenship is socially constructed and dynamic, changing over time as a consequence of myriad political struggles. His analysis is interesting because along with juridical, political and economic practices (which are roughly equivalent to Marshall's civil, political and social components), Turner identifies also a cultural aspect of citizenship, where additional social rights might be developed in respect of culture [Turner, 1990 p.192]. This is not a new idea (although "cultural rights" did not form an explicit part of Marshall's model), but arose via explorations of the emergence of mass education in general, and higher education in particular, during the twentieth century, notably by Parsons and Platt [1973].

Not only did Turner specify the content of citizenship, but he also outlined a general theory of citizenship development. This is partially based around the differences between the mode by which citizenship can arise within a society; that is, whether it is conceded from above, or achieved from below. This distinction determines whether citizenship is "passive" or "active" in character. A second distinction lies in the relationship between the public sphere of the state and the marketplace,

TABLE 3.1: Citizenship vs. spheres.

The combination of how citizenship is achieved (from above or below) and the focus on the public sphere or on the non-public spheres provide a combination of four distinct political environments. (Source: Turner, 1990, p. 200).

	From Below	From Above
Public Sphere	Revolutionary contexts	Passive democracy
Non-Public Spheres	Liberal pluralism	Plebiscitary authoritarianism

and the non-public spheres encompassing home, family and individual moral and religious beliefs. The concept of spheres is the focus of Chapter 4. Turner's model (see Table 3.1) combines these two dimensions of citizenship to specify four political contexts in which citizenship rights can be created or institutionalised.

Revolutionary citizenship (top left corner in the table) occurs in situations where there are demands from below, and an emphasis on the public sphere. The private world of the individual is regarded with suspicion. In liberal pluralism (bottom left), interest group formation leads to movements towards rights from below, but revolutionary social protest is contained by a continued stress upon the individual and individual rights. These active citizenship forms contrast with passive democracy (top right) where citizenship rights are conferred from above, and citizenship becomes a strategy for the regulation of class conflicts by public or government agencies. The fourth situation, of plebiscitary authoritarianism (bottom right), involves management by the state of the public sphere, where citizens are periodically invited to select a leader who is then no longer accountable to the electorate on a daily basis. In this situation, the non-public spheres provide an important refuge from state regulation.

Citizenship and pervasive computing systems

The theories of citizenship we have presented were originally developed to account for the emergence of various types of citizenship rights and obligations at times when societies were undergoing the turmoil of rapid industrialisation. Yet, they may also be of use in respect of more recent, rapid social transformations that have

Citizenship in ancient times

The concept of citizenship was coined in Greece in the 6th century B.C. In the democratic polis, citizenship reflected both a status and a function. Citizenship was an inherited and exclusive quality cementing the community and conferring on a minority of people a privileged access to decision making. Being a citizen, or *politis*, in Athens was an essential condition of existence since non citizens (youngsters, women, slaves, metics, and barbarians) were considered as non humans. Socrates preferred dying than leading a life in exile as a metic deprived of his citizenship. Such a definition of citizenship based on a *ius sanguinis* principle (the law of the blood) reflects the holistic and primitive nature of Greek societies, and it also introduces an original link between belonging to a community and the political function.

In contrast to its Greek predecessor which is highly exclusive and restrictive, the conception of citizenship in the Roman Republic and Empire has an open and assimilative nature. Citizenship was no longer defined through tradition and ethnicity but asserted by jurisdictions and law which distinguished between three different rights: *ius conubii* (individual laws), *ius commercium* (economic rights), *ius suffragio* (political rights). Citizenship became an instrument of power via integration. It accompanied the expansion of the Roman Empire and was granted to every newly conquered tribe. Such a revolution was due to the universal and assimilating nature of the Roman regime whose ambition was to wield and expand its leadership over the world. In doing so, the Romans spread law and individualism which are two core notions of modern definitions of citizenship.
(Source: http://www.chez.com/bibelec/publications/international/p1.html)

occurred (and are currently occurring) as a result of new technologies [Agre, 1997; Castells, 1995]. Turner's thesis in particular, with its emphasis on dynamic social practices rather than just the content of citizenship rights and obligations, lends itself to this type of consideration.

Unfortunately, the pervasive systems community, despite the far-reaching effects, consequences and potential of pervasive technology, has not taken up the ideas presented in this chapter. This could be a result of the fact that the pervasive systems community is quite new, and so far has been most heavily influenced by hardware designers and middleware designers. Some HCI researchers, however, have been very interested in the issues we have discussed here.

For example, issues of usability in respect of particular groups such as the elderly [Hanson, 2001; Tilley, 2003] and disabled people [Edwards et al., 1995] have been linked with the concept of social citizenship. Citizenship has also been discussed in respect of the provision and dissemination of local and national government infor-

mation to the general public [Catarci et al., 2000; Donnelly & Merrick, 2003], and has also been linked to concerns about rights to privacy and private information [Ackerman et al., 1999; Rezgui et al., 2002]. Meanwhile, Dearden and Walker [2003] have begun to examine issues of citizenship and technology in respect of "civil society".

There can be many design implications for pervasive systems based on the theories and views we have discussed so far. For instance, freedom of speech and the right to own and control one's own property are examples that can influence our designs of pervasive systems, just like they influence our design of urban spaces and buildings. If they are to become part of everyday life then they must support the ways and rules of everyday life. Furthermore, as different societies have different values (what is acceptable, how to behave in public, laws), the pervasive systems we design should reflect those differences.

What we have tried to do here is to lay out the basic characteristics of a citizen. Again, we may not know much about users the users of a publicly available pervasive system, but we do know a lot about citizens. *Users don't have rights, but citizens do.*[16] For the purposes of designing for pervasive information access, we must focus on the citizen aspects rather than the user aspects.

In the latter half of this chapter we explore a design implication which stems from a citizenship right: the right of having access to information. We believe that this right should be reflected in our designs of pervasive systems. In terms of design, we discuss the provision of a publicly available pervasive system as a public service. To do this, first we discuss public services in general and try to distil some public service characteristics that we can use in the design of pervasive systems.

---

16. Ultimately, users have rights because they are citizens.

3.3   Pervasive systems as public services

We believe that all citizens should have access to a public pervasive system. This means that a wide variety of people, including young and old, male and female, and people of various abilities, ethnicities, etc. should be able to use the system. This issue is addressed by the notion of universal access [Stephanidis, 2001] or universal usability [Shneiderman, 2002]. The goal of these approaches is to make computer systems accessible and usable by everyone, much like television, electricity and cars. This noble goal also has roots in legislation such as the US Communications Act of 1934, which attempts to ensure that facilities are provided without "discrimination on the basis of race, colour, religion, national origin, or sex" (section 1, 47 U.S.C. 151). Furthermore, universal access would allow government-based electronic facilities to be used by the public, at its most simple level including, for example, electronic voting.

Apart from the user interaction issues that universal usability addresses, other issues are raised by the citizenship rights described earlier. These imply that *people require adequate and unbiased information circulating in the public sphere*[17] *in order to make informed decisions and to take part in the everyday democratic process.* This is recognized by legislation and has resulted in the concept that access to information is a public good [Green, 2002].

Are pervasive systems a public service by definition?

A public pervasive system with sufficiently wide area coverage may be regarded as a nationwide carrier of information which is accessible by the public. In our society, access to information is seen as a public good [Green, 2002]. In western societies, a nationwide carrier of a public good is considered to be a public service. So, truly pervasive systems with widespread public access may be viewed as essentially a public service [Kostakos & O'Neill, 2004a]. This perception in itself can have great

---

17. For a discussion of the public sphere, see Chapter 4, page 77.

consequences for the way in which the system is used, what people expect from it, and indeed what people demand of it.

We now turn our focus on existing public services and discuss how they are defined, what people think of them, and how they operate. It can be helpful to the reader to read the following section with pervasive systems in mind. Thus, whenever we refer to the "users of public services", the reader should keep in mind that this could be substituted by "the users of public pervasive systems". In Section 3.5 on page 72 we present an analysis of how useful is this analogy between public services and public pervasive systems.

3.4   The characteristics of public services

The specific details of what consists of a public service varies depending on what it supplies, to whom and by whom. There are at least three different approaches to defining the meaning of a public service [Krajewski, 2001]:

- Services considered as public or for the common good. Such services may include health, education and transportation. This view focuses on what is supplied.

- A service provided to the general public. This is often understood as the notion of a "universal service obligation", i.e. the obligation to supply the service universally at affordable terms, without distinguishing between the costs of supply in different regions. Thus, this view focuses on to whom and under which conditions the service is supplied. It should be noted that most services from the first category are often attributed with a universal service obligation.

- A service provided by a public entity, such as the government, a governmental agency, or a public enterprise. This view focuses on who supplies the service.

Devolution and participation   Two issues relating to public services are devolution and participation. For the past 30 years, there have been numerous attempts to introduce measures to increase dev-

olution and participation in public services [Tunstall, 2001]. Despite the frequent use of these terms, explicit definitions are hard to find.

Definitions of participation range widely. For instance, the UN Development Programme requires people to have "constant access to decision-making and power". Other views are wider, only requiring participants' views to be requested in order to "consider them before a decision is made" [Tunstall, 2001].

Defining devolution can be problematic because of the debate on what is being devolved: power, responsibility, legitimacy, resources. Decentralization and devolution have been used interchangeably although decentralization refers to power shifts within a single organization, while devolution requires power shifts to autonomous organizations.

Obligations

Across the EU, obligations on public service providers have been imposed [Harrison & Woods, 2001], such as security of supply and obligation to supply. The public takes reliable public services for granted. Apart from being reliable, public services must be universal, i.e. all of the public should be equally entitled to benefit from it. Although some customers may be more desirable than others, the obligation to supply demands that everyone has access to the service.

Advanced economies

In advanced economies, however, what defines a public service is increasingly complex and blurred. In seeking to define the role of public services, it is worth remembering why they exist: to *improve the quality of people's lives*. However, this definition is too broad, and could be interpreted to include Tesco's supermarkets or Virgin Holidays. Some key characteristics of public services are [Stoker & Williams, 2001]:

• They rely (directly or indirectly) upon an element of taxpayers' money to be provided.

- They accept an extended type of accountability, and are subject to a form of democratic scrutiny.

- They have a defined customer base. This means that most public services are unable to choose their customers, and most customers are unable to choose their public service supplier.

In a recent New Local Government Network (NLGN) commissioned survey [Stoker & Williams, 2001], the public was asked what defines a public service. The two top definitions of public service were: "available for everybody to use" (40%) and "important to the whole community" (38%). Furthermore, 23% of respondents believed that a key definition is the management of the service by central government, and only 4% thought that a public service had to be provided free at the point of use. This agrees with other opinion surveys that have found that the role of the private sector in public services is a low salience issue. This may relate to the fact that profit is not a barrier to the delivery of effective public services.

Functional characteristics

Apart from the economic and political characteristics of public services, there are some functional characteristics that can be useful in our analysis. Public services tend to provide a stable, static product that does not change very often [Patterson, 1999]. In fact, changes to public services need to go through the public's scrutiny.

This relative stability and coherence has some very important implications. The products and services persist over a long period of time, usually spanning more than one generation. As a result, they become embedded in the way of life of individuals, who then reinforce this way of life upon their children [Kostakos & O'Neill, 2004c]. Most children are told not to put their fingers in the electricity sockets, not to dial random numbers on the phone, to call 999 if an emergency occurs. Tenets such as these become so embedded in our lives that they tend to pervade the social and cognitive environments of people. This is well illustrated in emergency situa-

tions, when people may be overwhelmed or stressed, yet are still able to summon the emergency services or get access to electricity and water.

Offering a pervasive computing system as a public service, and submitting it to the general functional characteristics described above, would result in a system that pervades the social and cognitive environments. People would use this system so many times, and would have enough experience with it, that they would be able to use it even in the most extreme situations.

3.5   Implications for the design of pervasive computing systems

Traditional public services, such as the telephone, tend to pervade the social and cognitive environments as a result of people using them repetitively and from a very young age. A pervasive system that is offered and perceived as a public service could reach its full potential by pervading the physical, social, and cognitive environments to a similar extent.

A close inspection of successful public services reveals a number of common characteristics. For instance, the expectation of receiving the same service regardless of physical location or who is using it. This, for example, does not imply that the service cannot be personalized; indeed personalization plays a key role in user satisfaction. Instead, what the users expect is that personalization is available to others too.

Increased use, familiarity, reliability

In order for pervasive systems to benefit from the frequent use and familiarity that other public services enjoy, care must be taken that the existing paradigms of public services are followed.

For instance, reliability has to be a key characteristic. Any attempt to provide a service to large numbers of people is bound to run into problems if the users are expected actively to contribute to the continued operation of the system, or any

part of it. This could also result in users feeling that they are actually contributing more than they get in return. When using a publicly available resource, we expect to be treated in the same way as everyone else, not better or worse. Furthermore, we expect the service to be reliable and not dependent upon the users for its operation. This can be seen in public transportation (buses, trains etc) and public services in general (hospitals, police, fire brigade, etc).

Centralised structure

We observe a somewhat centralized structure when it comes to delivering public services. This approach is appropriate for delivering uniform and consistent services. This has resulted in the development of notions and ideas that are applied to new forms of services, as they come along. Good examples are the concepts of a police "station", a telephone "centre" or an Internet "provider". Furthermore, not one of these services actively relies on its users for its day to day operation. Users may enjoy the services without much work. The stability and consistency of a centralized service provider have been preferred over a flexible decentralized system in which the user has increased responsibilities. This can be the case for pervasive systems as well.

Maintaining the required stability and uniformity for a service can best be achieved with a centralized way of providing the service, at least when it comes to pervasive computing. Also, the users should be treated uniformly, regardless of physical point of access, real identity of user, social status of the user, etc. Furthermore, the services need to be simple enough to be used by anyone, regardless of their age, education, gender and race. Also, dedicated and specialized personnel should be available for repairing damage and faults to the system, much like the telephone repair personnel who are responsible for fixing problems with the phone network.

Of course, in terms of technology, implementation and design, pervasive systems are nothing like e.g. the bus service, both from the provider's and the user's perspec-

tive. However, we have identified a number of common characteristics that are found in almost every successful public service. Therefore, we should try to incorporate such characteristics in our pervasive system designs, or at least provide design characteristics that cause users to perceive the new technology as yet another form of publicly available service.

3.6 Summary

In this chapter we argued that it can be beneficial to view the intended users of a publicly available pervasive system as citizens. We then proceeded to define the term citizen according to both classic and modern ideas and principles. Based on these conceptions and debates over citizenship, we have argued for a social good model of provision for public pervasive systems. This model is also referred to as public services, and so we explained this term and related notions and concepts. Besides the political and economic characteristics of public services, we pointed out that functional characteristics could provide a source of design ideas for public pervasive systems. We also argued that by offering pervasive systems as a public service, these systems could reach their true potential of pervading the social and cognitive environments (in addition to the physical environment). Furthermore, throughout this chapter we emphasised that users don't have rights, but citizens do. This may not have direct design implications, but it can have direct consequences, as we saw in the opening of this chapter.

In the following chapter we describe the second element of our framework - spheres. Specifically, we discuss the notions of public, social and private spheres, notions which we have mentioned briefly in this chapter (Table 3.1 on page 65). Furthermore, the issues of designing pervasive systems as a public service are of relevance when we discuss the notion of the public sphere.

CHAPTER 4

INFORMATION SPHERES

4.1 A top-down approach to describing tasks

We have argued that while domestic pervasive systems are typically optimised for quite narrow, specified purposes, public pervasive systems need to be much more flexible, in order to offer useful, usable computing resources to the public. This raises challenges to our efforts to guide designers of public pervasive systems. How can designers possibly account for any arbitrary activity that any particular user may wish to perform using a public pervasive system? To begin with, a system could offer specific services, thus limiting the number of activities the users could perform using the system. In time, the system could be extended to cover many more activities, thereby making the system useful in more situations.

But this type of bottom-up approach is unlikely to allow our system to reach its full pervasive potential. It is simply impossible to account for all possible activities that all possible users could wish to perform using a public pervasive system. Our proposed top-down approach, on the other hand, attempts to define such activities abstractly and to offer support in reasoning about and designing from these abstractions. The established HCI design focus of task is, unsurprisingly, more suited to the design of conventional desktop software applications and the domestic pervasive systems of Table 1.1, supporting tightly specified purposes. Our approach replaces the concept of task with the concept of spheres.

What are the spheres?

In operationalised terms, a sphere is *a pool of information*, or a *resource*. As we describe in this chapter, the type of information that exists in a sphere defines the classification of the sphere itself. Furthermore, the notion of spheres, especially the "public sphere", carries a lot of semantic value mainly from the domain of sociology. The public sphere has been related to public discourse, the degree of freedom within a society, and the degree to which citizens can participate in the democratic process. [Green, 2002].

Citizens and spheres

For instance, Malina [1999] examined the relationship between citizens and spheres, drawing attention to the role of the media, particularly new media such as the Internet, as "spheres of public debate" [p.23]. The emergence of computing systems has effected a transformation in respect of our relationship with media forms, from the paternalistic, one-to-many characteristics of the traditional media to the possibility, at least, of more democratic and participatory media forms. Malina drew attention to the possibilities provided for broadening aspects of democratic practice. These possibilities are dependent, however, upon factors such as whether the types of information necessary for democratic participation are packaged as an easily (and cheaply) available "social good", or sold as a costly "consumer

product" [Malina, 1999, p.38]. In arguing for pervasive systems as a public service [see Kostakos & O'Neill, 2004c], we propose a "social good" model. The spheres element of our framework obliges us to consider the activities to be supported by a pervasive system in terms of the types of information required by citizens.

The 3 types of spheres

Drawing on concepts from sociology and social policy, we have identified the *public* sphere, *social* sphere, and *private* sphere as three general categories of the types of activities that citizens might be involved in and for which they might require, use and share different types of information. These spheres allow us to categorise specific information and associated activities or services offered by a public pervasive system. For example, we can say that the activity of sending a message to a friend and the corresponding digital service that allows us to do so fall in the social sphere, whereas the activity of looking up a train timetable falls in the public sphere. This assignment of information and services to specific spheres can help in the design of pervasive systems. Examples of this can be seen in Chapters 8 and 9. For now, we describe in detail each of the three types of information spheres we have identified.

4.2   Public Sphere

The term public sphere was introduced by Habermas [1962], who is linked to the Frankfurt School of social thought. The Frankfurt School was inspired by Marx but argued that he had not given enough attention to the influence of culture in modern capitalist society. They studied what they called the "culture industry", including film, television, radio, newspapers and magazines. They argued that the spread of the culture industry, with its standardized and simplistic products, undermines the capacity of individuals for critical and independent thought. Habermas was influenced by these themes, and developed them in a different way. He analysed the development of media from the early eighteenth century up to the present day, tracing out the emergence - and subsequent corruption - of the public sphere.

The public sphere is a conceptual area of public debate in which issues of general concern can be discussed and opinions formed. It has also been defined as the space in which citizens deliberate about their common affairs, and a site where social meanings are generated, circulated, contested and reconstructed [Fraser, 1995]. The public sphere, according to Habermas [1962], was first developed in the salons and coffee houses of London, Paris and other European cities. People would meet in these salons to discuss issues of the moment, using as their primary means for debate the news sheets and newspapers which had just started to emerge in that period. Of particular importance was political debate, and although only small numbers of the population were involved, such salons were vital to the early development of democracy. The public sphere, in principle, involves individuals coming together as equals in a forum for public exchange of information and debate.

Hence, the types of information residing in the public sphere include the many types of information of common interest and to which access is not restricted. The list of examples is infinite but includes access to train timetables, news reports, debates, voting schemes, or even laws and court decisions. The degree to which the public sphere is controlled, monitored and constrained reflects the degree of freedom of access to information, freedom of speech and freedom of civil and political rights within a society.

Public pervasive systems have the potential to offer citizens opportunities and means to participate in the public sphere. There are currently technologies and systems optimised for each of the above examples of public sphere information, but these are in the tradition of conventional applications and the domestic pervasive systems of Table 1.1. For example, government websites provide comprehensive libraries of legislation while Internet newsgroups and chat rooms facilitate debate

and gossip on all manner of subjects. But access to these resources is still restricted to a relatively small subset of the public with the wealth, knowledge and equipment to avail of them. For a fully working public sphere, it is necessary that citizens have access to the public sphere with a minimum number of obstacles. These may include, for example, economic obstacles (having to pay a fee to access the public sphere, having to buy equipment to access the public sphere), physical obstacles (unavailability of resources to access the public sphere in the physical environment to which one has access), educational and other obstacles.

The designers of large-scale truly pervasive systems should not aim simply at providing access to publicly available information. Rather, pervasive systems should aim at reinforcing, helping, sustaining and feeding into the public sphere. This would attribute importance to pervasive systems as they would be contributing to everyday discourse. A one-way communication channel between public information and citizens is not enough to maintain the public sphere; there must be two-way communication and interaction between the public sphere and the citizens who wish to contribute to it. We are not suggesting that this two-way communication allows for complete anarchy in the public sphere (imagine everyone being able to modify what everyone else has contributed). Rather we perceive this two-way communication as the bare minimum for exchange of ideas and information which would sustain the public sphere.

4.3 Private sphere

The "non-public sphere" consists of two parts. The first category is the private sphere, which deals with completely private issues and information whose owners would not want to be accessed by others at any point. The second category is the social sphere, which we describe in Section 4.4.

**The Public
Sphere Project**

The Public Sphere Project (PSP) is a CPSR (Computer Professionals for Social Responsibility) initiative to help promote more effective and equitable public spheres all over the world. Doug Schuler, a longtime CPSR activist and Seattle Community Network co-founder, is the Program Director. The PSP is an outgrowth of the "Shaping the Network Society" symposium convened in Seattle in May, 2000. The Public Sphere Project is intended to provide a broad framework for a variety of interrelated activities and goals including event organizing.

The project's objectives are to advance our understanding of opportunities and challenges of the public sphere for democracy, education, social justice, economic development and environmentalism, and to develop and act on strategies for creating and strengthening an equitable and effective public sphere. More information can be found at http://www.cpsr.org/program/sphere/.

We can envision the private sphere as a bubble around each one of us, containing information that nobody else can access, and which follows us around wherever we go. Information that belongs to an individual and is deemed completely private is said to belong to the private sphere. In Chapter 1 we noted the need to allay the fears of citizens about the potential dangers and abuses of pervasive computing. These fears are most apparent when considering the private sphere. With pervasive technologies there exists huge potential for conflict between the capacity of pervasive systems to make information accessible from everywhere at any time and the privacy requirements and restrictions on private and sensitive information.

For example, we may view a person's everyday wallet as a (non-digital) technology that provides access to part of that person's private sphere, including for example her bank account. This "wallet-technology" allows people to access private information with relative privacy, and also allows them safely to store information.   If we then provide a "digital wallet" that allows one to access similar private information from her PDA, her car, any phone booth and so on, the range of possibilities increases with a truly pervasive system, providing access to digital wallet services from more and more locations, devices and situations.[18]

Existing research has addressed issues of privacy, in either traditional computer systems, mobile technologies, networks, but also at a theoretical level [Ackerman et al., 1999; Rezgui et al., 2002; Westin, 1970]. This work is quite relevant to our discussion (after all, privacy was an issue long before we starting thinking about pervasive systems). For instance, the definition of privacy as "the ability to decide for themselves when, how, and to what extent information about them is communicated to others" [Westin 1970, p.7], is very similar our notion of a private sphere (in the sense that the private sphere has a boundary which one can control).

However, our goal here is not to develop a fully fledged theory of privacy, nor to contradict existing research. Instead, we believe that by proposing the concept of the private sphere, we can put concerns about privacy at the heart of the design of pervasive systems. The notion of spheres is a step towards formalising the design requirements for pervasive systems, identifying research that has addressed privacy requirements, and proposing a coherent way of introducing these concerns in the domain of pervasive computing.

The designers of pervasive systems need to take into account the sensitivity of such private information and provide designs that prohibit the nature of pervasive systems from compromising the privacy of such information. Having a range of options and locations to access private information means that the same information is not only accessible from different settings, but also by different technologies. As we describe in Chapter 5, different technologies create different interaction spaces, with varying degrees of privacy. One of our goals is to describe which interaction spaces (and thus which enabling technologies) are suitable for specific situa-

---

18. There is a distinction between the information we access and the technology we use to do so. It so happens that with traditional wallets the information and technology are always in the same physical location. This does not have to be the case with digital technology.

tions, thus supporting design decisions in the development of truly pervasive systems that support the private sphere.

4.4   Social sphere

The social sphere describes activities and information of a semi-public nature. The information within a social sphere is neither private nor public. It is useful to think of such situations as not being private because more than one person is involved in the creation and exchange of information. Furthermore, these situations are not public either, mainly due to social or physical constraints such as having to pay for a ticket at the cinema in order to watch the film. Further similar constraints include the ability to deny entry to your home, personal relationships, and even physical constraints such as the inability to listen to someone who is physically far away. We can manipulate many of these constraints using computing systems. For example, computing resources can be used to listen to someone in another country or to control access to a film. Hence, we can use computing systems to help define the social sphere.

In recent generations, the telephone has been the technology most associated with the development and maintenance of a social sphere. The phone is a way of extending one's private boundaries beyond the individual or the home to family and friends who are welcome to call at any time. The phone is used to capture a friend or household and bring them into an elastic, psychological domain of social space. A private call has the effect of relocating the other psychologically within the social sphere. A person may access more than one social sphere simultaneously, and may be actively involved in more than one, and may pass information between these social spheres and a private or public sphere.

The nature of a social sphere is much more dynamic than either the public or private spheres. Social spheres are created and destroyed ad hoc, in various kinds of

situations. For example, a group meeting, a conversation, a note on the refrigerator are all instances of a social sphere. These are examples of "pools of information" to which only a few people have access. There can be more or less permanent social spheres, which may take the form of shared repositories of information and resources, accessible by a limited number of people. But the vast majority of instances of social spheres are ad hoc, day to day instances (playing a board game, watching TV with friends, going to the movies).

There has been a lot of research on supporting group tasks, on developing groupware, as well as studying in detail the collaborations that take place amongst members of groups. For instance, the locales framework [Fitzpatrick et al., 1996; Fitzpatrick et al., 1998] is a framework based on sociological observations for designing CSCW systems. This framework deals with shared group resources and "worlds" which are similar to our notion of a social sphere. As we have said about the private sphere, we do not seek to develop a theory of group work. Our goal is to understand how to design for such situations of semi-publicness in a pervasive computing setting.

The approach we have taken to researching the design of pervasive systems does not involve the development of a comprehensive theory either of privacy or of groupwork. Instead, we have tried to map out these areas and relate them to the practice of designing pervasive computing. In Chapter 1 we explained how the domain of pervasive computing is currently very disjoint and lacking substantial theory. Although researchers in pervasive computing are aware of privacy concerns as well as the need for pervasive computing to support groups of people, it is still not obvious how existing work in other areas can be leveraged and brought into pervasive computing in a coherent fashion. We believe that great value can be realised by proposing concepts (such as the spheres) which can help in the design of pervasive

systems, and then relating them to existing work and pointing out possible directions for future research within pervasive computing.

In the following sections we show how the notion of spheres can assist in the design of pervasive systems.

4.5   Implications for the design of pervasive computing systems

So far, we have described the basic characteristics of the public, social and private spheres. Many of these characteristics have been explored in other fields, such as sociology, social policy and philosophy. In the latter half of this chapter we describe how the ideas we have presented so far can help us in designing pervasive computing systems.

To do so, first we derive some more tangible and operationalised *characteristics* for the three kinds of spheres, and also explain the *relationships* amongst them in a more detailed and practical fashion. Having done so, we then discuss how the concept of spheres is a useful way of *thinking about* pervasive systems from a citizen's point of view (but also from a designer's perspective).

Characteristics of the spheres

We have already said that the spheres can be conceptualised as pools of information. The public sphere, for instance, contains information on public discourse and information of a public nature. Such information is located in a wide range of physical and digital locations. The same can be true for the private and social spheres as well. It is up to the technological design to provide a coherent way of presenting this fragmented set of information as an information pool, and allowing access to it.

But what do we mean by "accessing" an information sphere? Access to a sphere can happen in *physical* form, *digital* form, or as a *spillover*. We propose three straightforward ways in which one may access an information sphere.

First, a sphere may be accessed *physically*, in the sense that the technology being used to access the information is non-digital. For instance, reading a handwritten note, seeing a poster, or listening to someone next to you talk, are all ways of physically accessing a sphere. Furthermore, writing on a paper, or talking to someone are also ways of accessing (but this time not consuming but contributing to) a sphere.

Accessing a sphere in *digital* mode means using digital technology (i.e. technology that uses electricity) to access the information. Accessing a website, a newsgroup, and an on-line chatting room are all examples of this. Furthermore, posting a message and sending an email are also ways of accessing a sphere by contributing to it.

The third and final mode, which is also the most interesting, is accessing a sphere by accident, or by *spillovers*. For example, when two people are having a conversation on the bus, others can overhear them by means of a spillover. When a private conversation is posted on the web, others have (unauthorised perhaps) access to this information as a spillover. When using your laptop, those who can see your laptop's screen may gain access to information as a result of a spillover. These situations can happen both during physical or digital access of a sphere.

Furthermore, it is not necessarily the case that spillovers are malicious (for instance the people on the bus who can overhear the conversation are not necessarily malicious agents), but we can certainly think of situations where people with malicious intents can take advantage (or even cause) spillovers.

The physical way of accessing information spheres has (mainly) to do with physical constraints, which are controlled and manipulated by the architecture of the environment. The digital way of accessing information spheres can be controlled and manipulated by the interaction spaces which are created by the digital technology.

These two different types of space are discussed in Chapter 5. Understanding these two different sources of constraints (or possibilities) can help the designer in choosing a solution or remedy to a problem (or spillover) accordingly, as we will see in the case studies later.

Relationships between the spheres

We envision the existence of one public sphere, which every citizen is entitled to access. Furthermore, each citizen has her own private sphere, which contains their private information. Finally, citizens participate in various social spheres, either dynamically or on a continuous basis.

A person, therefore, has access to more than one sphere simultaneously and over time. The issue that we now need to address is how the simultaneous access of more than one sphere by a person affects the relationship between spheres. We have said that people should ideally have unlimited two-way access to the public sphere. We have also said that access to the private sphere is ideally restricted to the "owner" of the private sphere. Finally, people might want to share information and resources amongst a specific group of people, in the form of a social sphere. These requirements and limitations are introduced by the very nature of the three kinds of spheres.

We now examine situations where information is moved from one sphere to another. Let us consider a simple situation where a group of people are sharing information within a social sphere, and the participants choose to share and contribute information by sharing some of the contents of their own private spheres. In this case, person X would have simultaneous access to her private sphere and the social sphere of the group. To successfully share information, X would have to access information from her private sphere and feed it into the social sphere. She would have to access her private sphere in such a way that spillovers are avoided, as they would compromise the privacy of her personal information. Furthermore,

another person, Y, might find some of the shared information interesting. He could therefore take some information from the social sphere and feed it into his private sphere. Again, this would have to happen without any spillovers to avoid compromising private information.

If in this example we consider the public sphere instead of the social sphere, not much would change. Accessing the private sphere would again have to happen with care, only this time the potential for spillovers could be greater, and the consequences potentially more severe.

Spillovers are once more of concern in situations where the public sphere is being accessed simultaneously with a social sphere. To continue with our example, a group member might have found something interesting in the public sphere, and would like to share it with the other group participants. To do so would require accessing the public sphere in order to get the interesting information, and then accessing the social sphere in order to contribute this new information. This time, spillovers could be caused by any of the participants in the social sphere, so in this sense it becomes harder to control the privacy of information in a social sphere: to do so would require knowing how each of the participants accesses the information, and making sure that no spillovers occur while doing so.

So far we have talked about accessing spheres, and what the designers need to be careful about. We have done so without referencing any of the other factors that affect the accessing of spheres. One such factor is the physical location where one is present.

We have presented this idealised view of the spheres for three reasons. First, we wanted to demonstrate in principle the ideas and concepts we have introduced, and explain how they interact with each other. Secondly, we have not finished the

description of our framework, and what remains to be shown in Chapter 5 is of direct relevance to the ideas we have discussed here. Therefore, a more thorough discussion of the ideas presented in this chapter will be possible once all three aspects of our framework have been introduced. Hence, in Chapter 6 we present a discussion of all three aspects of our framework. However, the third reason for discussing spheres in the present chapter is because we believe that the notion of spheres, on their own, can provide a way to think about and conceptualise pervasive systems.

Reasoning using the spheres

Earlier we discussed the problems caused by various systems having their own view of what the user is trying to do, through the example of someone wishing to share their photographs with friends (see page 45). Each of the systems involved in carrying out this task has its own view of its world, its own metaphors and objects for supporting its view of the user's task. This leads to problems of consistency: one moment a photograph is a file, then it is a thumbnail, then it is a workbook, then it is on the clipboard, then it is an email, then it is a website, etc. There is no consistency in the task of a user or with the digital objects relating to the task. The nature of pervasive systems could exponentially increase such problems. Accessing information using any device from any location sounds like a useful idea, but as we have discussed in Chapters 1 and 2 many potential problems exist.

We acknowledge that a public pervasive system will be composed of various systems, and probably most of them will be created by different developers. We therefore see the need to create metaphors and objects at a higher level and more abstract than the system level. Most suitable is the information layer: information itself transcends all systems and subsystems.

In a pervasive environment, information is envisioned to travel seamlessly across systems and services, to reach its final destination or to support a specific task that

someone is carrying out. If, therefore, we envision a pervasive system that is composed of such seamless systems, at what level should we aim to describe "the system" to the users? If, on the other hand, the components of a pervasive system are self-oriented in how they view the world and the tasks that people carry out, then how ought one to model the system from an end-user's perspective?

Our answer to both question is the same: model the system based on the information that it contains and transfers. Why is it helpful to think in terms of information spheres? We can consider this both from a citizen's point of view and from a designer's point of view. The private sphere is a way of declaring one's own private (digital) boundaries within this everywhere anytime access to information. Citizens do not have to worry about all the various mechanisms and systems that they could possibly need to use in order to access their private information. They don't need to learn new metaphors and a new language in order to get access to their information. Similarly, from a designer's perspective, providing and maintaining a private sphere is what they need to worry about. Obviously, the mechanisms and technology to do so could be wide-ranging. With this metaphor in place, the systems, software and hardware, can all be adapted to support the access of private spheres. "Anytime anywhere access to information" is too broad; *access to the private sphere without spillovers* can provide much needed guidance as to what citizens require of a pervasive system. In some cases, spillovers may be a required effect. Since we can design to avoid spillovers, we can also design to create spillovers, if indeed they are required.

The situation with the social and public spheres is similar, but as we move "up" the scale (i.e from private to social to public) the privacy requirements become more relaxed. Identifying the different social spheres in which one is participating is a good way of separating out the various resources to which one has access. Pervasive

systems aim to offer access to the same resources from any location, using varying technologies, and at any time. Therefore, the experience of accessing the same resource under different circumstances could vary. This is exaggerated if we imagine many people accessing shared resources, both within a specific group and outside this group. Therefore, the notion of a social sphere, or a name to a social sphere can provide helpful memory aids, and be a way of controlling the experience of people. In situations where private and public information is being accessed (such as our hospital case study in Chapter 8) we propose that the interface informs the user of the type of information being accessed.

Accessing the public sphere should notify citizens regarding decreased privacy, or to notify them that they are currently accessing resources that don't belong to them, or they are in an "unsafe" digital location. This should be seen in relation to the previous two kinds of spheres; we can consider someone traversing private and social sphere resources using pervasive technology. Should they be notified if they somehow start browsing public sphere resources? If so, how should this notification be done? Possibly this could be a context-aware function. For instance, when using a private device, it can be unwanted to access the public sphere. Or in certain locations it can be unauthorised to access the public sphere (such as in an exam room). Whichever is the case, the notion of a public sphere lends itself to these kinds of considerations.

4.6   Summary

Throughout this chapter we have raised issues with implications for implementation. In addition, we argued in Section 4.2 that access to the public sphere should ideally be two-way, and not just a one-way downward feed. This has implications for the kind of technology used. For instance, two-way communication would

require both input and output technologies (software and hardware) at the "point of use" of pervasive technology.

Furthermore, designers also need to consider accessibility issues. Certain input and output technologies may result in certain underprivileged groups of people who have limited access to the public sphere. Of direct relevance to this are the research areas of universal access and universal usability, which we have discussed in Section 3.3.

We have also said that the private sphere should be accessible only by its owner. To this extent, designers need to rethink user authentication and user recognition, and figure out how to provide a secure model without affecting usability. We have experimented with embedded biometrics, and have proposed the use of biometrics where natural physical interaction occurs [Garzonis et al., 2004]. By doing so, and at the same time extending the network protocol headers, we believe that user authentication can be achieved with minimum cognitive load and explicit input from the users.

The contents of a private sphere, as we have already noted, may be dispersed in various locations. Therefore, we need to provide a coherent way of presenting the contents of a private sphere. This also relates directly to the need we have identified for providing users with high-level metaphors and objects. The private sphere should be the main and strongest object or metaphor. Some technologies are attempting to do this for the web. For instance, the Microsoft .Net Passport system[19] allows users to have their contact details in "one" place, and access them from various websites in order to, for instance, fill in a form. Although simplistic, it is a good indication of how we envision similar functionality for the private sphere. Designers need to adapt and extend this functionality in order to provide

19. See http://www.passport.net

access to the contents of the users' private spheres from any part of the system, under varying conditions, and using different resources to do so.

Carrying on from providing access to the private sphere, we have also raised concerns about possible spillovers, and the need to minimize them. In relation to this is the notion of interaction spaces which we describe in Chapter 5. In terms of spheres, designers need to make sure that unauthorised access to spheres is minimised.

Finally, we have raised some issues that should be of concern to the designers of human-computer interfaces for pervasive systems. We mentioned the need for users to be aware of the different social spheres or public sphere which they are accessing. These issues are at the human-computer interface level, and existing methods and techniques from other domains could be brought in.

The concept of spheres was discussed in relative isolation. We need to keep in mind that these ideas are useful in conjunction with the other two elements of our framework: citizen and space. In the following chapter we discuss the third and final element of our framework - space. Doing so will allow us to engage in a more fruitful discussion of the ideas we have seen so far, and relate them to the design of pervasive computer systems. The application of the ideas we have developed so far will become apparent in Chapter 6 were we discuss our design tool.

CHAPTER 5

THE IMPORTANCE OF SPACE

5.1   Places, spaces and domains

The Concise Oxford Dictionary defines *place* as "a particular position or point in space", and *space* as "a continuous area or expanse which is free or unoccupied". Norberg-Schultz [1971] takes the view that a space is the physical manifestation of something, where as a place is the subsequent interpretation that a person has of the space. Harrison and Dourish [1996] propose that we are located in space, but we act in place. According to geographical theory, space refers to abstract geometrical extension and location, while place describes our experience of being in the world and attaching meaning, memories and feelings to physical location. Thus, place incarnates the experience and aspirations of people [Tuan, 1971]. Perhaps the rela-

tion between space and place has best been described by Joseph Joubert: "Space is to place as eternity is to time". There are subtle differences between the two terms, but in everyday language these terms tend to be used interchangeably. Despite our (mis)use of these words to refer to simply physical location, physical locations have much more to them. Physical locations have a purpose, they provide opportunities for specific circumstances to arise, they have values and a history. These characteristics are very hard to describe, but they are experienced by the occupants of physical locations.

Just as our use of the words place and space devalues their meaning, so in trying to construct and reflect the notion of domain, many pervasive systems today use location as a substitute. Unfortunately, the notion of location lacks many important qualities, such as the social dimensions, that are intrinsic to the notion of a domain. In trying to identify the social dimensions of a domain, some have studied the differences between physical location and their respective social understandings [Harrison & Dourish, 1996]. A domain has embedded understandings and protocols of what is regarded as appropriate behaviour. Domains have values attached to them. Domains tend to convey cultural meaning and frame our behaviour. In addition, the presence of others in a domain has an effect on how we behave and what we perceive.

These issues and characteristics become very complex when considering public pervasive systems. With such a multitude of locations being covered, each one with its own peculiar characteristics, how can we design a system that will take into account all these different domains? Can we do better than simply monitor the physical location of people and artefacts?

In the third element of our framework, we propose a top-down approach that categorises all possible spaces into three main groups: *public spaces*, *social spaces* and *pri-*

vate spaces.[20]   These terms are borrowed from sociology, e.g. [Green, 2002], and have influenced a number of researchers. For instance, Hall [1969] specified four interpersonal distance zones: intimate, personal, social and public. The first two map to our notion of private space, while the latter two directly map to our notions of social and public space. These notions carry with them the qualities of a domain, a great number of characteristics and understandings that are peculiar to each society or social group, while at the same time highlighting the importance of physical location.

Interaction spaces

We make a further distinction between the spaces created by our physical environment and the interaction spaces created by artefacts including computing and communications devices [O'Neill et al., 1999]. We define an interaction space as the volume of space within which the device or artefact is usable. Interaction spaces depend on the type of technology used, as well as the physical characteristics and affordances of the technology. Similar to spaces defined by the physical environment, interaction spaces may be private, social or public. For example, in Figure 5.1 the plasma screen positioned in front of the two people creates a social interaction space that includes both of them. The person on the right is wearing headphones, which create a private interaction space for him. The other person's PDA can create different types of interaction spaces, depending on its position and orientation. The owner of the PDA may choose to tilt the PDA towards himself, thus leaving the other person outside the interaction space created by its small screen, or he could choose to position it in such a way that they can both use it. Interaction spaces are the means by which we can use digital technology to access the information spheres.

---

20. Although we call them spaces, they are places in the sense described by Harrison and Dourish [1996].

FIGURE 5.1: Various examples of interaction spaces.

The plasma screen offers a social interaction space, headphones offer a private interaction space, while the PDA can do both depending on its orientation.

Interaction spaces are also the means by which digital spillovers are created (see page 84). For instance, if someone is working with her PDA in a bus, others may be included in the interaction space that the PDA's screen creates. On the other hand, headphones can create an interaction space that includes only the person who wears them. Thus, the appropriate design of interaction spaces can minimize spillovers.

Similarly, the design of conventional (architectural) spaces has an impact on physical spillovers. Information that needs to be kept secret gets locked away in safes while locations such as cinemas (which we discuss in Chapter 9) restrict physical access to their screens by locking the doors and placing obstacles.

We now proceed to describe the different categories of spaces and interaction spaces. We should keep in mind that interaction spaces are created in physical spaces. Hence, what we say on spaces also applies to the respective interaction spaces. For instance, what we say on public spaces also applies to public interaction spaces. One important detail we need to keep in mind, however, is the following.

Different types of interaction spaces can be created in the same physical space. For instance, a private interaction space (created for example by a vibrating pager in one's pocket) can be created within a public space. To reason about this private interaction space, we need to apply our ideas from private spaces, not public spaces (although this private interaction space is within a public space).

5.2   Public, social and private spaces

The notions of spaces carry with them a great number of characteristics and understandings, which are peculiar to each society or social group. For each space or physical location, a characteristic set of behaviour rules may be defined [Barker, 1968]. Public spaces are places that belong to the community, i.e. a square may be a publicly owned space and thus is a public space. In such locations, for example, you would have the police, not private security. They are spaces where you can sit down and chat with a friend without having to buy something, in contrast for example to a privately owned café. On the other hand, private spaces are owned by someone, which can be used in whatever way the owner sees fit. Private spaces promote a sense of security and privacy. Public and private spaces are recognised by law precisely because of the importance that people attach to them - the physical locations as well as the notions, trends, rights, and common issues that pervade these locations.

For our purposes, the distinction between public, social and private has implications for access to that space. Our categorisation does not reflect how people behave in the space. Furthermore, the distinction between public, social and private spaces are not clear-cut. Public, social and private spaces are not simply defined by their geographical coordinates. Therefore, it is not helpful to try to categorise "pure" locations, i.e. "is a house a social space?", "is a park a public space?",

etc. It is up to the designer to decide in which category a space falls, as well as on the specific situation which may dynamically change.

Public spaces   Public spaces offer free access to everyone; i.e. a town square is a public space, carrying with it norms, expectations etc. concerning people's activities and behaviours in that space. The concept of public *interaction* spaces describes situations where interaction spaces exist in public spaces and provide unrestricted access, much like the big screens in New York's Times Square. An artefact that makes a resource available to everyone present in a public space generates a public interaction space. For instance, a large screen showing a football match in a public square creates a public interaction space.

Increasingly, the notion of public spaces has been squeezed and gradually colonised by the market, and has been commoditised. Now there are advertisements in public spaces, even in schools. In trying to define a normative framework, i.e. what should be the ideal, a public space is where information is provided without adver-

FIGURE 5.2: Public spaces.

Public spaces can be parks, streets, beaches, squares, or any other type of physical location that allows anyone to come in or come through. Such locations have strong expectations and norms that govern our behaviour.

tisements, and everybody has the ability to put information in, and take information out. The main issue is basically an exclusion issue: if some organisation can pay money to use an area, then people who cannot pay the money are excluded.

There currently exists the perception that we live in an increasingly divided society. The issue of privatised public space and its effects has been a key issue for urbanists, sociologists and geographers [Biddulph, 1993; Kayden, 2000; Putnam, 2000]. Such work highlights a shift in emphasis from the public to the social and private spheres, the increasing trend towards privatisation of the public realm, and the pulling apart of communities resulting in a "striking diminution of regular contacts with our friends and neighbours" [Putnam, 2000]. Public space where people from all walks of life can meet and interact is arguably more important than ever, therefore, for social exchange and the strengthening of community bonds, but is under siege in many places [Lees, 1994]. The observed withdrawal from public space (and the public sphere) and the higher degree of withdrawal in the more urbanized community has been attributed to the observed insufficiency of functions [Buchecker, 2003].

Private spaces   On the other hand, private spaces are spaces owned by an individual, which can be used in whatever way the owner sees fit. Private spaces promote a sense of security and privacy. Artefacts that are usable by only one person create private interaction spaces. For example, headphones create a private interaction space, even if the person who wears them is in a public space.

The concept of private spaces changes from society to society, as well as with the passage of time. For instance, a few centuries ago in Europe bedrooms did not have doors to separate them from the rest of the house. In contrast to this, today's bedrooms are seen as a private space, where one feels safe and has control over what happens. In fact, a number of teenager bedrooms have signs outside the door with

messages like "Go Away" or "Leave Me Alone". Although such a message can be seen as comic or trendy, they do highlight the fact that people have strong feelings about their bedrooms as private spaces.

Social spaces   Finally, social spaces are those spaces that are neither private nor public. Examples of such spaces are homes, cars, buildings etc. Groups of people interacting collaboratively with artefacts are usually within a social interaction space. Consider a group of children playing a board game in a public park. The interaction space that is created by the board game is a social interaction space; the children are included in it, but not everyone in the park is.

It is important to stress that public, social and private spaces are not simply defined by their geographical coordinates. Therefore, it is not helpful to try to categorise "pure" locations, i.e. "is a house a social space?", "is a park a public space?", etc. We want to point out that the criteria for categorising a location need to consider the

FIGURE 5.3: Private spaces.

Typical private spaces are bathrooms and bedrooms. However, any physical location where only one person is present and has control over this location should be thought of as a private space.

values attached to the location and the things that are happening there. Similar to the boundaries between information spheres described earlier, the boundaries between the 3 types of spaces we propose are also quite dynamic and fluid. A good example is "the park", which can be seen as a public space or a social space (e.g. a group of friends playing a board game in the park). Another example are benches on a street (a situation we examine in Chapter 9) where they are available to every-one (thus are a public space), yet a couple of friends exchanging information while sat in the benches might consider the bench as a social space. Finally, our categori-sation reflects access to the spaces, not how people behave or are expected to behave in public, social and private spaces.

5.3 Our approach to the design of pervasive computing

The essence of our approach to the design of pervasive computing is the *effective integration* of spaces (physical location + social dimensions) created by the built

---

FIGURE 5.4: Social spaces.

Typical social spaces are meeting rooms, night clubs, and living rooms. These are locations where several people are present, but access is restricted to the public due to physical or social constraints.

environment with interaction spaces created by computing resources distributed in that environment.

Experiences with human-computer interaction to date across a wide range of settings from aircraft flightdecks to the office environment have demonstrated the wisdom of playing to the respective strengths of humans and computers. Computers are good at storing and retrieving information, constant monitoring and other monotonous tasks, and performing complex calculations. Humans, on the other hand, are good at identifying patterns, spotting changes, extrapolating from knowledge and experience, and responding to new situations. Hence, human-computer systems at their best can be complementary.

With the coming of the pervasive technology paradigm, we are now interested in embedding computers in the environment. To understand how this can best be done we have focused on the relationship between the built environment and computer systems. Once again, playing to the strengths of each can result in an effective and complementary system.

On the one hand we have the build environment and architecture. One of the goals of architecture and urban design is to manipulate physical spaces in such a way as to provide greater functionality to people, and to allow them to do things quickly, effectively and with minimal obstacles [Bentley, 1985]. Thus, spaces are designed in such a way to facilitate people in what they try to do.

On the other hand we have the digital components (software and hardware) that we wish to embed in the environment. Many of the applications of computer systems may be viewed as complementary to architecture: performing complex calculations and data manipulation and exchange in order to overcome physical constraints. This is illustrated by the use of computer systems for communication and informa-

tion exchange across great distances, in effect redefining social spheres by creating new social interaction spaces.

In developing successful pervasive systems, these two complex designed systems, the built environment and computer systems, can benefit from each other by tackling each other's weaknesses and playing to each other's strengths. In operational terms, we claim that *architecture manipulates physical spaces, while computer systems manipulate interaction spaces*. We need to do this whilst minimizing the number of spillovers that occur both in spaces and interaction spaces.

5.4   Learning from architecture and urban design

The built physical environment is a system that almost everyone has used for the extent of their lifetime, and which has been studied for many millennia and through different eras. The pervasive presence of modern human-made physical artefacts like roads, bridges, paths, buildings and homes provides an example of a truly pervasive system. Architecture could be considered as a pervasive system or technology. It would be wise to study this well-established pervasive system, and make sure that the systems which we create - the systems that we want to be part of everyday life - do not conflict with or challenge architecture. As a first step, we summarise here some of the general approaches that architecture takes in designing public spaces. Like the virtual spaces projects we discussed in Chapter 2, the ideas and approaches described in this section can be adapted and applied to the design of pervasive systems.

At the most basic level, shelter from weather conditions and the degree of exposure to such natural elements plays an important role in the design of public spaces. Especially in northern countries, where the winds and temperatures can be threatening, people enjoy the shelter of surrounding buildings. The same applies to extremely hot climates.

A second important element in the design of public spaces is the degree of safety experienced by citizens. For instance, long narrow passageways and streets can be quite threatening unless enclosed by occupied buildings. In general, busy places tend to feel safe. Another potentially disconcerting experience is being lost in an unfamiliar city. Orientation therefore becomes crucial and is best served by recognizable and distinct features like monuments and landmarks. People may become uneasy in situations where streets intersect at odd angles or form an irregular network, as well as in tunnels, subways and underground passageways that contribute to loss of orientation and direction [Chapman, 1996].

In focusing on the aspects of place that give it its appeal, Logie [1954] proposes a number of "devices of urbanism" which are basic characteristics found in urban landscapes, either by accident or by conscious design. The significance of these devices lies in how they are perceived. For example, progressions are quite common, in the sense that streets are a type of progression. Another device of urbanism is the element of surprise, which can be important in avoiding the monotony caused by progressions, as well as creating psychological stimuli. Finally, contrast in form, colour and texture is common, as well as the contrast of scale between buildings and humans. Ideally, an area should be made up of a series of positive contrasting spaces that are clearly defined and unambiguous. These series may include static spaces, focal points where people may meet, undertake activities or just rest. Dynamic spaces are created by linear streets and alleys which act as linkages in the whole structure [Krier et al., 1979].

A number of further suggestions for urban design have been presented by Bentley [1985]. Some of the proposed key qualities are *variety* (the range of people, buildings and activities), *robustness* (the ability to change the use of space), *visual appropriateness* (the degree to which the visual cues make it clear what options are

available), *richness* (the range of sensory experiences like smell, lighting and touch), *personalization* (the ability of inhabitants to customise their environment), and *permeability* (privacy, and how many paths are available and how many of those are obvious). Similar ideas have also been considered by Collins et al. [1965] in a set of proposed principles for the relationships between buildings, public spaces and monuments.

A compact and practical set of guidelines for the design of public spaces has been proposed by the Project for Public Spaces (PPS) [2000]. They argue from empirical studies that successful public spaces are *accessible*, they cause people to engage in *activities* in them, they are *comfortable*, and finally they are *sociable* locations. These findings support our point that spaces are much more than physical locations. There is a wealth of activities, social interactions and social understandings present. A pervasive computing system available in a public, social or private space should promote all of these characteristics in order to function in harmony with the space itself.

Another effort at providing guidelines has been made by DETR[21] and CABE.[22] They suggest that designs for successful public spaces have a number of common elements.

- *Character*: Places with character have their own identity. They promote character in townscape and landscape by responding to and reinforcing locally distinctive patterns of development, landscape and culture. In terms of nationwide pervasive systems, should they be the same everywhere, or should they differ from country to country or city to city? Differences could be in terms of interface, services

---

21. UK Department of the Environment Transport and Regions
22. UK Commission for Architecture and the Built Environment

offered or back-end functionality. As we discuss in Chapter 9, the ultimate answer to this question can be given by local communities and authorities.

• *Continuity and enclosure*: These are places where public and private spaces are clearly distinguished. This helps promote the continuity of street frontages and the enclosure of space by development which clearly defines private and public areas. This directly relates to our discussion of private, social and public spaces, and once more highlights the importance of ensuring that the users are aware of where they are in both physical space and interaction space.

• *Quality of the public realm*: Places should have attractive and successful outdoor areas. Public spaces and routes that are attractive, safe, uncluttered and work effectively for all in society, including disabled and elderly people. How can a public pervasive system enhance the quality of the public realm? We examine this aspect in Chapter 9, where we discuss ways that a public pervasive system could do so.

• *Ease of movement*: Places should be easy to get to and move through. This promotes accessibility and local permeability by making places that connect with each other and are easy to move through, putting people before traffic and integrating land uses and transport. Pervasive systems can help citizens move through physical space by providing helpful information, as we discuss in Chapter 9.

• *Legibility*: Places should have a clear image and be easy to understand. Legibility can be achieved through development that provides recognisable routes, intersections and landmarks to help people find their way around. Landmarks have been used in virtual environments to assist users in navigation. We should be developing ways for pervasive systems to make use of landmarks.

• *Adaptability*: Adaptable places can change easily. Adaptability can be promoted through design that can respond to changing social, technological and economic

conditions. We have already highlighted the need to pay attention to these conditions in the design of pervasive systems. If places should be able to change easily, then the pervasive systems embedded in these places should also be able to change easily.

- *Diversity*: Diverse places offer variety and choice. Diversity and choice can be promoted through a mix of compatible designs and uses that work together to create viable places that respond to local needs. Once again, we should be thinking of ways for pervasive systems to help provide wider diversity and more choices to people in public spaces.

Oldenburg's concept of the "third place" [Oldenburg, 1999] is also relevant here. Third places are "social condensers" which enlarge and reinforce civil society by providing a space in which people of a community can meet to interact with others. A successful third place has a number of important characteristics. They are *free or inexpensive to enter* and to purchase food or drink. They are highly *accessible* so that ideally people can reach them comfortably on foot. They should have "regulars" expected on a daily basis. Everybody should feel *welcome* and it should be easy to get into conversation. They should be *inclusive* without formal criteria of membership or exclusion and unconcerned with status.

Finally, we must consider how the ideas we have presented may be implemented in a variety of societies, ranging from western modernized societies to traditional village societies to far eastern societies. Echoing Alexander [1975; 1977], we argue for the crucial role of local expertise and participation when it comes to deciding what is appropriate in terms of designs and buildings, given that space has a social logic to it and that what is thought of as an appropriate structure is influenced by the structure of society [Hillier, 1984]. However, according to Chapman [1996], "Inevitably, our origins in Western Europe precondition both our values and aspi-

Christopher Alexander

Christopher Alexander is an architect and professor emeritus at Berkeley. He is very interested in design, and computer scientists who read his books are impressed by the parallels with designing software. Christopher Alexander was born in Vienna, Austria in 1936. He graduated with degrees in mathematics and architecture from Cambridge University and with a PhD in Architecture from Harvard University. For his doctoral dissertation, Alexander developed a computer program that attempted to analyse and create new environments based on logical programmatic analysis. This interest in creating new environments would mark all of his future work. Eventually his confidence in mathematical methods as a basis for better design declined and he utilized empirical research to create patterns. Disenchanted with computer-driven design, but more than ever interested in what made certain places work both spatially and psychologically, Alexander developed a theory of "fit" in terms of what he called "patterns". This theory suggested a means for creating successful places that blended the application of logic with collective experience.

Pattern theory inspired many, but also failed to lead consistently to beautiful buildings. In the late 1980's Alexander started to develop a further theoretical basis for good design based on a careful definition of "wholeness", or a kind of deep and abiding beauty.

Although most of his buildings have effectively supported his theories, Alexander has mainly influenced the architectural profession through his writings and teaching rather than through his completed buildings. Due to a softening in his stance, his critics now accuse him of embracing ornamentation and craft at the expense of modern technology.

(Source: http://www.greatbuildings.com/architects/Christopher_Alexander.html)

rations. Nevertheless […] the principles and ideas we have discussed are applicable

to cities and villages world-wide". (p.153).

5.5   Implications for the design of pervasive computing systems

The ideas and implications for pervasive computing systems presented in this section fall into two general categories. This is a consequence of viewing pervasive systems from two different perspectives. First, a pervasive system may be viewed as a *functioning set of digital artefacts*. These digital artefacts, much like their physical counterparts [Brebner, 1982], have a huge impact on the way a space is used and perceived, and on the results and effects it has on people. In designing these artefacts, both digital and physical, we must anticipate their effects and try to tailor them and direct them according to our aspirations and goals. Secondly, a pervasive system may also be viewed as an *invisible part of or extension to* the physical environ-

ment. As such, a pervasive system must encompass our aspirations and goals both when viewed as part of the physical environment and when considered on its own.

The result of this duality of views may be illustrated by an example: successful public spaces, as noted above, offer comfort and security to people. The implication for the design of a pervasive computing system is that it should enhance the safety and security provided by the public space, while at the same time it should itself be safe and secure. This is similar to the notion that, for example, benches should be placed safely within a public space (e.g. not obstructing cyclists), but at the same time benches should be designed and built with safety in mind (e.g. no threatening corners, solid material, non-flammable).

A number of further design ideas and implications for pervasive (computing) systems may be drawn from the architectural design ideas presented in the previous section. As noted by the Project for Public Spaces [2000], successful public spaces are accessible, they allow people to engage in activities, they are comfortable, and they are sociable. From these four key characteristics follows a number of issues that should be considered in the design of public pervasive computing systems. In terms of allowing easy access, we should consider how the presence of the pervasive system is made "visible" or somehow manifested, so that people both in and outside the public space are aware of its existence. An example of how overlooking this issue can cause problems is the installation of wireless network access points in public parks. Initially people could not easily know if a location had wireless coverage or not. To overcome this, the installation of public wireless networks is often accompanied by the installation of signs and signposts to inform people of the presence of a wireless network. Although simplistic, signposts are better than nothing. However, we need to look into more efficient and accurate ways of manifesting the presence of a pervasive system both for the people in it and those outside it. This

becomes even more important in light of the popular view that pervasive technology should also be invisible [Weiser, 1991].

Easy access

People should enjoy easy access to a pervasive system. The first step in providing easy access is to allow for the easy recognition and identification of the system. The next step is to allow easy access both in terms of connecting or getting access to the system as well as using the system. The absolute minimum requirements should be expected of the users, and artificial requirements such as having a certain height, weight, age, special equipment or even special knowledge should be avoided. Conventional technology is a good place to look for examples. Public parks usually have water fountains which allow users to walk up to them and use them - no need for special equipment such as a cup or bottle, and no need for the users to intervene and fine-tune the system. In Chapter 7 we show how this requirement for easy access has affected the implementation of an interaction technique we have developed (see page 155).

Comfort

Pervasive computing systems should also enhance and augment the comfort provided by a public space. This means that any sensory, e.g. visual or auditory, manifestation of the pervasive system should be appealing to the owners and users of the public space, i.e. the public. Mechanical and electrical equipment traditionally is hidden in all but radical architecture and this is reflected also in conventional HCI notions of designing the user interface as an independent layer that floats serenely above the hidden maelstrom of code and network protocols and routers that provide the functionality of a system. However, we should also consider situations where the physical manifestation of the working of a pervasive system could assist in the learning curve of those using it. For example, the presence of cables could indicate the presence of the systems, or noise generated by the infrastructure equipment could indicate that the system is operational. In terms of wireless networks,

the base-stations providing access to the network could become physical markers denoting the presence of a network (instead of hiding them and installing signs).

The role of infrastructure

The importance of infrastructure was well demonstrated in the "Can You See Me Now" (CYSMN) game [Benford et al., 2003], where on-line players using the Internet were chased across a map of a city by runners who were moving through the real city streets, tracked by GPS and connected to the game by 802.11b wireless networking. It quickly became apparent that there were infrastructure problems such as GPS inaccuracy in tracking the runners, patchy wireless network coverage and frequent technical failures of components, cables, batteries etc. At first the runners suffered from these failings. Within a day, however, they had begun to develop their own models of the infrastructure and had learned to exploit the inaccuracies and idiosyncrasies of the system. For example, the runners developed tactics of lurking in GPS shadows and moving relative to the edges of wireless network coverage. This experience reinforces the view that infrastructure is often perceived by users and has effects on how a system is used.

Physical form, orientation and surprise

Successful public spaces attract all age groups of both genders, and this is something to which pervasive systems in public spaces should aspire [PPS, 2000]. In addition, a clearly represented, and in some cases manipulable, level of security should be provided by the public pervasive system, so that the public do not feel threatened or alienated by it. Also, a basic sense of orientation should be provided and supported by the system as a means of further enhancing the comfort and sense of security. Remembering that in creating public pervasive systems we are designing user experience for members of the public, the element of surprise could be considered as a way of stimulating people who navigate the available spaces, both physical and digital. The element of surprise and ambiguity in general has been proposed as a design resource [Gaver et al., 2003]. To follow from the previous examples, not all

cables of a pervasive system need to be visible (some areas could be wired without any indication - thus offering a surprise). This can help avoid the monotony of progressions (streets in the physical environment, interaction spaces in the digital sense).

Activities

Activities are a basic characteristic of public spaces. The fact that there are things to do gives people a reason to visit. In integrating with and augmenting the physical space, a pervasive computing system can improve the experience of visitors by enabling its users to engage in activities, including group activities. We therefore need to design systems that support social interactions. Currently, benches and seats in public areas are placed in such a way to foster conversations between people, the formation of new friendships, and socialisation in general. Similarly, digital artefacts should be designed and deployed to foster and encourage such social interactions.

Support

Furthermore, because the success of group activities can be affected by how well the activities are being supported, it would be helpful for people to be aware of someone who is available to help or someone who is there to facilitate and assist in the activities. Much like public utilities have specialized personnel for various types of support (customer service, hardware problems), pervasive systems could employ similar support to their advantage.

Interaction spaces

We have described how an artefact such as an interactive computing resource defines an interaction space within which a person's activities are supported by the artefact. These interaction spaces may be public, social or private. They are created within and combine with the public, social and private spaces defined by the surrounding architecture to form a context for the person's activities.

Implications for the design of public pervasive systems include the ability for a user, given the resources available in a particular setting, to define interaction spaces suited to her moment to moment purposes and activities. Hence, we need interaction techniques that support, for example, the definition of a private interaction space at one moment and a smooth transition to defining a social or public interaction space in the next moment, all while the person is physically located in, for example, a public space. Additional requirements include the ability to use common interaction techniques across a very wide range of devices with varying physical characteristics, thus freeing the interaction technique from the physical form of the system. In Chapter 7 we do just that, by describing an interaction technique we have developed, and showing how the design implications we have just discussed have influenced the creation of this interaction technique.

5.6 Summary

The third and final part of our framework, space, highlights the importance of space in the design of pervasive systems. Here, we provided a classification of spaces and interaction spaces into private, social and public, reflecting also our classification of information spheres. We have also shown how ideas from urban design and architecture can be applied to the design of pervasive systems. Finally, we described our high-level approach to the design of pervasive systems, which consists of effectively integrating spaces (manipulated by architecture) and interaction spaces (manipulated by computer systems).

We have now provided a description of the three aspects of our framework: citizen, sphere and space. Next, in Chapter 6, we discuss our framework in full, and present our design tool which embraces the three elements of our framework.

A DESIGN TOOL FOR PERVASIVE COMPUTING

In this chapter we build on the framework presented in Chapters 3, 4 and 5. The framework identifies three key elements that can be used to understand and to consider requirements for truly pervasive systems. In this chapter we describe how we have operationalised the framework to provide a tool for decision making for designers of pervasive systems. In Chapter 8 we illustrate the use of this design tool, applying it in the real world context of a hospital Accident and Emergency department.

Having described each of the three elements of our framework in detail, we can now begin to see how each element interacts with each other. We can also develop our approach to designing pervasive systems and build on our analysis so far.

In Chapter 5 we discussed research within Architecture that builds on architecture and urban design in order to design better virtual environments. This body of knowledge is also useful for the design of pervasive systems, and that is why our space element builds on this work. However, space does not exist without people in it, and any understanding of space is incomplete without an understanding of the people in it. For our citizens element we have proposed that citizenship rights should be taken into account, and by doing so we arrived at a view of pervasive systems based on a public service model. Furthermore, we argue that the design and provision of a wide-spread pervasive system should be coupled with social responsibility. We believe that our framework is well suited to coping with important contemporary social issues of space, and the effects on people and society.

The ultimate objective for pervasive systems can become clearer by considering some analogies. The designers of a payroll system seek to improve the payment process for both employers and employees. The designers of a patient database seek to improve the provision of health services. The designers of train scheduling system seek to improve the perceived quality of train service as well as to optimise its operation. What should the designers of public pervasive systems as we envision them seek? In Chapters 3, 4 and 5 we described a number of issues on space, spheres and citizenship, and problems that have been raised in domains outside Computer Science and Human-Computer Interaction. Yet, these problems become our problems when we seek to deploy our computer systems throughout society, in every part of our daily life, in every object we use daily.

The three elements of our framework lie at the heart of some of the issues that could potentially become the goals that public pervasive systems seek to address or accomplish. To successfully build pervasive systems then, we need to understand ourselves, how we behave in our societies, what we expect and how we seek to improve our lives. We also need to understand the built environment in which we have chosen to live, and how we seek to have it improve our daily lives in many ways including shelter, comfort, sociability, and even fashion. Moreover, we must understand the notion of the public realm or the public sphere. Therefore, to the extent that pervasive systems can potentially improve our lives on a daily basis, we propose that the three elements of our framework need not only to be understood, but also operationalised in terms of design.

6.1   From framework to design tool

Our framework leads us to consider pervasive systems in terms of three key elements: citizens, spheres and spaces. Spaces, in turn, are separated into space and interaction space. In building an applied science of HCI, we wish to go beyond the theoretical base and to operationalise the framework in a form that designers can readily use. The design tool produced by this operationalisation is presented in the initial diagram in Figure 6.1 and explicated in our description of the process of examining this diagram and proposing changes. We discuss this process in Section 6.3. As an aid to visualising the elements of our design tool, Figure 6.2 instantiates Figure 6.1 with photographic examples.

In the following sections, we use the initial version of our diagram to introduce the notion of *connectors* (the lines that connect one point of the diagram with another) and explain what they mean. We then explain how a pervasive system can be represented using *instances of the diagram* and conversely how our design tool can be used to describe a pervasive system that we wish to design. At that stage it should

be clear how to read a diagram and how to create a diagram. With this understanding, we proceed to enhance the semantics of the diagram in Figure 6.1, in such a way that it conveys more information visually. To do this, we discuss the connectors. We do so for all three groups of connectors - one group between each consecutive pair of columns. We describe every connector, give examples of what it represents, and explain some of the issues that arise with the presence of each specific connector in an instantiation of the diagram. We then visually code each connector according to its nature, so that designers (or anyone else who wishes to deal with such diagrams) can easily recall the semantics of connectors simply by looking at the diagram.

The diagram   Looking at the initial version of our diagram in Figure 6.1, we can see that there are four columns. The first two columns represent spaces and interaction spaces,

FIGURE 6.1: An initial version of our design tool diagram.   The elements of designing large scale pervasive systems, shown with their relationships.
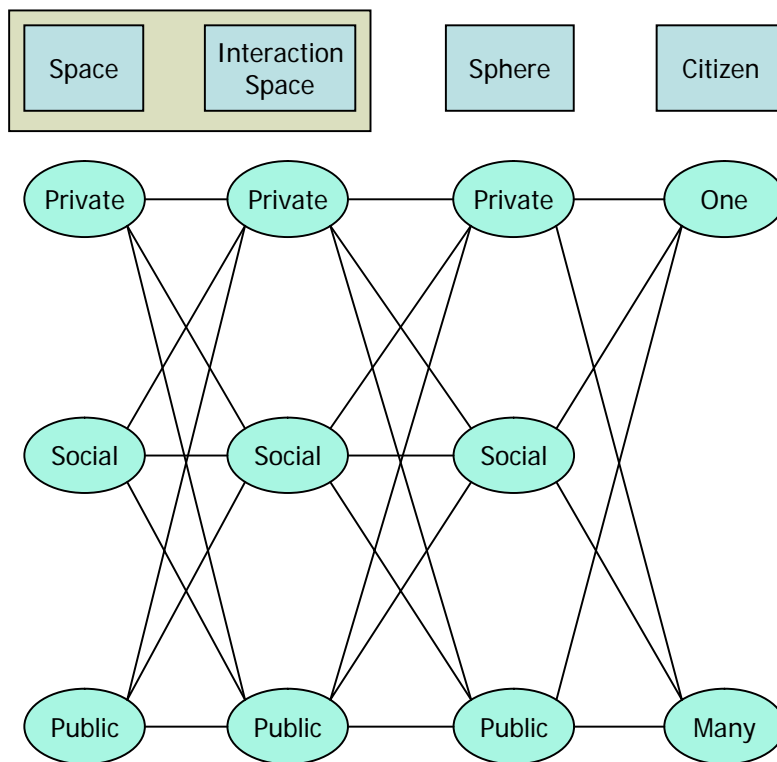
FIGURE 6.2: Our
design tool diagram
shown with examples.
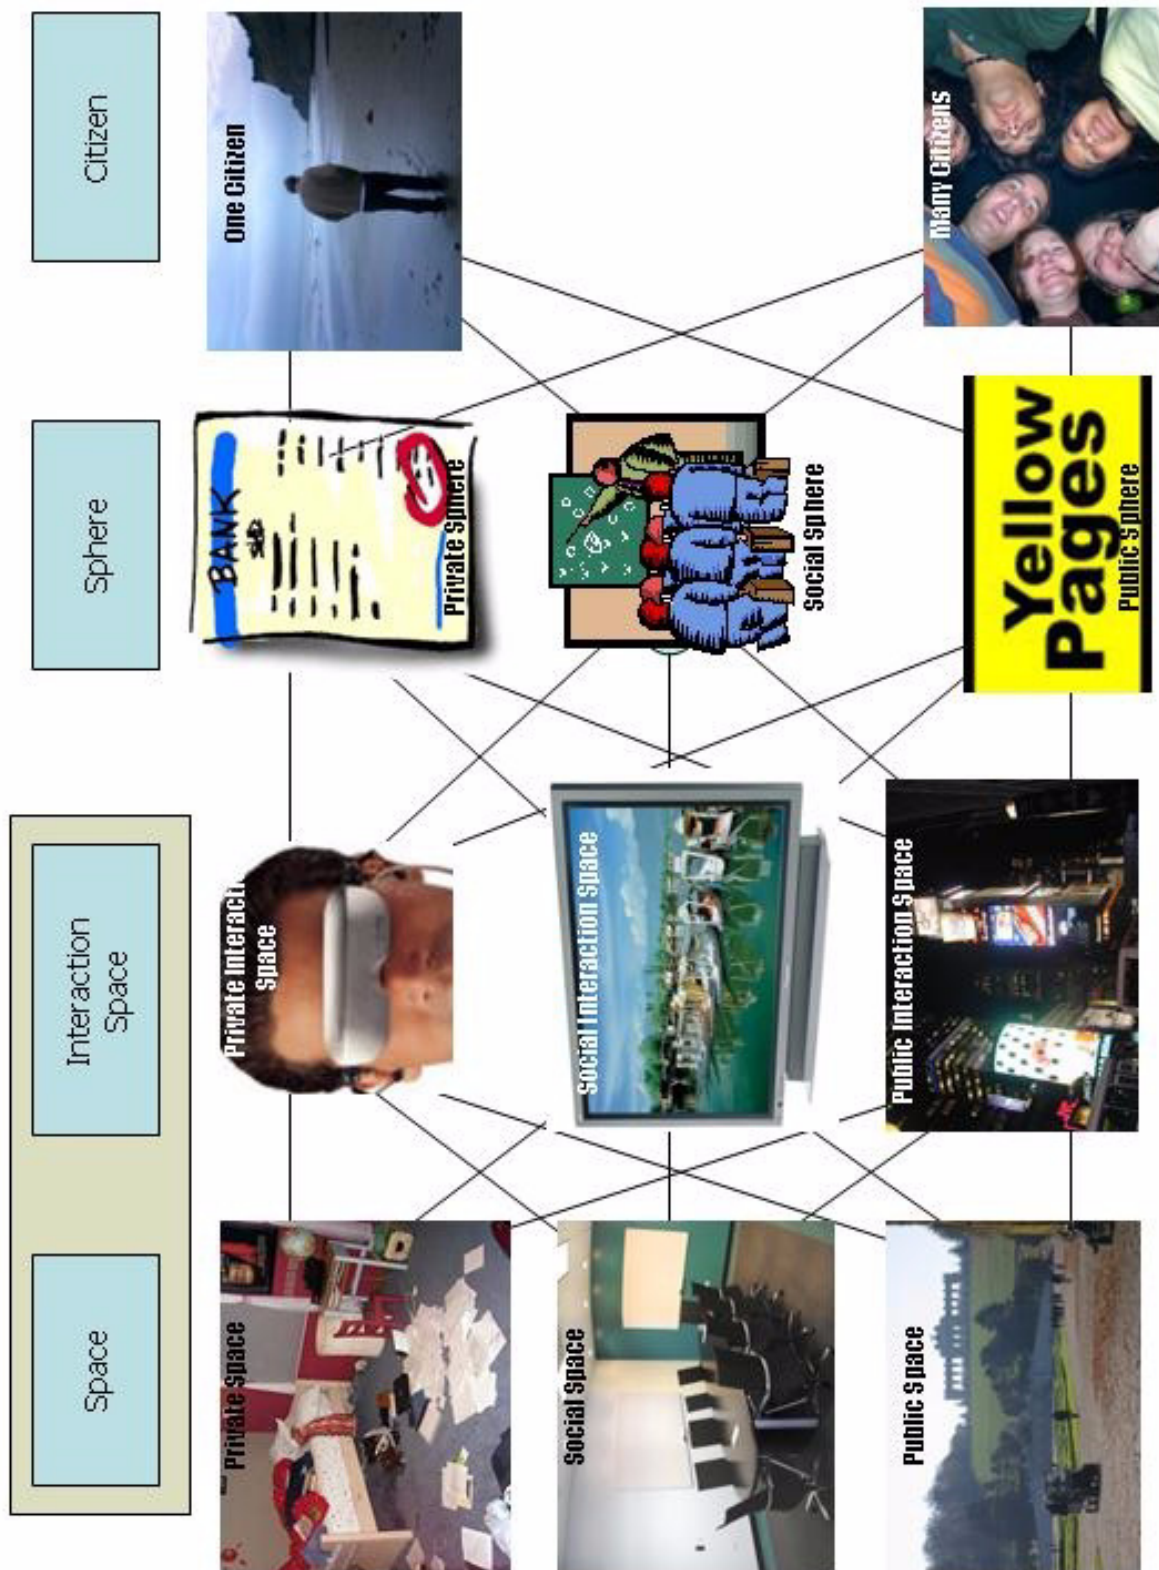
which combine to form the spaces element of our framework. This element of our framework was discussed in Chapter 5. The third column represents the spheres element, which we discussed in Chapter 4. The rightmost column represents the citizen element of our framework, which we discussed in Chapter 3.

Each of the first three columns has three rows (or points), which are shown as ellipses. These reflect the classification we have used throughout the description of our framework: private, social, and public. The citizen column consists of only two points: one and many. These represent the presence of one person or more than one person, either within a space or an interaction space. This column is placed next to the spheres column to indicate that in both spaces and interaction spaces, citizens access information spheres.

Our design tool can be used to examine existing settings and artefacts or to evaluate envisioned designs and proposed artefacts. In either case, using the design tool will help the designer to decide which devices and technologies are suitable for delivering which pieces of information to particular citizens in particular settings. Furthermore, we can reason about mixing traditional technologies (such as paper-based posters, people speaking to each other, etc) with digital technologies, an approach that should not be dismissed, since many traditional artefacts are much better at what they do than their digital counterparts. As we will also see, the design tool can identify potential spillovers (see page 84), situations where physical interaction is not possible, as well as situations with potential cognitive overload for those using the system. For now, let us explain how a pervasive system can be represented using our diagram.

Let us consider an imaginary pervasive system, call it PerfEx, and let us describe how it can be represented using our diagram. What we propose to do is to study a specific aspect of PerfEx that is of interest to us. Let us say that this system enables

a user to access *train timetables* from the *privacy of their home*. This information is delivered by means of a *small personal device*. This device also allows *many users* to engage in *various group discussions* about the train timetables and quality of service.

In the above simplistic description of PerfEx we can find encoded the following features: public sphere, private space, private interaction space, social sphere, one citizen and many citizens. This information is encoded in the emphasised terms. We could rewrite the description of the pervasive system by substituting those emphasised terms by our own decoded terms, as follows: PerfEx enables *one citizen* to access the *public sphere* from *private spaces*. The information is delivered by means of a *private interaction space*. The interaction space also allows *many citizens* to access various *social spheres*.

We have transformed the original description of PerfEx into a more generic description using the terms of our framework for the following reason: this new description is the textual equivalent of the diagram instance shown in Figure 6.3.[23] Both the textual description and this diagram instance describe PerfEx. The connectors between each of the points mean that these aspects are related because of the setup of the system. For example, because the public sphere is accessed from a private interaction space, there is a connector between these two points in the diagram instance.

A first point to note is that the description of PerfEx using our own terms is generic. This means that slightly different systems can have the same diagram instance and textual description using our framework. For instance, if in the original description we changed "privacy of home" to "privacy of toilet", not much would change in Figure 6.3. Also, if we changed "small personal device" to "earphones", again not

---

23. We have simplified the process of creating a diagram instance for illustration purposes. For a complete account of this process, see Section 6.5.

FIGURE 6.3: The diagram instance representing PerfEx.

A diagram instance of our design tool contains points which are used by the system, and connects those points which are used together and points that are related because of the system setup.
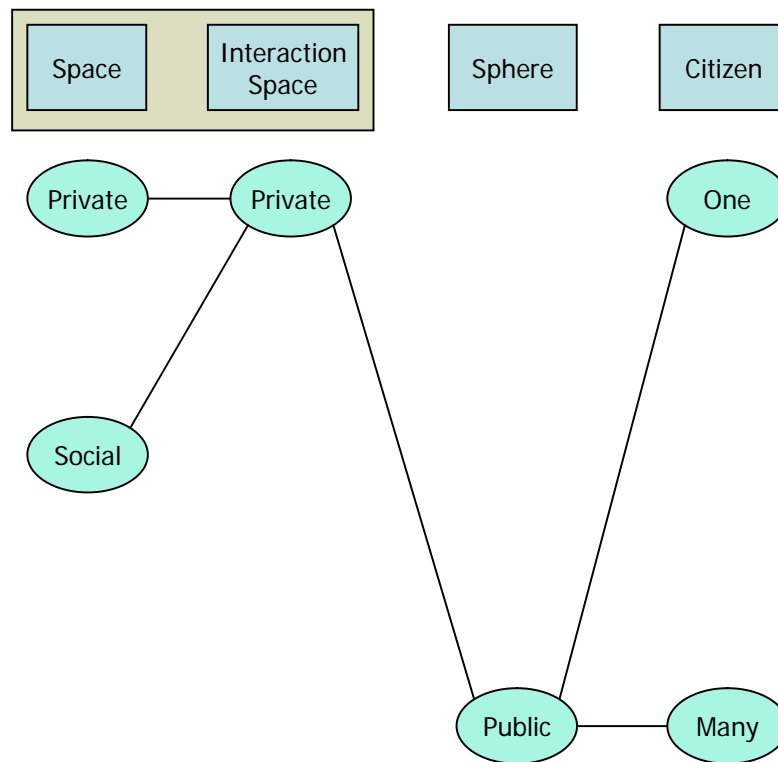
much would change. The reason behind this is that our diagram distils what we know about the pervasive system, and only shows the information and aspects of the pervasive system that are reflected in our framework.

What this means is that going *from* a diagram *to* a specific system we have many choices and options. Therefore, to specify a pervasive system using our diagram and then try to actually build it, we have to transform terms like "private interaction space" into particular technology, or "social space" into physical spaces. There are many possible ways of doing so.

We believe that this polytypic nature of our diagram is a strength of our design tool. It can encode various possibilities into one diagram, but can also convey the important characteristics of a pervasive system. It is not too restrictive and allows for various courses of action and the exploration of slightly different implementations from the same specification. But also it is not too vague, as it encodes a lot of infor-

mation in the diagram instances. Each of the points on the diagram carries a lot of meaning, as we have discussed in Chapters 3, 4 and 5. But also, a lot of meaning may be carried by the connectors between specific points on the diagram. The following section discusses the information that can be conveyed by the connectors.

6.2   The connectors

So far, our diagram has consisted of points and connectors, and we have not distinguished between different connectors. We now do so, by using dotted lines to indicate connectors where the designer of a pervasive system must be particularly careful, where conflicts are likely to arise between the demands and affordances of the different elements of our framework and where particular activities or information access may not be supportable in particular settings by the available technologies. To identify such situations, we traverse through all the connectors and present a short discussion about each one.

**Spaces &
Interaction
Spaces**

Regarding the connectors between spaces and interaction spaces, it is important to note that interaction spaces are ultimately *bound by the physical and social constraints* of spaces. However, by exploiting appropriate technologies, interaction spaces can span various types of spaces. For instance, a public interaction space generated by sound from loudspeakers can span private, social and public spaces. This, however, *does not support direct physical interaction and manipulation* by the participants in the interaction space. The loudspeaker creates a one-way channel that can indeed bring the public interaction space defined by the loudspeaker into, for example, an individual's private space but turns the individual into a passive recipient of the information being broadcast. Similarly, a social interaction space could span private spaces. For instance, a chatroom could be accessed by participants who are in their own private spaces - such as in their bedrooms. Again, we see that there is a

lack of direct physical interaction, since by definition only one person can exist in the private space, but many have access to the social interaction space.

Therefore, we represent the downward sloping connectors (reading Figure 6.1 from left to right) between spaces and interaction spaces with *dotted lines* to highlight the *caution* required when attempting to design a pervasive technology to bring a public interaction space into social and private spaces, as well as social interaction spaces to private spaces. The designer should make sure that if these connectors are present in the specification of a pervasive system, then physical interaction should not be a requirement for the task at hand.

The more upward the slope is (again reading Figure 6.1 from left to right) between spaces and interaction spaces, the more care needs to be taken not to *intrude on privacy.* Here, the technology must be what we term "insulating". A good example of this is the link between public space and private interaction space. In this case, the technology must ensure that the freedoms and norms of the public space are not carried into the interaction space, thus jeopardising privacy by means of spillovers. For example, headphones are good at creating private interaction spaces within public spaces. Also, devices such as phones and PDAs with small screens, despite the criticism they sometimes receive for being unusable [e.g. Kostakos & O'Neill, 2003], well serve the purpose of allowing an individual to interact with them within a private interaction space from which other people are excluded.

Another example is the connector between public space and social interaction space. This connector is exemplified by holding a group discussion in a public space. The technology to support this should be insulating, and avoid spillovers of information. As we have said, interaction spaces are ultimately bound by the physical and social constraints of the physical spaces they span. So in this last example, the good old-fashioned technology of "speech" could suffice, since physical con-

straints do not allow speech to travel very far, and thus a group can have a discussion in a public space with relative privacy.

The least problematic connectors between spaces and interaction spaces are the horizontal ones. Private spaces to private interaction spaces can support direct physical interaction, but at the cost of *not being able to span multiple spaces*. Social spaces to social interaction spaces can have a many-to-many relationship. Thus more than one social interaction space (or other types of interaction spaces as well) can be present within a social space, but if the various interaction spaces are unrelated in terms of task, sphere, or intentions, then the participants will be broken up into groups, thereby reducing interaction amongst them. Public spaces to public interaction spaces are exemplified by shared public spectacles, such as an open street performance. In such cases, *all participants need to be included* by the interaction space, something which can be quite a challenge for the technology being used.

We also note that a space can have more than one interaction space in it. We have already noted this for the social space to social interaction space connector. This could potentially create *cognitive overload* for those within the spaces. The same holds for private and to a smaller degree public spaces. In the case of public spaces it is unusual that we are required to get the attention of *everyone*.

Interaction Spaces & Spheres

In the connectors between interaction spaces and spheres, upward slopes denote insecure design options, i.e. design options that have the *potential to undermine the privacy or security* of information because of spillovers. For example, a public interaction space being used to access private sphere or social sphere information is insecure. Doing this is analogous to showing the contents of one's wallet on a wall display, or posting the contents of one's bank account on a public web site. Using social interaction spaces to access the private sphere raises the same issues. An example of this would be to make available someone's private information in a

social space such as an office. However, the restrictions causing the information to be social could be of minor importance, and as such privacy and security are not undermined.

Private interaction spaces offer the least problematic connectors, but at the cost of physical interaction: by definition, a private interaction space has only one person in it. Therefore, besides being able to access the private sphere in a secure manner, accessing a social sphere from a private interaction space takes a toll on the possibilities for physical interaction either within the participants of the social sphere or with any resources and artefacts that are being used for the task. The same holds for accessing the public sphere from a private interaction space, in which case it can be hard or impossible for all those accessing the public sphere to have direct physical interaction with each other. This is problem for services and functionality where this interaction is required.

We also highlight the one-to-one relationship between social interaction spaces and social spheres. If more than one sphere is accessed (thus breaking the one-to-one relationship), there is a risk of fragmenting the resources required to sustain the activity at hand. Some activities require access to various sources of information and this is not precluded, but designers need to be cautious whenever more than one social sphere is allocated to a social interaction space. The cocktail party effect describes this [Aoki et al., 2003], where it can be very difficult for someone to follow or participate in more than one social spheres which simultaneously exist in the same social space.

Finally, when public sphere information is accessed from non-public spaces there should be a *consensus of common interest*. For example, public information shown within a movie theatre should be of relevance to those present. This is also the case when accessing private information within a social interaction space. We have

noted that this is not particularly secure, but it is a way of sharing private information, effectively making it non-private. In such cases there should be a common interest or common activity in order to avoid situations where unwanted information is provided uncontrollably.

Spheres & Citizens

Our design tool also has connectors between citizens and spheres, in order to encourage the designer to explore how people actually look for information and contribute to the information spheres. Once private information has been made social or public, it is hard - even impossible - to reverse that movement. In terms of the diagram in Figure 6.1, information only flows downwards, hardly ever upwards.

Having many citizens access a private sphere raises privacy concerns, since by definition only one person should be granted access to a private sphere. Additionally, a citizen can have access to many social spheres simultaneously but will generally attend to one at a time. The typical example for this is again the cocktail party effect, where it can be very difficult to tune into more than one discussion simultaneously. Whilst accessing more than one sphere, citizens can transfer information between them, however this transfer may only be downwards.

Citizens are physically present within physical spaces. Therefore, we could conceptualise connectors going from the citizen column to the space column. Ultimately, however, citizens access information spheres, so rather than having to visualise this by tracing the connectors from citizens to spaces to interaction spaces to spheres, we directly represent how citizens access spheres. Furthermore, we have said that interaction spaces are bound by the physical and social constraints of spaces. A factor affecting and contributing to these constraints is the number of citizens present within the spaces. Therefore we can consider the presence of citizens within spaces and its effects.

Figure 6.4: The
final version of our
design tool diagram.

This is the final version of our design tool diagram. The dotted connectors indicate that special attention needs to be given. For a more detailed explanation of the dotted connectors, see Section 6.2.
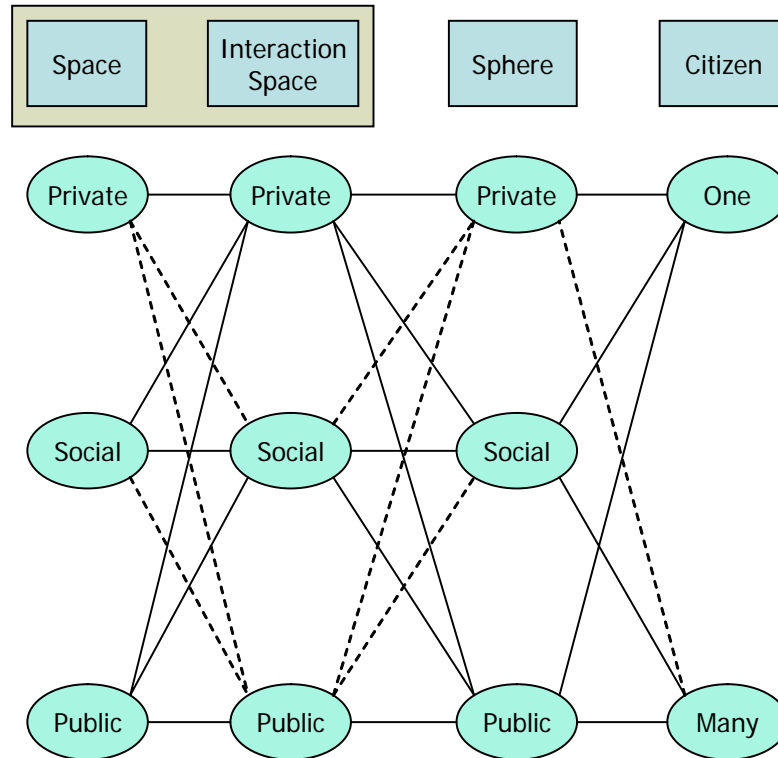
In this section we have described in detail the semantics of the connectors used in our design tool diagram, and have extended the diagram of Figure 6.1 in order visually to convey the details we have discussed in this section. The resulting diagram is shown in Figure 6.4. This the final version our diagram, which we can instantiate as part of the design process, is our design tool. Next, we explain a method for using our design tool visually to inspect the design of a pervasive system. Our method consists of four steps, each of which we discuss over the next four sections.

6.3   Step 1: Generate a list of artefacts

Our design tool is used in four steps, although more iterations are possible and encouraged. By the end of every cycle, the designer should have new instances of our design tool diagram. These diagram instances can then be re-evaluated until a satisfactory system design has been reached.

The first step in our method (❶) is to generate a list of "artefacts" that are of interest to us for the purposes of the specific application and real world setting for which we are designing. This list should contain artefacts that fall within three general categories:

- locations;
- technologies;
- information.

Locations

Locations are all those artefacts that refer to physical spaces. These can be rooms, buildings, halls, corridors, homes, parks and generally any physical location that is of interest to us. The locations that are of interest to us are those where our technology will be installed and used, as well as those locations where citizens will be.

Technologies

Technologies refer to all those artefacts that create interaction spaces. For instance, posters, display screens, flyers, headphones and whiteboards all fall into this category. It is important to note that these artefacts are not limited to digital or electronic technologies. The technologies that we should be interested in are those that will be used to deliver and provide access to the information and resources that are relevant to our system. Although one could argue that almost any artefact in the physical environment creates an interaction space, we only need to focus on those artefacts that are actually part of our system.

Information

Finally, the information category includes all the information that is identified as relevant to the specific design situation. Depending on the setting, this can be a very diverse category, ranging from personal alerts and phone books to general announcements about train times, hospital waiting times, or even advertisements. The grouping of the information resources should be done based on the category of information (public, social, private).

6.4 Step 2: Produce a check-list

Having produced the categorised list of artefacts, the next step (❷) is to produce a check-list for each artefact. This check-list should include those points (represented by the ellipses) on our design tool that are associated with (for example, generated by) the artefact.

Items falling under the locations category produce checkpoints under the space and citizen columns. This is because physical locations are only capable of generating spaces, but as we have also discussed they act as containers where people are present.

Next, technologies produce checkpoints under the interaction space and citizen columns. The reason, as we have previously discussed, is that interaction spaces are generated by technological artefacts, and can include or exclude people.

Finally, information artefacts produce checkpoints under the spheres column. This is quite straightforward, as information artefacts will belong in one of the three types of spheres.

To better understand this process, let us consider the PerfEx system we described on page 120, and generate a list of artefacts along with their checkpoints. The main artefacts we can draw out from our initial description of PerfEx are: the train timetables, home, small personal device, and group discussions. The list of artefacts we have generated, along with the checkmarks for each artefact, are presented in Table 6.1. In this table, we see that the train timetables artefact belongs to the public sphere. Also, we have identified home as a location that offers both social and private spaces, depending on the situation, and which can have either one or many citizens in it. The personal device creates a private interaction space which is used by only one person. Finally, the group discussions belong in the public sphere.

6.5   Step 3: Create diagram instances

With a list of artefacts and their associated points on our design tool diagram, we next have to identify the various subgroups of artefacts that exist in our list (❸). Our goal then is to generate an instantiation of our design tool diagram for every meaningful group of locations, technology and information that are related in the real world. The way meaningful groups are defined depends on the situation we are designing for and the aspects of the system we are exploring.

We can achieve this by creating groups of related artefacts within our list. These groups define ways in which artefacts are related to each other. For instance, the technology that delivers information, the information itself, and the location where the technology is installed can belong in the same group. Depending on the level of resolution at which we are studying our system, we should get one or more instances of our diagram. If, for example, we are studying a very specific aspect of our system, it is quite likely that all the artefacts in our list are related, and thus we generate only one diagram instance. On the other hand, we could have a more complex situation where the artefacts we have listed are not all related, but rather form "sub-groups". In this case we could have two or more diagram instances to describe the pervasive system we are studying, as shown in Figure 6.5. Also, it

TABLE 6.1: An initial checklist of artefacts for PerfEx.

For each artefact we have identified in our list, we add a check mark under the categories of Space, Interaction Space, Sphere and Citizen which the artefact reflects. Note that 'P' means public, 'S' means social, 'Pr' means private, 'M' means many, and '1' means one.

Artefact	Space			Interact. Space			Sphere			Citizen	
	P	S	Pr	P	S	Pr	P	S	Pr	M	1
train timetables							✔				
home		✔	✔							✔	✔
personal device						✔					✔
group discussions							✔				

FIGURE 6.5: A list of artefacts could generate many subgroups.

A list of artefacts can generate two or more groups. In this figure, we have identified two groups of related artefacts. Note that artefacts can be shared between groups, such as the "travel times" information artefact in this example.
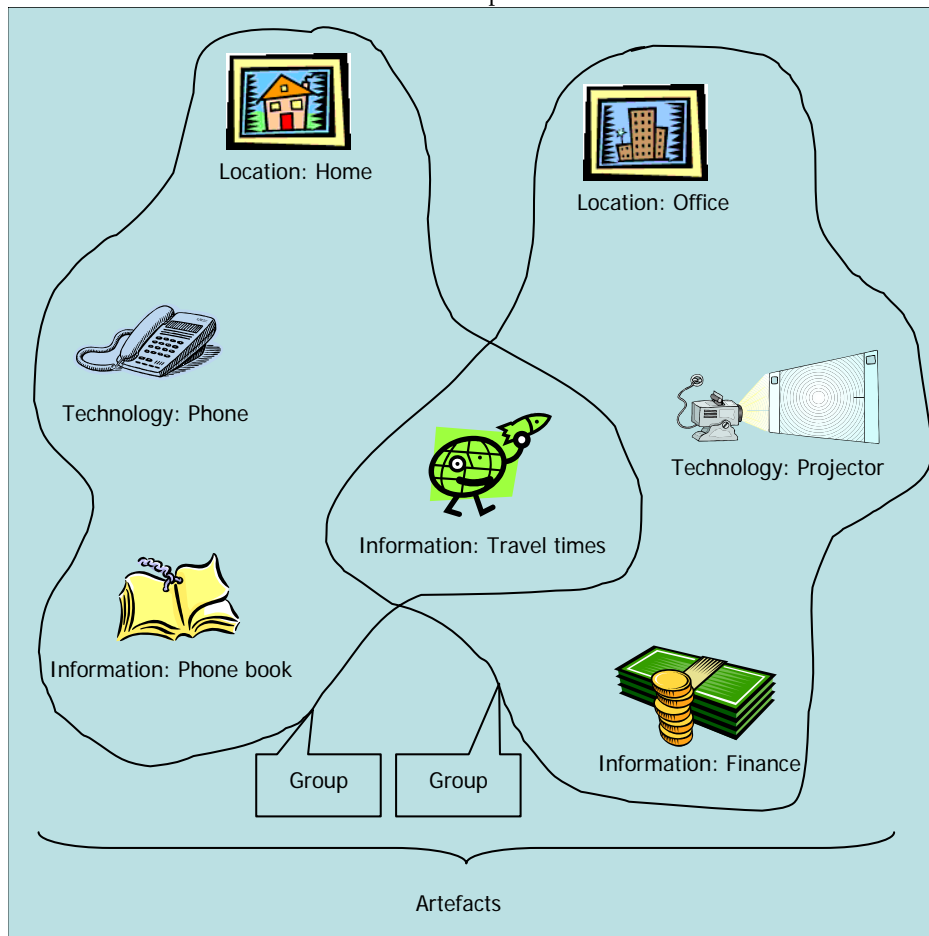
could be the case that a specific artefact is present in more than one group. Finally, in situations where more than one instances of a particular type of sphere is being accessed (for instance more than one social sphere is being accessed) the designer may wish to separate each of those instances into a separate diagram depicting the sphere and the interaction space used to access the particular sphere.

We now have one or more groups of artefacts, each of which has at least one checkpoint in the four columns of our design tool. At this point we are ready to draw instances of our design tool diagram. We create one diagram for each of the groups of artefacts we have identified. To do this, we simply go through each of the artefacts in a group, and draw all those points (ellipses) that have a checkmark in our

table. Then, we draw the connectors between those points that are actually related in our system design. Note that we should use dotted lines in our diagrams in exact as we did in the diagram in Figure 6.4.

To continue with our example of the PerfEx system, we have already created the initial checklist, as in Table 6.1. Next, we have to identify potential groups in this list. As it stands, all the artefacts in our list form one group. However, we could think of situations where these artefacts created two subgroups. In fact, as we show in Section 6.6, thinking of ways to separate out the groups can offer solutions to potential problems. We could, for instance, propose that the discussions be of a more secure nature, thus falling under the social sphere. In this case we could, for example, assume that the discussions are of no relation to the train timetables and therefore we have two groups: the first three artefacts are one group, and the last three artefacts are another. The definition of groups reflects the requirements of the system.
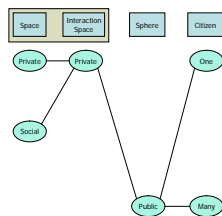
But as we just noted, currently the artefacts in Table 6.1 only create one group. We have already seen the diagram instance for this group in Figure 6.3. A smaller version of this diagram is shown in the sidehead. We are now ready for inspecting the diagram, identifying potential problems, and proposing changes and solutions to the problems.

6.6   Step 4: Making the changes

We now (❹) have to work through the elements and connectors in the design tool, paying particular attention to the dotted lines. This should be based on our understanding of each of the elements that have been plotted on the diagram, as well as each of the connectors that have been drawn. For a full description of the connectors, see Section 6.2.

After inspecting the diagrams, the designer is at a stage where she can make design decisions and recommendations. Most of the design decisions to which the design tool will help guide the designer are classified as follows:

- Change the technology being used;

- Redesign the physical spaces;

- Relocate artefacts;

- Reconsider the links between artefacts.

Change the technology being used

We may manipulate the existing interaction spaces to overcome problems identified. This manipulation can be done by modifying existing technology. In Section 5.3 we described our approach to designing pervasive systems in terms of spaces and interaction spaces. We said that the effective integration of spaces and interaction spaces are key to the success of a pervasive system. By changing the technology being used, or manipulating and tweaking existing technology, we can effectively manipulate interaction spaces. This redesign option is one of the easiest and most flexible to implement, especially compared to redesigning physical spaces. Also, this option can be easily applied in situations where a system already exists and is being used.

In terms of our design tool, the effect of this option is to manipulate the points under the interaction spaces column, shifting them up or down according to our design decisions.

Redesign physical spaces

Carrying on with our design approach, the latter half consists of redesigning physical spaces. This can be much harder than redesigning the technology, and in some cases may be altogether impossible. However, we need to acknowledge that in the past architecture *has* been affected by the technology being used inside buildings, and has adapted and responded accordingly to take advantage of this. For instance, the introduction of the elevator allowed architects to design high-rise buildings, and

the introduction of the telephone made convenient separate offices located next to each other. Therefore, although this option seems quite inappropriate in the short term, in the long run the redesign of physical spaces will play a crucial role in the success and acceptance of pervasive technologies.

In terms of our design tool, the effect of this option is to manipulate the points under the spaces column, shifting them up or down according to our design decisions.

Relocate artefacts

Pervasive technologies include embedded technologies and services, and to a large extent they are dependent on the location in which they are made available. Relocating artefacts suggest that the designer has identified some of the artefacts as having been misplaced. The designer can suggest that they should be moved to a different location. This can generate a requirement to create a new physical location for particular artefacts, and in some cases may dictate the type of physical spaces that ought to be created (private, social or public).

Relocating artefacts (essentially relocating technology) can have two effects: it may introduce a new location artefact in our design specification, or it may cause a change in the composition of the groups of artefacts we have identified. In either case, the diagram instances that describe our system will change and by doing so we can overcome some of the problems with our design.

Reconsider the links between artefacts

The most generalised approach to suggesting changes and improvements to a pervasive system design and specification is to reconsider the links between artefacts. The previous approach - relocate artefacts - was a straightforward way of doing so. But depending on the composition of groups of artefacts in our specification, the designer could deem it appropriate to move an artefact from one group to another, break up a large group of artefacts into two or more groups, or even abolish an arte-

fact altogether. This can result in new instances of our design diagram, which can be examined and evaluated in iteration.

6.7   Summary

In this chapter we gave a full description of our design tool, which takes the form of the design diagram that operationalises our framework. We have also described our method for using the design tool to inspect and improve designs. The examples we have used in this chapter have been quite basic for the sake of clarity. The whole process we have described here can and should be iterated. Having made design decisions and recommendations, we can go back and follow the process again, using the design tool to help us explore any new arrangements we have proposed.

At this point we have developed and operationalised our ideas, and we are now ready to apply them. In Chapter 7 we show how our ideas can have an impact at the interface level of system design. So far we have discussed the delivery of information in varying interaction spaces. However, the interaction method itself (i.e. keyboard, mouse, touch screen) can sometimes come in direct conflict with the interaction spaces used to deliver the information (see {I2}{I3}{I4}{I7} below). In Chapter 7 we show how we developed an interaction technique that can generate appropriate interaction spaces.

We then proceed in Chapter 8 to apply our design tool to a real-world situation: the Accident and Emergency department of a busy hospital in London. Following design objectives and recommendations based on our ethnographic study, we show how our design tool and method can be used to derive design solutions and alternatives and how to evaluate them. Having done so, we then present another case study in Chapter 9, where we set out to generate design recommendations using our framework. By studying specific locations in the city of Bath we generate a

number of objectives and recommendations which can be turned into concrete design solutions using the process we describe in Chapter 8.

As a way of helping the reader in understanding and using our ideas and tool, we reproduce here key features of the discussion in this chapter. This can be used as reference material, and may be convenient for following the discussion in the following chapters.

Issues between spaces and interaction spaces

{I1} Interaction spaces are ultimately *bound by the physical and social constraints* of spaces (page 123).

{I2} Downward sloping connectors (reading Figure 6.1 from left to right) between spaces and interaction spaces highlight that *direct physical interaction* may not be possible for the task at hand.

{I3} Interaction spaces can span various types of spaces, but this *does not support direct physical interaction and manipulation* by the participants in the interaction space (page 123).

{I4} The more upwards the slope is between spaces and interaction spaces connectors, the more care needs to be taken not to *intrude on privacy*. Here, the technology must be what we term "insulating" (page 124).

{I5} Private spaces to private interaction spaces support direct physical interaction, but at the cost of *not being able to span multiple spaces* (page 125).

{I6} More than one social interaction space (or other types of interaction spaces as well) can be present within a social space, but if the various interaction spaces are unrelated in terms of task, sphere, or intentions, then the participants will be broken up into groups, thereby *reducing interaction* amongst them (page 125).

{I7}   Public interaction spaces to public spaces are exemplified by shared public spectacles, such as street performances. In such cases, *all participants need to be included* (page 125).

{I8}   A space can have more than one interaction space in it, but this may potentially create *cognitive overload* for those within the spaces (page 125).

Issues between interaction spaces and spheres

{I9}   Interaction spaces and spheres connected with upward connectors denote insecure design options, i.e. design options that have the *potential to undermine the privacy or security* of information due to spillovers (page 125).

{I10}   Private interaction spaces offer the least problematic connectors, but at the *cost of physical interaction*: by definition a private interaction space has only one person in it (page 126).

{I11}   If more than one social sphere is accessed in a social interaction space (thus breaking the one-to-one relationship), there is a risk of *fragmenting the resources required* to sustain the activity at hand (page 126).

{I12}   When public sphere information is accessed from non-public spaces there should be a *consensus of common interest* or a common activity (page 126).

Finally, in the following chapters we will be referencing the above issues and the steps in our method quite often. To make it easy for the reader to understand which implication we are discussing, or which step in our method we are applying, we will be using **{I1}** to **{I12}** to reference the above implications. We will also be using the symbols ❶, ❷, ❸ and ❹ to indicate the method step we are involved in. These symbols will be placed whenever our discussion deals with:

- ❶: Generating a list of artefacts

- ❷: Producing a check-list

- ❸: Deciding on sub-groups and creating diagram instances

- ❹: Proposing solutions and making changes

AN INTERACTION TECHNIQUE

7.1 Keyboards forever?

To many, Milwaukee will always be known as the city of festivals and beer. But for a few, Milwaukee is known as the birthplace of the typewriter. The origins of the modern computer keyboard began humbly with the invention of the typewriter by Christopher Latham Sholes in September 1867. After completing his schooling, Sholes was apprenticed as a printer and four years later became editor of the Wisconsin Enquirer, in Madison. After a year, he moved to Kenosha to run the newspaper there and soon entered politics, serving in the state legislature. In 1860, he became editor of the Milwaukee News and later of the Milwaukee Sentinel. He later gave up that position to accept appointment from President Lincoln as collec-

tor of the port of Milwaukee. His new, less demanding job gave Sholes the time he needed to exercise his inventive genius. In 1864, he and a friend, Samuel W. Soule, were granted a patent for a page numbering machine. A fellow inventor-mechanic, Carlos Glidden, suggested to Sholes that he might rework his device into a letter printing machine and referred him to a published account of a writing machine devised by John Pratt of London. Sholes was so intrigued by the idea that he spent the remainder of his life on the project.

Sholes patented his first keyboard design in 1868. It was subsequently improved and marketed by Remington in the late 1870s. When initially introduced in 1877 the first typewriter models were utilized by typing with only two fingers. The development of 10 finger typing is attributed to a Mrs L. V. Longley in 1878. Shortly after, the concept of "touch typing" was introduced (attributed to Frank E. McGurrin, a federal court clerk in Salt Lake City), whereby typists would type without looking at the keys, having memorized their locations. Touch typists were so quick that Sholes had to redesign the keyboard (into our familiar QWERTY) layout to reduce the number of mechanical jams. These new techniques, and some celebrated typing competitions, demonstrated the worth of the new machine and led to continuing increases in sales.

Yes, but...   The keyboard has a long history of being used as a tool to help us in entering text. As times and technology changed, so did the keyboard. The keyboard's purpose has also changed, ranging from writing memos and Telex messages to writing C++ programs and playing games. Unfortunately, we cannot use the keyboard (and mouse) everywhere. Granted, the keyboard is currently one of the best and fastest ways of entering text into a computer. It has been around for many years, and we know how to design our software to work well with the keyboard and its faithful partner,

the mouse. The ergonomics of keyboards and mice have been studied, and various aspects of human physiology have been incorporated in their latest designs.

The main objection, however, to the use of keyboards and mice in pervasive systems is that they were not conceived for such use. Rather, they were designed as static input devices, with a one to one mapping between keyboard-mouse and the computer. Recently, we have seen the initial assumptions and understandings about the keyboard and mouse being stretched to their limits. Keyboards are being miniaturised and attached to people's sleeves, they are being made virtual and displayed on 5x5 cm screens, wireless mice are being used to control more than one computer and so on.

A new approach is required for a mobile and pervasive environment, developed with an appropriate set of requirements in mind. Desktop computers are typically used for word processing, spreadsheets and the creation of presentations, tasks that are not the primary focus of mobile, wearable and pervasive computing. Therefore, the input devices and methods used with desktop computers and for such tasks cannot be seamlessly transferred to a pervasive computing world.

How interaction relates to our framework

In terms of our framework, keyboard and mouse interaction falls short on many different fronts when used with varying interaction spaces. Both the keyboard and mouse have static physical dimensions, and therefore impose an extra set of requirements and limitations upon the system with which they are used. Shrinking the keyboard in size presents usability problems, since the physical form of humans (i.e. the size of their fingers, how well they can see, etc) comes in direct conflict with miniature keyboards.

Furthermore, the privacy offered by the keyboard and mouse cannot reflect and match the interaction spaces used to access information. There is only one way to

interact with a keyboard and mouse, and this has to be replicated across all possible situations. The interaction itself can come in conflict with the privacy that interaction spaces offer. For instance, a screen's private interaction space used with a keyboard could give away information simply by the clicking noise of the keyboard {I9}. Conversely, the keyboard or mouse may not work well with public interaction spaces where many people are involved {I6}{I7}.

Keyboards offer no flexibility either in terms of physical form nor in terms of privacy. Because the interaction spaces created by keyboards and mice cannot realistically be dynamically changed, there is no way of matching the keyboard's interaction space to the interaction space being used to access information.

In this chapter we describe an interaction technique which is suitable for interaction in a pervasive environment. We show how our design framework provided a basis for the design requirements of this technique, and how the issues we have discussed so far take form when implemented in an interaction technique. Specifically, we address the issue of how interface interaction can be related to our framework. To do this, we provide an interaction technique that can dynamically generate public, social and private interaction spaces in which subtle information about the interaction itself may be given.

For the sake of clarity, we first describe related work, and then describe the technical workings of our technique. With these workings in mind, it is easier to understand what our interaction technique is about, and it should become clear how our technique ties in with our design framework.

7.2 An introduction to stroke recognition

Given the inadequacies of traditional desktop input techniques, i.e. mouse and keyboard, in a pervasive computing environment and, even more so, with mobile and

wearable computing, there has been considerable research investigating alternative techniques. Prominent amongst these is gesture or stroke based input [Pirhonen et al. 2002]. Gesture and stroke interaction refer to interaction techniques where the physical movement of some specific item (a mouse, a hand) denotes input to the system. The difference we see between these two different terms is that gesture interaction can involve any type of physical movement, while stroke interaction typically describes physical movement in straight lines only. These types of interaction have formed the basis for many of the input techniques used with PDAs, whether in the form of touchscreen strokes to perform commands or in the form of alphabets, such as Graffiti on the Palm range of PDAs.

Moreover, stroke recognition predates PDAs by quite a while. One of the first applications to use some sort of stroke recognition was Sutherland's sketchpad [Sutherland, 1963]. The idea of mouse strokes as gestures dates back to the 1970s and pie menus [Calahan et al., 1998]. Since then, numerous applications have used similar techniques for allowing users to perform complex actions using an input device. For instance, design programs such as [Zhao, 1993] allow users to perform actions on objects by performing mouse or pen strokes on the object. Recently, Web browsing applications, such as Opera[24] and Mozilla[25], have incorporated similar capabilities. Guimbretiere et al. [2001] show how the FlowMenu system [Guimbretiere & Winograd, 2000] may be used with large wall displays. FlowMenu is very similar to pie menus. Unless the FlowMenu has been displayed, any pen stroke is interpreted as simple mouse input, using a simple down-move-up event model.

---

24. see http://www.opera.com
25. see http://www.mozilla.com

Sutherland's Sketchpad

In the early sixties, few computers ran "on line", as opposed to "batch mode". One notable exception was the TX-2 computer at MIT's Lincoln Laboratory. The Air Force paid Lincoln Laboratory to build TX-2 as a demonstration that transistors, themselves relatively new, could be the basis of major computing systems. TX-2 was a giant machine by the standards of the day, in part because it had 320 kilobytes of fast memory, magnetic tape storage, an on-line typewriter, the first Xerox printer, paper tape for program input, and most importantly, a nine inch CRT. The display, a lightpen, and a bank of switches were the interface on which Ivan Sutherland based the first interactive computer graphics. In 1963, his PhD thesis, "Sketchpad: A Man-machine Graphical Communications System", used the lightpen to create engineering drawings directly on the CRT. Highly precise drawings could be created, manipulated, duplicated, and stored. The software provided a scale of 2000:1, offering many acres of drawing space. Sketchpad pioneered the concepts of graphical computing, including memory structures to store objects, rubber-banding of lines, the ability to zoom in and out on the display, and the ability to make perfect lines, corners, and joints. This was the first GUI (Graphical User Interface) long before the term was coined. In 1988, Ivan E. Sutherland received the ACM Turing Award. (Source: http://www.sun.com/960710/feature3/sketchpad.html)

There is a number of current open source projects that involve the development of stroke recognition, including Mozilla, Libstroke,[26] X Scribble,[27] and WayV.[28] The latter is a library created for recognizing characters as well as strokes. It is based on a technique called point density analysis which uses matrix mathematics. In its latest version it has included a second "backup" method for recognizing strokes, which implements a form of directional recognition. This method imposes an n by n matrix on the stroke and assigns every stroke point to a cell in the matrix. By comparing the relative position of two boxes which contain consecutive stroke points, a sequence of directions is produced, and is used to assist the point density analysis algorithm in the recognition of the stroke.

---

26. See http://www.etla.net/libstroke/libstroke.pdf
27. See http://www.handhelds.org/projects/xscribble.html
28. See http://www.stressbunny.com/wayv/

The Mozilla browser uses a simpler technique. Each point of the gesture is compared to the previous one, and one of four directions is generated (U, D, L, R), while discarding consecutive Us, Ds etc. Then, the sequence is compared against a table of stroke signatures, and if no exact match is found, then only the last 2 and then the last 3 elements of the direction signature are used. If that fails, then the signature is processed for diagonals, simply by replacing consecutive Ls and Ds by '1' (for diagonally left-down), Rs and Ds by '3', Ls and Us by '7', and Rs and Us by '9'. Then, this modified signature is checked against a table for matches.

Learning techniques have been applied to stroke recognition, with some success. For instance, Boukreev[29] has implemented stroke recognition using neural networks. This technique involves recording the path of the stroke, smoothing it to base points, translating it to the sines and cosines of the points' angles, and then passing these values to a neural network. The neural network will try to recognise the stroke, and in the process of doing so, will actually improve its recognising capability.

Stroke interaction for pervasive computing

Its range of uses over the past three decades illustrates a key characteristic of stroke recognition as an input technique: it is not tightly bound to a particular device. Our aim in this research is to exploit this characteristic to develop an input technique that can be used seamlessly across the wide range of devices in a mobile device-populated, pervasive computing world.

The diverse characteristics of such devices, and potential future devices, impose key requirements on such an interaction technique. At one end of the scale, the user may wish to interact with a device as limited in processing power and surface area as a smart ring or credit card, perhaps using a stylus to make the gestures. At the other end of the scale, the user may wish to interact with a wall-size display, perhaps

29. See http://www.generation5.org/content/2001/gestureapp.asp

using the smart ring itself, or indeed using just the user's hand, to make the gestures in the air. Furthermore, the user may wish to interact with a private, social or public interaction space, and the interaction technique itself should not conflict this {I2}{I3}{I4}{I7}{I9}. The choices that users make can be based on the degree of privacy that each method provides, since the interaction spaces created by every technique are distinct.

In the following sections we present our technique for recognizing input strokes which can be used successfully on devices with very low processing capabilities and very limited space for the input area. The technique is based on the user's denoting a direction rather than an actual shape and has the twin benefits of computational efficiency and a very small input area requirement. This means that physical constraints do not impair the interaction technique itself. We have demonstrated the technique with mouse input on a desktop computer, stylus and touchscreen input on a wearable computer and hand movement input using real-time video capture. Finally, we discuss how our technique makes use of our concepts of interaction spaces and public, social and private spaces.

7.3   The directional stroke recognition algorithm

As its name implies, this is a technique for recognizing strokes based solely on their direction. Other characteristics of a stroke are not used. For instance, the position of a stroke is of no importance, nor are the relative positions of several strokes. This enables our technique to be completely free of physical form, and thus not be affected by, or affect, the physical form of the system itself.

Recognising a single stroke

The first step of our stroke recognition method is to collect the input data. Typically, the data for a stroke performed by the user is a set of coordinates. Our method regards the input stroke as an ordered set of lines. Each line consists of a "fromPoint" and a "toPoint". Using these two coordinates, we can calculate the

direction of each line as shown in Figure 7.1. For consistency purposes, every direction is represented by a two letter combination. Therefore, North, East, South, and West are represented as NN, EE, SS, and WW respectively.

Humans tend to be more accurate at drawing vertical and horizontal lines than diagonals [Pirhonen et al. 2002], especially when on the move. Therefore, by adjusting the relative angle for acceptance, e.g. a variation of 25 degrees for diagonals and 20 degrees for other strokes, we may accommodate for inaccuracies in stroke directions.

At this stage we have a stream of "directions", for example: "SS, SS, SS, SS, WW, WW, WW, WW, NW, NW, NW". The next step is to remove noise from this stream. This is achieved by setting a threshold as a percentage of the length of the whole stroke. This threshold is applied by removing any sequence of identical directions that does not reach the threshold. So, for example, for a threshold value of 10 percent and a stroke recorded as a stream of 40 directions, any contiguous sequence of fewer than 4 identical directions would be removed.
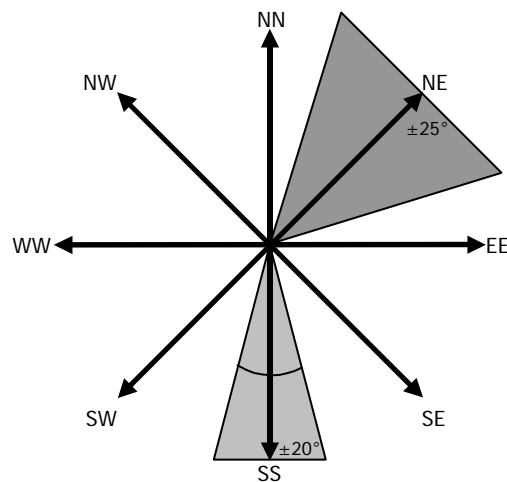
FIGURE 7.1: Calculating the direction of a line.

For consistency purposes, every direction is represented by a two letter combination. Therefore, North, East, South, and West are represented as NN, EE, SS, and WW respectively.

It is worth noting that this method performs very badly when given a stroke which is a curve. Because a curve is a sequence of lines which continuously changes direction, our method would calculate that the whole curve is noise, and thus would not be able to recognize it.

Having removed the noise, we then reduce adjacent appearances of a given direction to just one occurrence. At this point, we are left with a "signature" that looks, for example, like "SS, WW, NW". Using the signature that we have derived from the stroke, we can execute predefined operations. Some sample strokes along with their signatures are shown in Figure 7.2.

Recognising multiple strokes

The next development of our technique was to recognize gestures that consist of more than one stroke. In order to allow users to perform multi-stroke gestures, the GUI has to allow for a short "timeout" period, in which the user is able to stop drawing a stroke and start drawing a new stroke. For example, in the case of pen input, the user would be able to draw a stroke, lift the pen, and within the timeout period start drawing a new stroke. In this case, a special symbol may be used in the input stream to denote that the pen was lifted.

FIGURE 7.2: Some strokes and their signatures.

Note that strokes like NN-SS and EE-WW-EE do not need to form an angle, but are shown like this for illustration purposes.
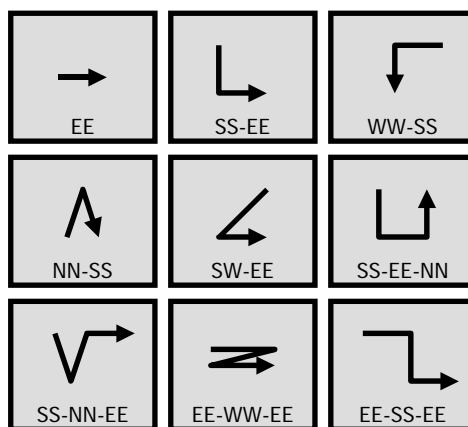
Having allowed for gestures consisting of more than one stroke, we have introduced an interesting characteristic. Now, different gestures may map to the same signature. Thus, a signature and, in turn, an operation can have more than one way of being accessed. For example, a stroke that looks like L and a gesture that looks like a cross may have the same signature, as shown in Figure 7.3.

This flexibility is beneficial when users may be working with multiple devices, each with different form factors and characteristics. In the case where screen size is limited, users may choose to decompose a gesture as finely as they wish, even into separate single-line strokes. The single-line strokes may be performed on top of each other, thus requiring less space on the screen. On the other hand, users with enough space may choose to perform one long composite stroke in order to save time, or in order to allow others to see that the interaction is taking place (This relates to our discussion on providing appropriate interaction spaces; see page 153.)

It may be argued that the number of possible operations is limited when many strokes are mapped to the same signature. Although this is true, we believe that in many cases the flexibility provided by our method outweighs the need for a plethora of different operations. The optimum solution is probably to allow the user to

FIGURE 7.3: Different strokes with identical signatures.

For any possible stroke, users may decide to break up the stroke into any number of substrokes, and then perform each substroke independently, regardless of its relative position to the rest of the substrokes.
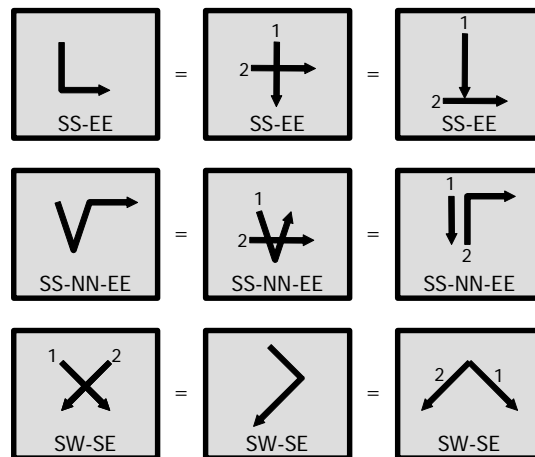
choose between different gestures mapping to the same or different signatures, per-
haps according to the visual similarity of gestures as described in [Long et al. 2000].
Rubine [1991], for example, demonstrates an approach to training a system for var-
ious gestures.

7.4   Real-time video capture of hand movements as input

In principle, our stroke recognition method deals with pure coordinates and noth-
ing more. Therefore, any input technique can work with our stroke recognition
method, so long as there is a meaningful way of deriving a set of coordinates from
the input technique. This supports our aim of developing a flexible interaction
technique that can be used across multiple devices and platforms (Figure 7.4). For
example, in the converging world of mobile and pervasive computing, our user may
at one moment wish to interact with her PDA using a common set of gestures and
in the next moment move seamlessly to interacting with the wall display beside her
using the same set of gestures. At one moment the PDA provides the interaction
area on which the gestures are made using a stylus {I4}; in the next moment, the
PDA itself becomes the "stylus" as our user makes the gestures in the air with the
PDA while interacting with the wall display {I7}. As a proof of principle, we imple-

FIGURE 7.4: Stroke recognition can be used with various techniques.

Any method and input technique that can produce a meaningful set of coordinates may be used with our stroke recognition technique.
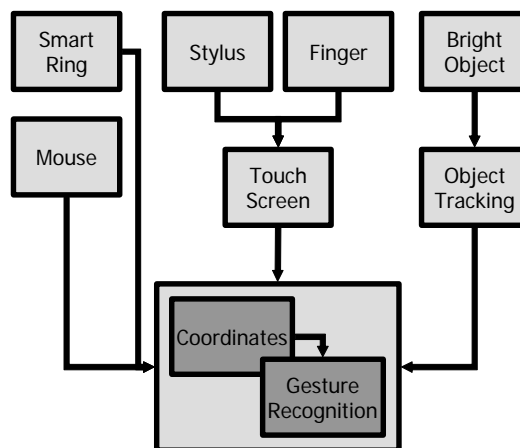
mented a real-time object tracking technique that we then used along with our stroke recognition algorithm as an input technique.

For our prototype we implemented an algorithm that performs real-time object tracking on live input from a web camera. The user can select a specific object by sampling its colour, and the algorithm tracks this object in order to generate a series of coordinates that describe the position of the object on the screen, or to be precise, the position of the object relative to the camera's view. We then pass these generated coordinates to our stroke recognition algorithm, which proceeds with the recognition of the strokes.

Due to the characteristics of our stroke recognition method, the coordinates may be supplied at any rate. So long as this rate is kept steady, the stroke recognition is very successful. Thus, despite the fact that our object tracking algorithm is not optimal, it still provides us with a useful prototype.

Object recognition is performed using HSB (Hue - Saturation - Brightness) sampling. The object to be tracked is described in terms of HSB based on its colour. We then apply a varying threshold of approximately 5% to the live video input, which results in certain pixels being identified as belonging to the object. We then perform a second pass in order to identify which region has the highest density of object pixels, and from this we derive the object's centre. These coordinates are then passed on to the stroke recognition algorithm.

Lifting the pen An issue that we had to address was how to allow for an act corresponding to lifting the stylus from a touch screen. Our initial approach has been that the user can, for example, hide the object within the palm of her hand. This does not limit functionality in terms of stroke recognition - remember that any signature can be performed as one long stroke - and contributes to the similarity of input methods

The HSB colour model

HSB stands for Hue, Saturation and Brightness. According to this model, any colour is represented by 3 numbers. The first number is the hue, and its value ranges from 0 to 360 degrees. Each degree represents a distinct colour. First there is red (0 degrees) and then there are all other colours (for example yellow at 120 degrees, green at 180 degrees and blue at 240 degrees), up to violet. All the rainbow's colours are represented here. The second number is the saturation. It represents the amount of colour or, more exactly, its percentage. Its value ranges from 0 to 100, where 0 represents no colour, while 100 represents the full colour. Finally, the third number is the brightness. You can enhance the colour brightness by adding white, or you can reduce it by adding black. In this case zero represents white and 100 represents black. The more this value tends to 0, the brighter the colour is. The more this value tends to 100 the darker the colour is. Other popular colour models are RGB (Red Green Blue), CMYK (Cyan Yellow Magenta blacK) and Lab.
(Source: http://www.wowarea.com/english/help/color.htm)

across platforms and devices. A potentially less cumbersome solution to this would be to allow for an "invisible" light, such as infrared, to be emitted from a small hand-held object and to be used by the object-tracking camera. Such an object could be a dedicated input device for interacting with pervasive computing facilities and could emit light when squeezed or held at a certain orientation. Alternatively, and more in line with our general vision of integration across mobile and pervasive devices, a PDA or other device that can provide an input area (e.g. a touchscreen) for gestures could also act as a "stylus" by emitting infrared on user demand. This would allow for seamless transitions, using a common gesture alphabet, between interacting with a mobile device and interacting with surrounding pervasive computing devices.

7.5   An interaction technique for public pervasive systems

We have developed the directional stroke recognition technique as a way of exploring our design framework and deepening our understanding of pervasive systems.

In this section we give a synopsis of how our technique can work in a pervasive computing environment, as well as how it makes use of our concepts of interaction spaces and public, social and private spaces.

Switching between devices

The flexibility of our method allows switching between input devices and methods with no need to learn a new interaction technique. Any object or device that can provide a meaningful way of generating coordinates and directions can provide input to the gesture recognition algorithm (Figure 7.4).

Some important characteristics of this technique, which closely reflect the concepts of our framework, include the ability for users to choose the scale and nature of the interaction space they create, thus influencing the privacy of their interaction {I2}{I3}{I4}{I7}{I9}. In addition, the physical manifestation of our interaction technique can be tailored according to the situation's requirements. As a result, the technique also allows for easy access, literally just walking up to a system and using it, with no need for special equipment on the part of the users. We now describe these aspects in more detail.

Providing appropriate interaction spaces

An important characteristic that reflects our concern with private, public and social spaces and interaction spaces is that the user of our technique can choose the interaction space she creates. This ability and flexibility to tailor the interaction space according to our wishes and needs can greatly enhance sociability and the sense of safety in a public or social space. For example, people may choose to create a public interaction space by making large strokes with their hands or an object in the air in order to, for instance, change the channel on a big screen television {I7}. This allows everyone watching the TV to be aware that someone is interacting with it and not, for example, to assume that the TV is randomly switching channels. Likewise, a group of people may use a common artefact, such as a digital game board with stroke recognition, to create a social interaction space. Finally, a private inter-

action space can be created through the use of small strokes on or with personal devices such as PDAs, smart rings and smart cards. This would be appropriate for use in situations where a private interaction space is used to access the information {I4}.

At any moment a person may choose to revert to any of the possible modes of interaction spaces, thus including or excluding at will other people within the same physical location. This encourages people to invite others into social spaces, a goal of projects such as HP's Schminky [Reid et al., 2003], and allow the easier coordination of activities in public spaces. Hence, this flexibility in selecting the interaction space can enhance the sociability of a space, while at the same time respecting privacy needs and ownership.

With our technique existing technology and design restrictions can be respected since our input technique can be used in all types of spaces and interaction spaces. Therefore, if in a specific setup we have technology that creates private interaction spaces, and the design requirements state that only private interaction spaces should be used, then our technique can be integrated in this setup, and be used only in a "private interaction" mode. In later chapters (see Figure 8.3) we give examples of design requirements and how we can manipulate interaction spaces to adhere to those requirements.

Physical manifestation

We have suggested that the physical manifestation of a pervasive system should be taken into account (see Section 5.5). Different systems in different settings have to explore how this can be done, and how to take full advantage of the effects of doing so. We argue that the full potential of the physical form should not be constrained by the interaction requirements, or vice versa. Hence, there is a requirement for an interaction technique that will transfer across systems, yet will not impose restrictions on these systems' physical forms. Also, the interaction technique should

respect the restrictions imposed by the system itself, and should not introduce discrepancies from these restrictions.

The Directional Stroke Recognition technique is flexible enough to accommodate a range of technologies (and their physical forms) yet provide the same functionality wherever used. Thus, issues concerning physical form may be addressed independently. Such issues include "How is the system represented or made visible?", "Are people aware of the system?", and "Are we aware of who is using the system?".

We are not suggesting that a single solution applies to all situations. For instance, in one situation you may want people to walk up to a specific spot in order to use the system, while in other situations you may want to avoid the potential congestion of doing so. In both of these examples, what is required is an interaction technique which has been separated from the physical form of the system. In contrast, current window-based interaction techniques are closely tied to physical form: mouse, keyboard and monitor. Consequently, they do not even transfer well to interaction with current PDA designs, much less innovative pervasive systems. The technique we have described goes a long way towards the separation of the physical form and interaction technique, which will allow us to take full advantage of the physical forms of public pervasive systems.

Easy access   The flexibility inherent in stroke interaction, and the separation we have achieved between physical form constraints and interaction constraints, has allowed us realistically to consider providing easy access in a pervasive environment. By easy access, we mean systems that do not require any special equipment on the part of the user, any intervention in the system's operation on the part of the user, and in general systems that the user can just walk up to (or into) and use.

We have demonstrated how our technique addresses these issues through test runs as illustrated in Figure 7.5. In this figure, a user is holding a highlighter pen in order to interact with the system. Other objects we have used include tennis balls, hats, books, etc. This flexibility has a direct impact on two different fronts. Systems can become more truly pervasive by encompassing a wider range of everyday objects and forms in their operation. But at the same time, they can also become more socially available, since their use does not have to be inhibited by economic or other factors - i.e. a user does not need to buy an expensive device in order to use a public pervasive system.

FIGURE 7.5: Stroke recognition used with object tracking.

Interacting with the DSR system using a highlighter pen. The user can sample the object to be tracked by specifying its Hue, Saturation and Brightness (Luminosity) as well as its approximate width and height. The system overlays a bright yellow circle over the recognised object.
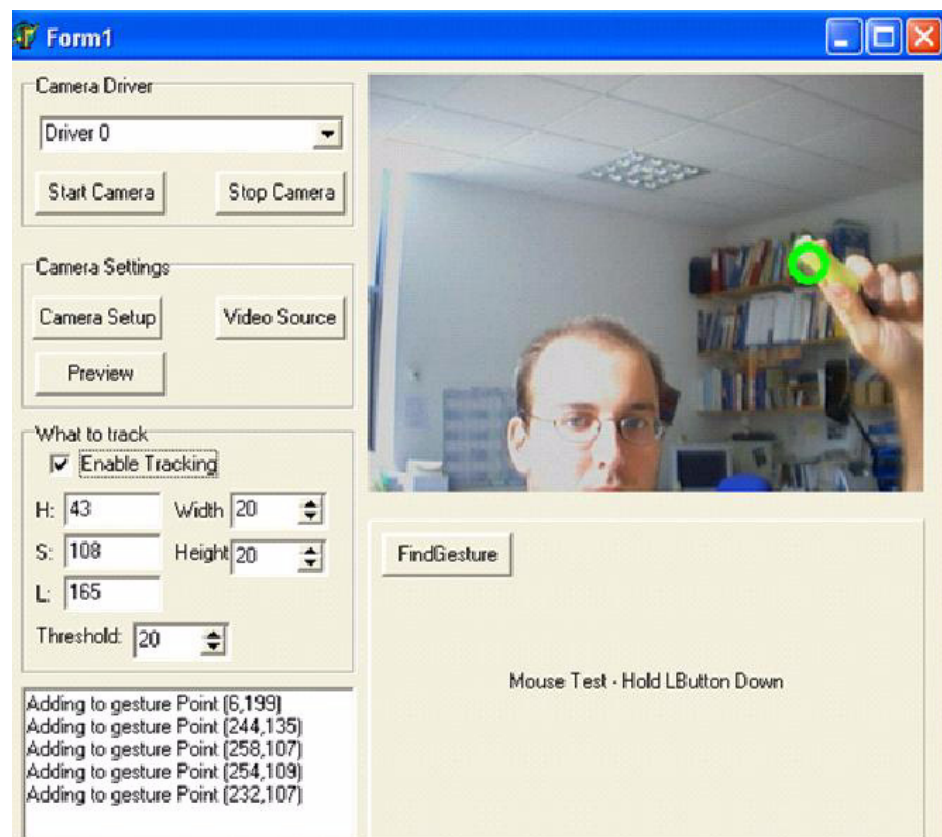
7.6   Summary

In this chapter we showed how the concepts of our framework have influenced the design of an interaction technique. At the end of Chapter 5 we listed some requirements that came out of our analysis of spaces and interaction spaces. Specifically, we said that public pervasive systems should enable users, given the resources available in a particular setting, to define interaction spaces suited to their moment to moment purposes and activities. We also said that in terms of the physical form of the system, it would be important to use common interaction techniques across a very wide range of devices with varying physical characteristics, thus freeing the interaction technique from the physical form of the system. Additionally, in Section 6.7 we mentioned the need to provide appropriate interaction spaces. An interaction technique's interact space needs to match the interaction spaces used to access the information.

In this chapter we have contributed to addressing the above issues by the directional stroke recognition technique. We showed how our technique addressed these important requirements for pervasive systems {I2}{I3}{I4}{I7}{I9}, as well as how it incorporated other ideas we have presented so far.

In the following chapter we will apply our design tool and method at a higher level: to a real-world case study of an A&E department in a London hospital. There, we apply our design tool to the findings of an ethnographic study in order to explore and evaluate design alternatives.

CHAPTER 8

CASE STUDY: DESIGNING FOR A HOSPITAL ENVIRONMENT

In this chapter we apply our arsenal at a higher level than in the previous chapter. We show how our design tool can be applied post hoc to explain problems that have already been identified in a real-life setting. Furthermore, we use our design tool to propose a new system and a priori identify potential problems with it.

The setting we tackle is the Accident & Emergency department of a London hospital. To do this, we have used material collected during an 18 month ethnographic study of this hospital. The fieldwork study itself,[30] reported in [O'Neill et al., 2004], focused on identifying ways of providing better information resources to the hospital patients. In many cases, the absence of this information had negative effects, including violence and assaults. In this case study we turn the fieldwork findings and proposals into design solutions.

Violence and abuse towards hospital staff are perceived to be sufficiently widespread and serious to necessitate the recent announcement by the UK Government of new national guidelines to help make hospitals safer environments for both staff and patients. Hospitals are now permitted first to warn and then to exclude without treatment patients or visitors who are violent or abusive. Despite such measures, the problem is reportedly still on the increase,[31] and seems to be particularly prevalent in A&E departments. Several incidents were observed or reported during our fieldwork. The majority of these incidents involved rudeness or verbal abuse of staff and other patients. However, even patients whose behaviour did not descend to these levels were frequently observed to show signs of annoyance, stress and exasperation, and these reactions tended to coincide with long waiting times.

As well as stress for the patients, continual requests for information about waiting times, even polite requests, also seemed to cause stress to the staff. The frequent need to respond to these requests was often distracting, interrupting their ongoing work. Such interruptions at times had the unfortunate effect of increasing the patients' waiting times still further. Hospital staff and managers are concerned that people have to wait so long for treatment, and are constantly reviewing the organi-

---

30. We are grateful to Dawn Woodgate who carried out the field work and provided us with data.
31. Miekle, J. Assaults rise despite crackdown. The Guardian, 27 March 2003.

zation of their work to try to reduce the problem. One consultant's comment was quoted in the press recently:[32] "There are times when I walk into the department and it is a total heartsink. Every cubicle is full, the corridors are full of people on trolleys, there is nowhere to sit and write up your notes, and you are trying to find space in the resus room […]". This concern is linked not only to the clinicians' obvious anxiety about the welfare of their patients, but also to negative reports in the media, and government directives which impose limits on waiting times but offer few suggestions for how to accomplish this in an aging facility with ever-increasing patient numbers.

8.1   Addressing problems in the A&E

Previous work has shown that certain characteristics of queues, such as uncertainty of waiting times and lack of information, can cause stress and antagonism [Stewart, 2002] and, more specifically, that urgent care patients who were told the expected waiting time for treatment and were kept busy while waiting, had higher percep-tions of satisfaction with their treatment [Naumann, 2001]. Maister [1984] sug-gested that customers who were given information about how long they would have to wait are less likely to be anxious about the wait. Dansky & Miles [1997] found that telling patients in an urgent care department how long they would have to wait was positively related to their satisfaction with the treatment.

Our study suggests that the provision of information of this type might be a useful tool not only for reducing stress and averting some violent or abusive incidents, but also in influencing patients' perceptions of satisfaction with their visit. In the A&E waiting area under study, there was some information on display, though nothing that related to likely waiting times, or the reasons why long waits occur for some patients. The only way patients could get this information was to ask a member of

---

32. Arlidge, J. Quick fix revolution. Evening Standard, 12 November 2002.

staff, usually a nurse or a receptionist. Some patients were observed to become angry when told that they faced a very long wait, especially if they had been waiting for some time already, and when other patients seemed to have been given priority over them. They were unaware of the reasons for this type of occurrence. The significance of the assessments of staff of the relative seriousness of a particular individual's condition, and the intricacies of triage categories for example, were not obvious to patients, and the prioritizing of some patients over others with little or no explanation was often observed to initiate negative remarks from patients who felt that they had been less favourably treated. It seems then that there are potential benefits to be had from making aspects of this information available to the patients. However, some of this information is private, and this must have a major impact upon the design of systems to provide such information.

8.2   A worked example - the case of the phone lists

In this section we apply our framework and design tool to data we have collected from our fieldwork study. Our goal here is two-fold; we wish to demonstrate that the ideas we have described so far, when applied to real situations can actually be useful in understanding and explaining the source of the problem at various levels. We also wish to show that real solutions to real problems can be proposed using our approach. The fact that previous research attributes violence to poor information provision is a very positive indication for us. A (pervasive) computer system that provides the right information at the right place and the right time, could solve many of the aforementioned problems of A&E departments.

As we have already described, the concepts of citizens, spaces, interaction spaces and information spheres - public, social and private - aid us in mapping from physical spaces to the technological artefacts that are available to us as developers and the forms of interaction we wish to support. Through the design of artefacts (dis-

play hardware, audio output, software interfaces etc), we define the interaction spaces within which people may interact with the information presented and, in some cases (i.e. public and social interaction spaces), with other participants. Through the identification of the information sphere into which a particular service falls, we can determine what kind(s) of interaction space - and, from there, what artefact designs - we require in order to provide that service in different settings, such as public, private or social spaces.

A simple scenario

To illustrate how our design tool can be used post hoc, in this section we go through an example scenario from the field data we collected during our ethnographic study. The scenario involves lists of telephone numbers and extensions of hospital staff. These lists were placed next to the reception, in the A&E waiting area. These lists were intended to be used by reception staff to help locate their colleagues.

As we described in Section 6.3, the first step in our analysis (❶) is to create a list of relevant artefacts. For the purposes of keeping our example simple, we present here a minimal list of relevant artefacts, consisting of: the telephone directory (information), the A4 sheets used to display the numbers (technology), the waiting room (location), and finally the reception area (location). Next, we create the check list for each item (❷), as shown in Table 8.1.

The telephone directory belongs to the social sphere, since the phone numbers and extensions are meant to be used by staff, and do not belong to the public domain. The A4 sheets of paper are a very flexible technology, and lend themselves to many kinds of interaction space. However, in this situation the A4 sheets of paper create a public interaction space {I1} and involve more than one citizen.

We have classified the waiting room as a public space because people can freely enter it whenever they feel appropriate (in most cases they do so when they are injured). On the other hand, the reception area is restricted from the public, as only staff are allowed to enter there and have access to whatever artefacts and information exist there. In both cases, more than one person is present within those spaces.

Next, (❸) we combine the artefacts into related subgroups. As things are, the "reception" artefact is currently unrelated to the rest of the artefacts as it is not involved in the delivery of information. Therefore, for now we can leave it out of our analysis. In other words, the group that we are currently interested in consists of all the artefacts besides the reception artefact.

We are now ready to inspect the current situation, and to do so we create an instantiation of our design tool that reflects the current situation (see Figure 8.1).

The first issue we identify is in the situation where the A4 sheets create a public interaction space. As we noted in Section 6.2 interaction spaces are bound by the norms and constraints of the spaces in which they are located {I1}. In our case, the A4 sheets are located in a public space, and they currently create a public interaction space which is used to access information in a social sphere. This could pro-

TABLE 8.1: An initial checklist of artefacts.

For each artefact we have identified in our list, we add a check mark under the categories of Space, Interaction Space, Sphere and Citizen which the artefact reflects.

| Artefact	Space			Interact. Space			Sphere			Citizen	
	P	S	Pr	P	S	Pr	P	S	Pr	M	1
Phone directory								✓			
A4 Telephone lists				✓						✓	
Waiting room	✓									✓	
Reception		✓								✓	

vide a privacy breach {I9}, since telephone numbers that are not meant to be public are currently being made public. Thus, an unintended harmful use of the current arrangement would be for a member of the public to abuse their access to the telephone numbers and cause problems by, for example, calling members of staff. (In fact, this occasionally happened in the situation.) In addition, there may be beneficial uses of this arrangement, such as staff who happen to be in the waiting area having easy access to these telephone numbers using the existing (paper based) technology.

In Figure 8.1 we have included two upward pointing arrows (❹). The first arrow carries the description of "relocation", and represents the design option of addressing the problem we have observed by relocating the artefacts, and thus involving a different type of space. The second arrow is described as "techniques", and repre-

FIGURE 8.1: A snapshot of our design tool.

A snapshot of our design tool which represents the use of A4 lists of phone numbers in the waiting area.
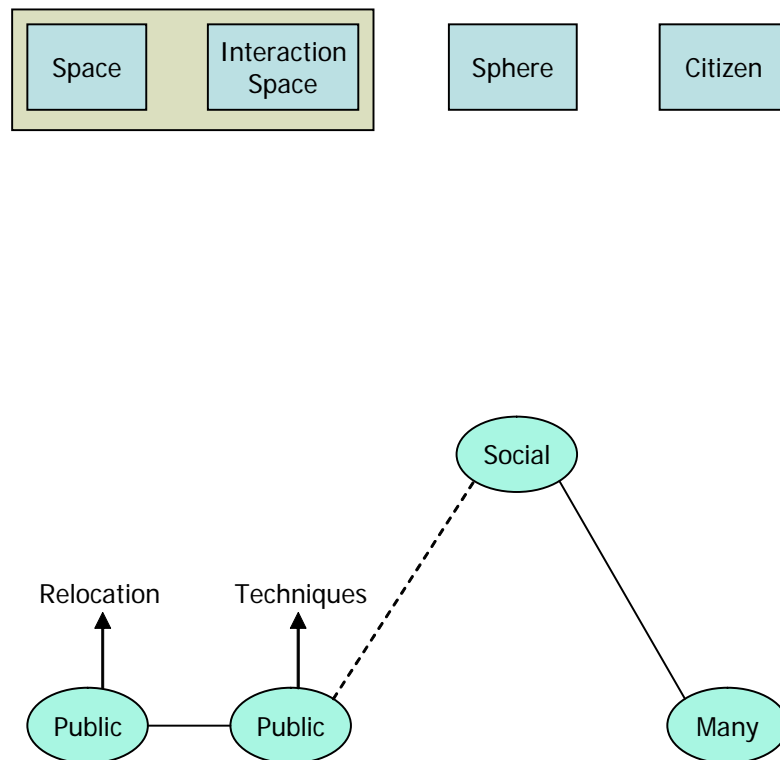
sents the design option of modifying or applying techniques to existing technology to manipulate the interaction spaces they create. For instance, the same technology (i.e. A4 sheets with phone numbers) could be distributed to each member of staff, thus {I1} using the A4 sheets to create private and social interaction spaces. In this case, however, the sheets would be mobile, and thus would exist in both public and social spaces.

We have noted that interaction spaces are bound by the spaces in which they exist {I1}. We can take advantage of this, and propose that the existing setting should be changed by relocating the A4 sheets in the reception (❹), in such a way that they are not visible from the waiting room. Taking this design option as our way forward, we can now iterate our design process. The option we have chosen results in changing the relations between the artefacts (❸), essentially redefining the subgroup in which we are interested. Our list of artefacts, shown in Table 8.1 remains the same. This time, however, the "waiting area" artefact is the one that we are not interested in, since it has no relation to the other artefacts (at this point in our analysis)(❸).

With our new subgroup in place, we can now study our new diagram instantiation, shown in Figure 8.2. Using our latest design tool instantiation, we are alerted {I11} to the fact that there is a 1 to 1 relationship that needs to be taken into account (see page 126). The fact that our (paper based) technology is static and constantly reflects the same information means that the 1 to 1 relationship is kept. However, this technology lends itself to unintended uses such as people writing messages or notes on the phone list. In such a case, the interaction space created by the A4 list provides access to more than one sphere: the phones list plus, say, instructions on how to deal with a specific computer problem. Finally, the many to many relation-

FIGURE 8.2: A new snapshot of our design tool.

A new instantiation of our design tool. In this case the A4 telephone lists have been placed in the reception area.
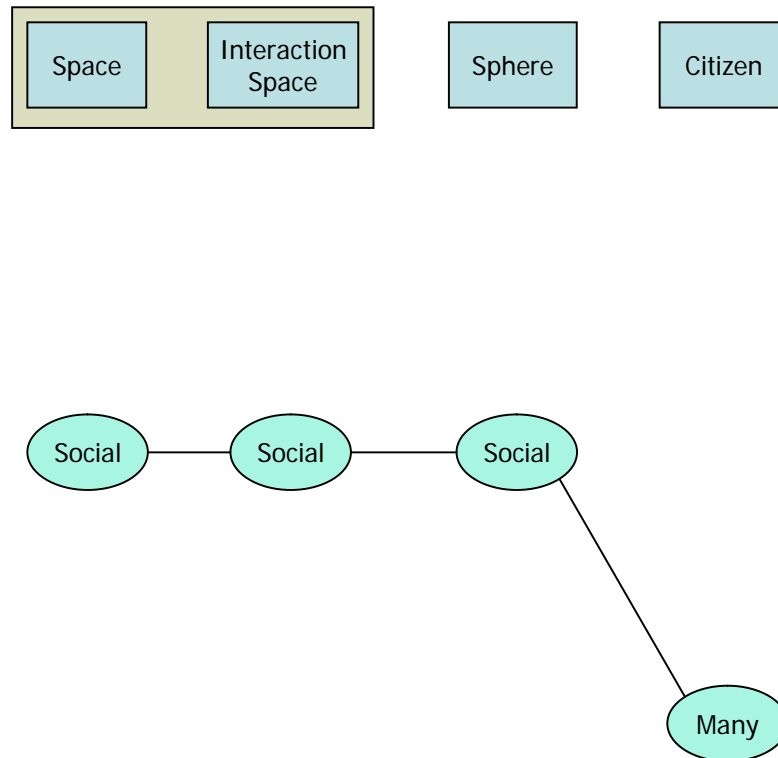
ship between social spaces and social interaction spaces {I6} is of no relevance to us since we are dealing with only one interaction space.

Instead of choosing to relocate the A4 sheets, we could have proposed to use completely new technology (❹). For example, we could have proposed to provide each member of staff with a personal electronic device that provides them with access to the telephone numbers. This device would be portable, and could be used in any room in the hospital. Furthermore, we would require these devices to create a private interaction space, in order that only the owner of this device would have access to the information that the device provides {I9}. In this case, our analysis (❷❸) would yield the following table and diagram (Table 8.2 and Figure 8.3.).

TABLE 8.2: Our checklist after new technology is introduced.

A checklist for each of the artefacts we have identified. In this case, the A4 telephone lists have been replaced by personal electronic devices.

Artefact	Space			Interact. Space			Sphere			Citizen	
	P	S	Pr	P	S	Pr	P	S	Pr	M	1
Phone directory								✓			
Personal device						✓					✓
Waiting room	✓									✓	
Reception		✓								✓	

FIGURE 8.3: A new snapshot of our design tool after new technology is introduced.

A new instantiation of our design tool. In this case the A4 telephone lists have been replaced by personal electronic devices.
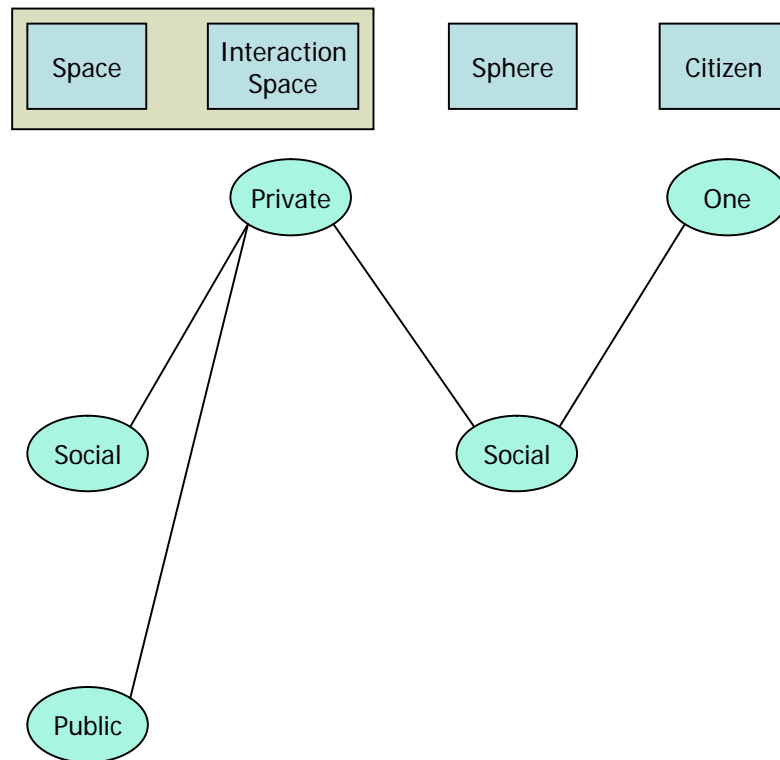
We have proposed that the upward slopes between spaces and interaction spaces require "insulating" technology (see page 124) {I4}. In this case, we would have to test that the personal electronic devices satisfy this criterion. For instance, when a member of staff used the device in a public space (e.g. in the waiting room), the

information displayed by the device, as well as the interaction techniques used, would have to "insulate" the user from the environment, thus denying anyone else access to the information.

Finally, we note that this simple example does not involve multiple participants exchanging information within a social sphere. Should that have been the case, we should observe that the location artefacts have different citizen checkpoints from the technology artefacts. In this case, we should be alerted to the many to one relationship between the social sphere and one citizen. Participants may access more than one sphere simultaneously, and this could be problematic if many participants had to collaborate for a particular activity.

8.3   More examples

The example we have just worked through has served the purpose of getting the readers acquainted with our use of language, and understanding how we work through the application of our design tool using our method. This example was meant to be simplistic, and does not address the real problems of information provision for the patients. Furthermore, this example did not lend itself to applying all of the possible solution paths that we have listed in Section 6.6.

In this section we wish to display the full power of our design tool and method. To do this, we have to forego the level of detail which we adopted for our previous example. The examples we work through in this section are situations which we have identified as important and likely to be the cause of the stress and assaults we described at the beginning of this chapter.

Information leaflets and posters

We begin our examples by addressing a problem we have reported in [O'Neill et al., 2004]. In our ethnographic study, one major problem which resulted in draining

staff time was the provision of public information to patients. Quoting our earlier work [O'Neill et al., 2004]:

> "It is useful first to summarize the existing types of information available to patients and visitors in the waiting area of the A&E department. At the time of our study, this took the form of posters and signs displayed on the walls, and also leaflets that could be taken up and read by interested parties and taken away should they wish for future reference. Typically, the posters and leaflets related to health information or health promotion advice, alerting people for example to the symptoms and dangers of particular serious diseases and conditions, or drawing attention to various charities and facilities offering help and advice to individuals and groups. The posters were designed to be eye-catching, featuring small quantities of information in large print, often in conjunction with photographs or other illustrations, whereas the leaflets were in small format, offering more detailed information. Many of the posters were somewhat faded and tired-looking, and some of the leaflet holders were empty. Leaflets in one area had become scattered over a tabletop and the floor. It was a further drain on precious staff time to keep displays featuring posters and leaflets in public areas in good order".
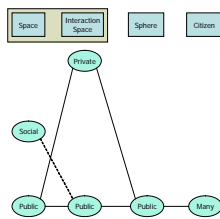
Apparently, the existing setup is unsatisfactory for A&E staff and this has taken its toll on patient satisfaction as well. It would seem reasonable for us to propose a digital solution to deliver the same kinds of information. In this section, we go through various development and analysis stages, propose and extend a pervasive system, and see how it satisfies the needs and problems we have identified. Specifically, we

• Describe the existing setup (i.e. posters and leaflets) using our design tool.

- Use this description to propose a digital instantiation of the system. This is identical[33] in terms of our framework, but has the characteristics of digital technology, such as flexibility, easy to update and maintain.

- With a description of a proposed digital pervasive system, we modify and extend our system to make full use of the new digital technology.

- Finally, we incorporate new functionality in our system in order to satisfy more needs, namely the needs for signage and direction giving within the hospital area.

This is our initial setup, same as Figure 8.4.

For our analysis, the list of artefacts (❶) that we are interested in consists of: posters, leaflets, reception area, waiting area, corridors, general health information, detailed health information. These artefacts are related as follows (❷❸): posters are used to show general health information, leaflets contain more detailed information, leaflets are placed in the waiting area, and posters are placed in the waiting area, reception area and corridors. The diagram that represents this situation is shown in Figure 8.4.

Our next task (❹) is to propose a digital system that utilises electronic technology, and which maps to the same diagram as the existing situation in the A&E department. To do this, the technology we need would have to create public interaction spaces in public spaces. This could be done by using large plasma screens or video walls. The use of sound could also be explored as a means of reaching a wider audience in terms of accessibility. A good example of how this can be done is in airports, where general announcements are made through speakers in the airport area.

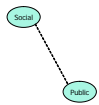To link a public space to the public sphere we can use plasma screens

An interesting issue we can pick up from the diagram is that there exists a link for which our analysis has predicted problems. Specifically, the link between social

---

33. By "identical" we mean that both systems would yield the same description using our framework. Therefore, they would be identical in those aspects that our framework describes, but can (and should) be different in terms of specific functionality and implementation choices.

space and public interaction space poses a problem {I2} for digital technology. Earlier, we have said that such a link prohibits direct physical interaction (see page 124). What this translates to in our situation is that we should avoid using technology that requires direct physical interaction with the users. An example of this would be (❹) a touch screen that requires input in order to proceed with a multimedia presentation. Because the technology would be located in a social space (reception) where access is restricted to staff, patients would only be able to receive information from this technology but not interact with it.
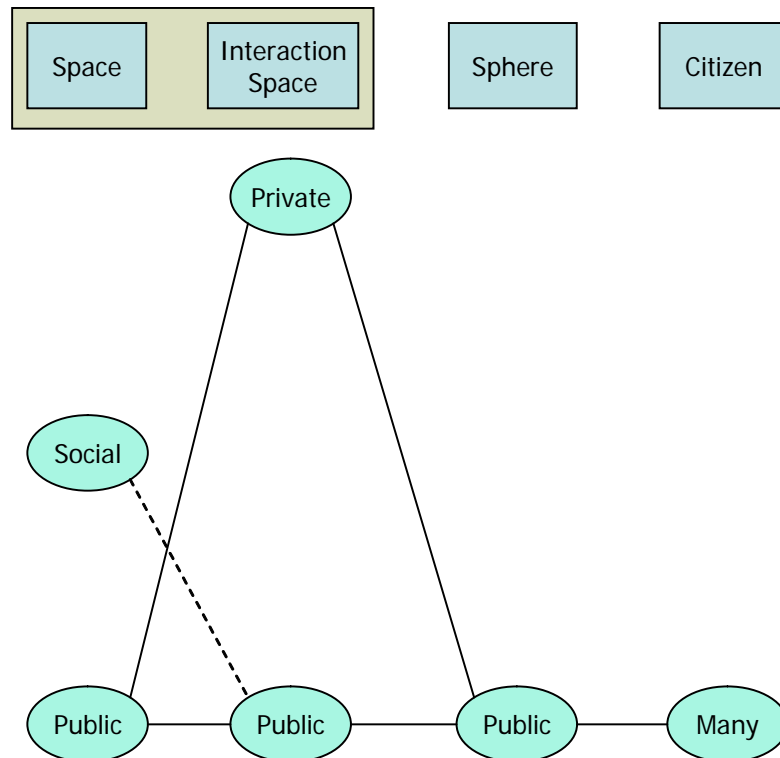
Besides the public interaction spaces, the current setup utilises private interaction spaces in the form of leaflets. It appears that this "leaflet-technology" is a very good design choice, and it would be very hard for a digital technology to surpass it. Leaflets are inexpensive enough that they are provided for free and patients may take them away if they wish. A digital alternative (❹) would involve use of some tech-

Touch screens in the reception are inappropriate {I2}.

FIGURE 8.4: The posters and leaflets setup represented using our diagram.

This diagram represents the existing setup of the A&E department's use of posters and information leaflets to disseminate public healthcare information to patients.
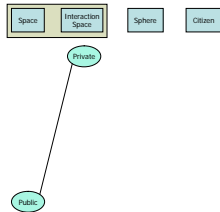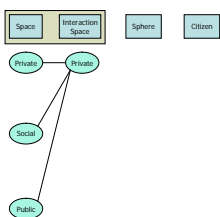
nology that would insulate people {I4} from the environment in such a way that they could focus on the information being delivered and be able to read it properly. Perhaps some kind of small screen attached to each chair (similar to those in passenger aircraft), or even headphones could be used to deliver this information. The new generation of "digital paper" could also be used in this case, although it would have many disadvantages over plain paper and would be better only in providing dynamic, live and easy to update information. However, this begs the question of who would constantly update this information, since hospital staff are very short of spare time as it is. In any case, these digital artefacts would probably be too expensive to be taken away, and would probably have to be fixed to some location to prohibit theft.

For a moment now, let us shift our focus away from spaces and interaction spaces and turn to the social issues involved in this situation. In the existing system we observe that the leaflets and the information they provide are given away for free - all that is required on the part of the patients to gain access to this information is literacy. This is a very good example of how our society values highly health care, and, as a society, we have decided that it would be best for the common good to provide such services free at the point of use.[34] On the other hand, current technology would prohibit a digital alternative to be offered in the same way. Almost any digital solution we offer would either impose (economic) requirements on the users or would not allow people to take the artefact with them.

For instance (❹), the patients' mobile phones and PDAs could be used as a means of providing private interaction spaces to deliver health care information. This option is quite attractive, since private information spaces would be created, and

To generate private interaction spaces we can use small screens or headphones.

The visitor's own private devices can also generate private interaction spaces.


---

34. Although some services are provided free at the point of use, someone has to pay. This is currently done via taxation, in schemes where those with more income pay more taxes than those with less income.

patients could take the information away with them. However, the requirement of owning a mobile phone or PDA is introduced, and this poses a barrier for those who cannot afford such a device. Furthermore, on top of literacy, those who wish to access the information also need to be computer literate and be able to configure network settings and troubleshoot configuration settings. It seems that technology adds so many new requirements, both economic and cognitive, that one wonders whether leaflets will remain the most suitable option.

Another attempt at providing a socially viable solution to giving away information for free can be made by adopting an informational point of view by focusing on the information spheres. Essentially, we want patients to get access to medical information that belongs to the public sphere. Currently, this information is co-located with the technology used to deliver it: printed words on a piece of paper. (Note that this issue of co-location has been discussed in Section 4.3). The alternatives (❹) we have discussed so far adhere to this practice: the plasma screens, video wall, mobile phones, headphones and private screens are all solutions where the information is co-located with the technology used to deliver it. However, this need not be the case.

There is an option (❹) which would decouple the information from the presentation device, and at the same time reduce the requirements, mostly economic, imposed on patients. This option would consist of giving patients a physical link to the information. This is equivalent to giving them a piece of paper with a URL written on it, but technologically this is also possible by giving patients i-Buttons, smart cards, RFID tags, or any other cheap device that can store a simple URL or something equivalent. These devices have been designed to be cheap enough to be given away for free.

Using physical markers (RFID tags) to access the public sphere generates unpredictable interaction spaces.
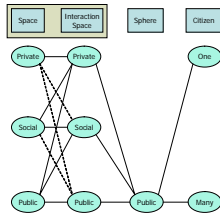
The patients, having been given this physical marker with the link to the medical information, could then go away and use any means they have available to access the information. In some cases this could be private devices like a mobile phone or a desktop computer. Likewise, people could use public devices (like the public library) or semi-public devices (like an Internet cafe) to access the information using the physical marker as a guide to the relevant information. This design option still imposes additional requirements on patients, but at least patients have the option of getting something for free. On the other hand, this design option does not produce the private interaction spaces we have seen in Figure 8.4. It would be rather difficult, perhaps impossible, to predict and trace the types of devices that people would use finally to access the information. Besides, the fact that the information in the original leaflets was delivered in a private interaction space should not constrain us: the accessed information's belonging to the public sphere relaxes further the security and privacy considerations.

So far we have discussed how an existing system (in our case a non-digital system) can be analysed and specified in terms of our framework. We have proposed digital alternatives to the existing system whilst trying to keep to the same specifications as much as possible. Next, we discuss how we can use our framework to extend the functionality that the new system could support, or even add new functionality to it.

Signage and directions

We now address another problem we identified during our study of the A&E department. Once again quoting our earlier work [O'Neill et. al, 2004]:

> "As well as healthcare and health promotion information, there were also directional signs, and notices warning of the consequences of violent or abusive behaviour towards staff and other patients, and about the procedures to follow should anyone wish to make a complaint. There was also a plan of

the hospital. Many of these items were juxtaposed with signs and notices directed mainly at staff, such as lists of names and extension numbers, which tended to overflow from the area behind the reception counter".

Signs and notices and information sheets were fighting for surface space on the walls of the A&E, and generally seem to be very disorganised. This has led to the privacy issues we discussed in Section 8.2, which resulted from the phone lists being posted in the waiting area.

In this section we address the problem of signage by analysing how the system we proposed in the previous section (all versions of it) could be used to deliver direction and signage information to the patients of the A&E department.

To offer adequate directional information, our system would need to be provided in all physical spaces of the hospital. Therefore (❶), this would include the reception and waiting area, treatment cubicles, corridors and toilets. In terms of information, we now wish our system to deliver information from the social sphere. Directions to locations belong to the social sphere because they are not restricted to just one person (so they are not private), but are neither public because outside the hospital they would be meaningless (they are thus restricted by physical constraints as we discussed in Section 4.4).

We have now added additional requirements to our system (Figure 8.5).
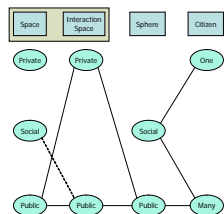
If we use the existing system as it stands (Figure 8.4), and simply use it to deliver the extra information of giving directions to patients (❷❸), then the diagram would look like Figure 8.5. What is new about this figure is that we have added the private space and social sphere ellipses to reflect the new design requirements. The private space has been added because our new requirements included the toilets within the range of the system, and we have classified them as a private space. Furthermore, directions to locations have been classified as belonging to the social

sphere, and that is why we have added it into the new diagram. We have also added the one citizen ellipse to denote that in the privacy of the toilets there will be one citizen within the space itself. We have also added the connectors between the social sphere and the citizens, because we have decided that everyone within the hospital should have access to this information.

However, what we have not done yet is to add the connectors between the private space and any interaction spaces, or between the social sphere and the interaction spaces. This is because doing so would denote a design decision, and we would like to postpone this until after our analysis.

Let us first focus (❹) on the private space ellipse, which is one of the new additions to the diagram. Based on the existing setup, there are two possible interaction

FIGURE 8.5: Adding requirement for delivering signage information.

In this diagram we see the fictitious system we presented in the previous section, but we have added the private space, social sphere and one citizen ellipses to reflect the new requirements. Notice that there are no connectors between the private space and any interaction spaces, nor between the social sphere and any interaction spaces.
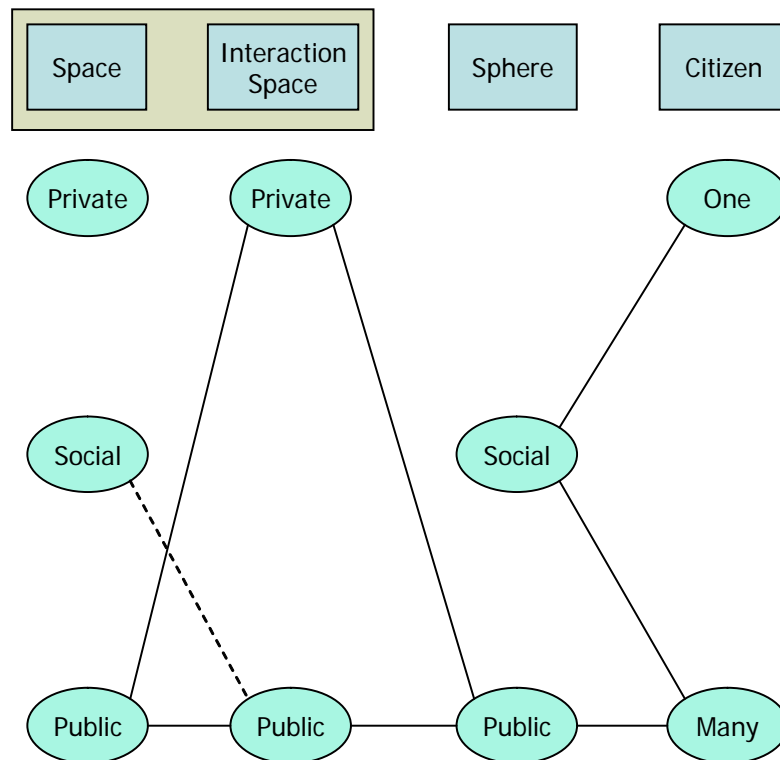
Given the existing setup, we have 2 alternatives for accessing information in private spaces. First, we can use plasma screens.
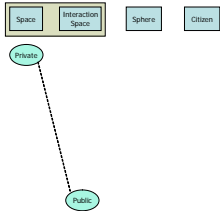
spaces that could be used to deliver the information: private or public interaction spaces. Using public interaction spaces would mean that direct physical interaction with the interaction space and the artefact would not be possible {I2}. To begin with, the plasma displays in the corridors cannot be used because they are not visible from the toilets. However, our public interaction space could take the form of voice messages using a network of speakers. In this case, however, it would be impractical {I3} to support interaction between visitors and our system. We would essentially have to offer the possibility for every visitor to control the messages delivered to everyone, which would result in chaos.
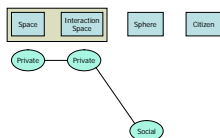
On the other hand, a private interaction space is much more suitable for delivering information in a private space {I2}{I4}. Currently, this is done by providing a hospital map inside the toilet cubicles. A digital alternative (❹) could take the form of a small screen installed in the toilet cubicles, while sound could also be used for accessibility reasons.

We can also use small personal screens.
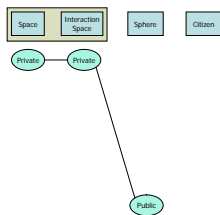
Now let us turn our attention to the social sphere ellipse which is another new addition to our diagram instantiation. As we have already said, the social sphere in Figure 8.5 has connectors to one and many citizens because we have decided that they should be able to access this information whether they are in private, social or public space. What now remains for us to decide is which of the available technologies will be used to access this information.

The small personal screens may also be used to access the social sphere.

Private interaction spaces could be used to access the social sphere information. On page 125 we said that this would hinder the possibility of direct physical interaction with other participants in the social sphere {I11}. In our situation however, direct physical interaction between the participants in the social sphere is not a requirement. Therefore, using private interaction spaces is a viable option.

FIGURE 8.6: Our system now delivers signage information to patients.

Continuing from Figure 8.5, here we present the design decisions we have made in terms of what interaction spaces are made available in private spaces, and which interaction spaces are used to access the social sphere.
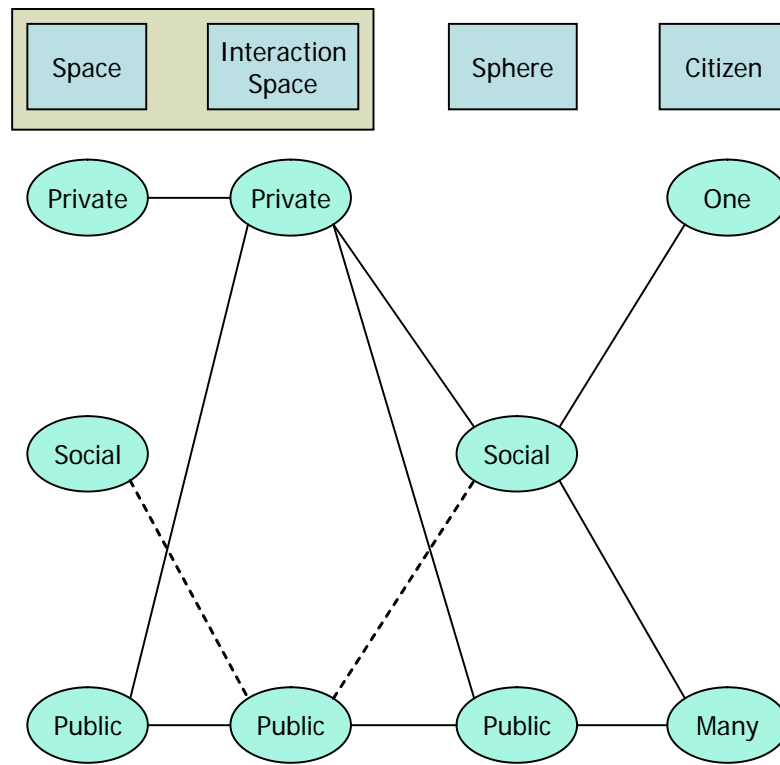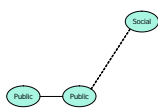

Using the plasma screens to access the social sphere is insecure.

Using a public interaction space to access social sphere information could be an insecure option according to our analysis on page 125 {I9}. However, we have already said that this information belongs to the social sphere simply because it would not make sense outside the hospital area, and therefore (❹) using a public interaction space to access it is not a problem. The resulting diagram instance (❸) of this analysis can be seen in Figure 8.6. Some further details, however, need to be clarified regarding the social spaces in this diagram.

In our requirements for this new extension to our system, we have said that everyone should have access to the social sphere information. Unfortunately, our hypothetical system does not include everyone within a public interaction space: the public interaction spaces are created by posters (or flat screens) in the reception, corridors and waiting area. This excludes most of the social spaces that we have

included in this example, for instance treatment cubicles and treatment areas. Currently, the distinction between the reception and all other social spaces is not reflected well in our diagram, so we decide (❸) to form two sub-groups, and study them separately.

Therefore, to make our analysis clearer, we choose (❸) one group of artefacts to include only reception as a social space, while the other group includes the cubicles and toilets but not the reception area. The diagrams representing the two groups we have created can be seen in Figure 8.7.

First, note that if we superimpose one diagram on the other we will get Figure 8.6. This happens because we have decomposed our system, and are taking different views of it. Also, the diagram on the left can be derived from Figure 8.4 (which represented our original system that had the reception as the only social space) by adding the social sphere ellipse (which represents our new requirement for delivering signage information and directions to locations).

The first thing we notice about the right hand diagram is that it has no connectors to the social space. As we have described and analysed our system so far, we have

---

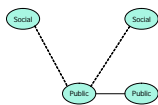**FIGURE 8.7:** Decomposing our system into two groups of artefacts.

Here we have created two subgroups for our system. Notice that if we superimpose one diagram on the other we get Figure 8.6. The diagram on the left has the reception area as a social space, while the diagram on the right has treatment cubicles and toilets as social and private spaces.
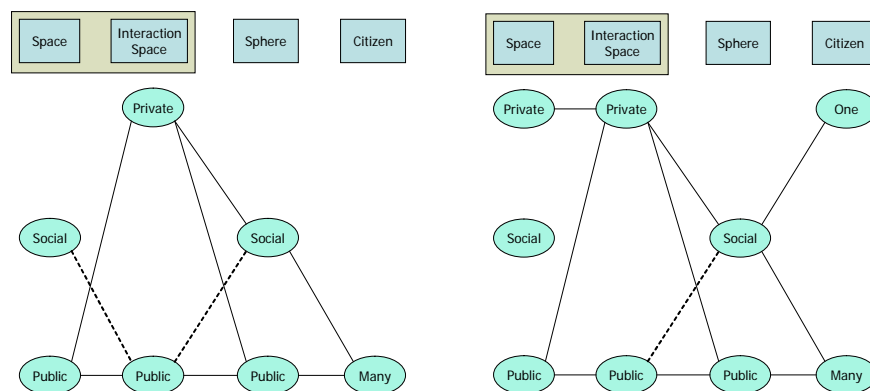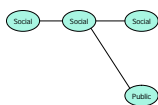
We now turn to social spaces. Based on our analysis so far, we can use either public or private interaction spaces.

not considered how information can be delivered in treatment cubicles and doctors' offices. Looking at the existing setup, we have two options (❹): to use either a public or a private interaction space to access the information. In the diagram we have also included the public sphere ellipse to indicate that our system is already delivering information of a public nature, namely healthcare information. We could have not included the public sphere ellipse if we wanted solely to focus on the signs and directions system, but we have decided to extend our existing system and not introduce a separate system only for delivering directions to visitors, patients and staff members.
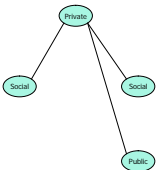
If we use the public interaction spaces created by plasma screens we get

Using the public interaction space is not necessarily an option (❹), depending on the design choices we have made. If, for instance, these public interaction spaces were created by plasma screens on walls, then the public interaction space ellipse should not be in the diagram because these interaction spaces do not span the social spaces in which we are interested. If, on the other hand, the technology we used to create the public interaction spaces in the corridors and waiting area did span the doctors' offices and treatment cubicles (such as with loudspeakers), then it would be an available option. In that case, however, direct physical interaction would not be possible {I2}, and thus interactive technology that requires input from the user (such as a prompt asking "Where do you want to go?") would not be a possibility.

If we use the public interaction spaces created by loud speakers we get

Similarly, the private interaction spaces in the right hand diagram of Figure 8.7 are those interaction spaces that we discussed on page 178, and which we dedicated to providing information in private spaces. As we said, they currently take the form of maps and are placed on the doors inside the toilet cubicles. Furthermore, we proposed (❹) that they could be replaced by small screens attached to the doors. For delivering signage information and directions in social spaces, both of these options are a possibility. We have noted that interaction spaces are ultimately bound by the

If we use the private interaction spaces (small devices) we get

physical and social constraints of spaces {I1}. Therefore, should a plasma screen be placed in a treatment cubicle or a doctor's office, it would create a social interaction space (since it would include more than one person, but not everyone would be allowed to enter that area). Therefore, we are proposing (❹) that a social interaction space be used to deliver the information in the social spaces we are considering (i.e. all of them besides the reception). This could either take the form of a small screen which people in the room could see, or a big plasma screen on the wall {I1} which could be more convenient, since the interaction space would include more people.

If we place the plasma screens in the offices, we will have social interaction spaces.

At this point we would like to recall the design process we have been through (❶❷❸❹), and to do this we can look at Figure 8.8, which includes all the diagrams we have created so far in this section. We have explored various technological alternatives (❹), we have proposed moving technologies into new spaces (❹), and we have also reconsidered the links between our artefacts (❹), and decided to create different groups to analyse the social spaces in more detail (❸).

8.4   Delivering treatment information to patients

We now return to the main problem we identified at the beginning of this chapter: providing patients with information regarding their treatment, such as how long they have to wait and what they need to do next. As we have suggested, the provision of information of this type might be a useful tool not only for reducing stress and averting some violent or abusive incidents, but also in influencing patients' perceptions of satisfaction with their visit.

In designing a system to deliver such information we can better address the requirements we set out, since there is currently no existing system that delivers treatment information to patients in the A&E. Furthermore, there is currently no technology

FIGURE 8.8: A summary of our design process.

Here we present a summary of the design process we have been through in this section. Diagram 1 represents the existing, paper-based system, and based on this representation we discussed possible digital alternatives (page 171). We then wanted to extend the system in order to provide directions to visitors, patients and staff. The requirements for this system are shown in diagram 2, while diagram 3 represents the design choices we made regarding private spaces. We then had to split diagram 3 into two diagrams (diagrams 4 and 5) because we wanted to have a finer look at social spaces by excluding the reception area (diagram 5). In diagram 6 we see the design choices we made for those social spaces.
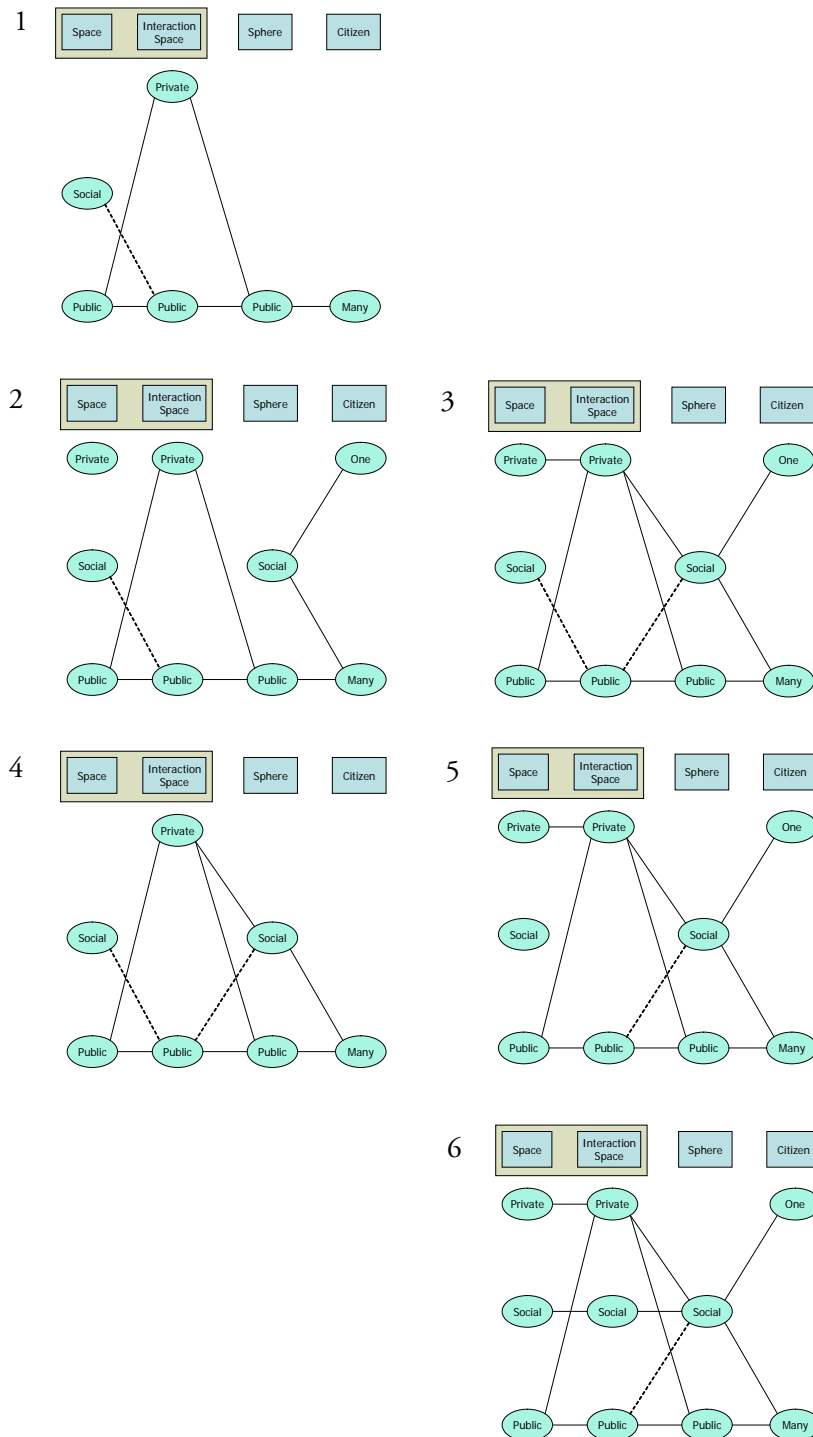
installed which could be altered to deliver this information without adding extra work for staff.

Requirements

Let us be more specific about the requirements for this system. We propose that this system should deliver to patients information about their treatment. This should be in the form of how long they have to wait before seeing a doctor, and why they have to wait for this period of time. Furthermore, the system should provide patients with an overview of their treatment by informing them of what steps their treatment will involve. Finally, this system should inform the patients of the status of their treatment, by telling them what they need to do or what has happened with, for example, a blood test result or an x-ray.

Another requirement has to do with how doctors treat patients. We would like our system to enable doctors and patients to communicate more efficiently and effectively about treatment decisions and choices, thus improving the patients' experience. This should probably be done while the doctor is visiting the patient and discussing with them their treatment.

Artefacts

For this system, the artefacts we consider are (❶):

- the waiting area, reception, corridors and treatment cubicles as *locations*,
- test results, x-rays, appointments, waiting time, treatment status and overview, medical records, as *information*.

These are the initial requirements. How can we best fill the gaps?
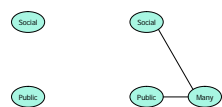
As noted above, we assume (❶) that there exists no technology in the hospital that can be of use to us. This allows us complete flexibility when it comes to interaction spaces, since we are not forced to use any existing interaction spaces (i.e. technology).

The diagram instance (❷❸) that reflects our system so far can be seen in Figure 8.9. The locations we have included in our list of artefacts have been classi-

FIGURE 8.9: The requirements for delivering treatment information to patients.

Here we see a diagram instance that represents the requirements we have set forth on page 184. Notice that there are no interaction spaces since for now we can assume that no technology has yet been installed to support our system.
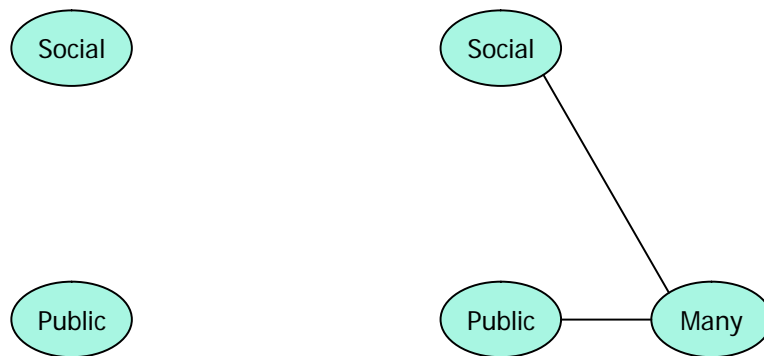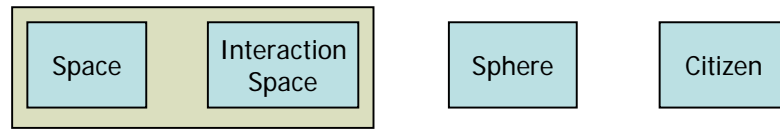
fied as either social or public. Most of the information has been classified in the social sphere, except for the waiting times information which could also be displayed in public as a means of giving an overall picture of waiting times without compromising any non-public information.

Furthermore, because we have assumed that none of the existing technology in the A&E department will be used for our system, the diagram does not have any interaction spaces yet. Designating which interaction spaces we would like to use involves design decisions, which we make next.

Choosing the technology

An obvious candidate for technology (❹) would be the plasma screens we discussed earlier {I1}. They could be used effectively to bridge the public spaces and the public sphere in our diagram. They currently create a public interaction space {I1}.

Plasma screens bridge public spaces and the public sphere.

Therefore, our first design decision (❹) is to use plasma screens deployed in the waiting area to display public information about waiting times. This information is of a public nature, and should not compromise any non-public information regarding any of the patients {I4}. For example, the plasma screens could show how many patients are waiting to be treated, the average time that each patient has to wait, and the urgency of the injuries of the patients currently waiting to be treated.

We have noted that patients were not aware why some people were being treated before them or had to wait for less time. This is the result of the hospital placing a higher priority on patients with urgent needs. Our system should relay this information and inform patients that some of those waiting have been given higher priority due to the nature of their injury or other factor such as age. Conveying this information does not necessarily have to convey the identity of the patients with higher priority, nor the nature of their injury. A way of doing this would be to include a message that informed patients that x number of patients have been given higher priority.

Using loud speakers we identify {I2}{I3}.
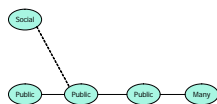
Another choice (❹) of technology suitable for delivering the same kind of information is loudspeakers. A network of such speakers could be used to make periodic announcements about the status of the queue and the average waiting times. Those speakers would create a public interaction space in both public and social spaces {I2}{I3}.

Now we want to provide access to the social sphere from public spaces.

Our next step (❹) is to choose a technology that would allow us to deliver personalised information to patients about their waiting time and treatment status. We have classified this information in the social sphere because the nurses and doctors are aware of the information - in fact they are the ones generating it. However, this information should be delivered to specific patients only, not everyone in the waiting room. Therefore, the public interaction spaces that the plasma screens create

The plasma screens are not suitable {I9}.
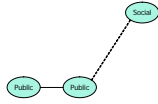
Leaflets, small screens and PDAs will do.
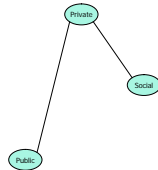
Using a kiosk to access the social sphere (i.e. blood test results) {I4}.
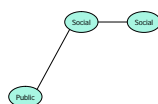
(❹) are not suitable in this case {I9}. To deliver this personalised information we could (❹) use a private interaction space {I10}. Thus, our discussion about leaflets, private screens, mobile phones and PDAs in Section 8.3 now becomes relevant. All those design choices would provide a private interaction space in which we could deliver personalised treatment information to the patients. In this case however, mobile phones and PDAs are not an attractive choice since not everybody has one, and thus not everybody would be able to receive personalised treatment information. On the other hand, private screens would be an effective, yet expensive solution.

Another option we could explore (❹) is the use of a physical marker to specific information. This option was explored on page 174 as a way of delivering public information to patients. In the present example, however, (❹) physical markers such as ibuttons or RFID tags could be used to gain access to personalised treatment information. This could be done by using a social interaction space created by something analogous to a kiosk (❹). The patients, having been given such a physical marker or tag could walk up to these kiosks and place their tag on a reader. Doing so would bring up the treatment information that is relevant to them. Although the use of private screens is the preferable {I10}, it is much more expensive than the kiosk option we have just described. The downside of the kiosk option is that they are susceptible to information spillovers, since kiosks do not offer complete privacy {I4}. Furthermore, these kiosks could cause a bottleneck if a large number of patients wanted to use them. However, they are a more economic option than installing numerous private screens.

So far we have discussed the technologies used in public spaces only. Our design decisions up to this point can be seen in Figure 8.10. The two diagram instances in this figure represent the two possibilities we have discussed so far: using private

FIGURE 8.10: Our design possibilities regarding technology in public spaces.

The left diagram instance represents the use of private screens to deliver personalised treatment information to patients. The diagram instance on the right represents the use of physical markers and kiosks to deliver the same information.
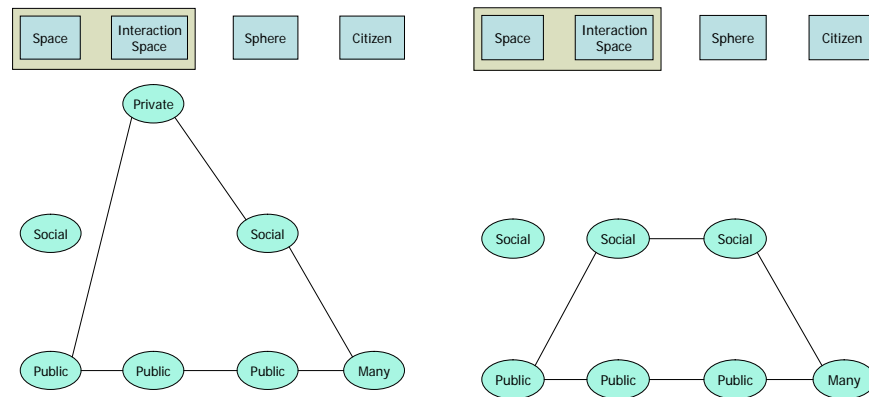
interaction spaces (by means of small personal screens or phones) and social interaction spaces by means of kiosks and physical markers.

By looking back at our analysis of the connectors in Section 6.2 we can infer a number of issues regarding the two different design choices. Looking at the left diagram instance, we are reminded that the connector between private interaction space and social sphere alerts us to the fact that this setup will not support direct physical interaction between the participants in the social sphere {I10}. In other words, using an insulating interaction space (such as a small screen) is not suitable for situations where many participants who are accessing the social sphere wish to interact directly with the device and the user of the device. These are situations analogous to the patient and doctors discussing test results and further treatment. Furthermore, as we have noted, because the private interaction space exists in a public space, care needs to be taken so that the technology does not allow information spillovers {I4}. In this case, this is important since the information being accessed belongs to a social sphere, and thus has a higher privacy requirement than information from the public sphere.

On the right diagram instance again we observe the need for insulating technology {I4}. In this case, the information being delivered is the same, so information spillovers need to be avoided. There are various implementation options (❹) for making a kiosk more insulating. For example, special glass screens can be used, similar to those used in ATMs, which allow the information to be seen from within a confined angle of vision. Another option (❹) would be to manipulate physical space, and simply make the kiosks' screens face towards the walls so that passers-by cannot see any information, or even introduce some kind of curtain or screen to the same effect.

An insulated kiosk is now more appropriate.

Next, we explore our design choices for technology in social spaces. The only social spaces we are considering (❸) are the treatment cubicles, where doctors and patients get together to discuss treatment options, and where nursing staff care for the patients. Currently there is no technology installed to support these collaborations. Instead, traditional speech-technology is used, i.e. the doctors and patients discuss the treatment plan. Privacy is secured in the following simple way: the doctor asks everyone in the cubicle to leave, so that doctor-patient confidentiality can be preserved. Asking people to leave a location is really an example of manipulating the physical space, and manipulating the links between locations and people.

We now explore our options for technology in social spaces.

In developing digital technology to support these interactions, we can start by investigating the options (❹) we have for delivering the information from a social sphere in a social space. Currently, we have public interaction spaces being used to deliver public information, but these interaction spaces are inappropriate for this new situation {I9}. Let us now consider (❹) using private interaction spaces. Using such a technology would insulate the user from her environment {I4}, which is not entirely appropriate for a collaborative task {I10}. We have said that we would like patients, nursing staff and doctors to use this technology to deliver bet-

Loudspeakers are inappropriate {I9}.

Private interaction spaces (small devices) do not support interaction between the participants.
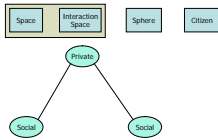
A shared display is more suitable.
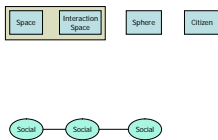
Using both private and social interaction spaces requires switching between the two (see Chapter 7).
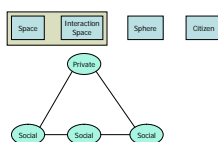
ter treatment to patients by allowing them to discuss treatment options. It would seem (❹) that a social interaction space is much more suitable {I10}{I12}. However, we need to acknowledge that in the process of the patient treatment, the doctors and nurses might need to access information that should not be seen by the patient. An example of this would be if the doctor decided to look at the records of patients with similar symptoms; the doctor might wish to check her schedule to see when she would be able to come back. We see therefore that more than one social sphere is being accessed by the doctor, and only one of them is the one that the patient is entitled to access.

Creating only social interaction spaces is not the solution in this case either. We have said that if the one-to-one relationship between social sphere and interaction space is broken, then resources are fragmented and we risk compromising information privacy {I11}. What we see as the best solution (❹) in this situation is to use both private interaction spaces and social interaction spaces. The setup could consist of the technology creating a private interaction space to be used by the doctors and nurses. They could use this technology to access information that should not be seen by the patients. Some form of personal device (❹) seems suitable for this {I10}. Additionally, some kind of technology (❹) could create a social interaction space, including the doctors, nurses and patients present within a cubicle. This technology would deliver information that the patient might want to see, such as possible treatment options, health status and health targets. Some kind of screen, such as a plasma screen (❹) would be suitable for creating such an interaction space {I1}{I9}.

An issue which becomes apparent in this setup is that the doctors or nurses would have to switch between the private interaction space and the social interaction space, and possibly control both. This is a potential application of the stroke inter-

FIGURE 8.11: The final version of delivering treatment information to patients.

In the first diagram instance we can see the complete system. We created two separate sub groups (diagrams 2 and 3) based on physical location. In diagram 2 is related to public spaces, while diagram 3 shows our design choices for social spaces.
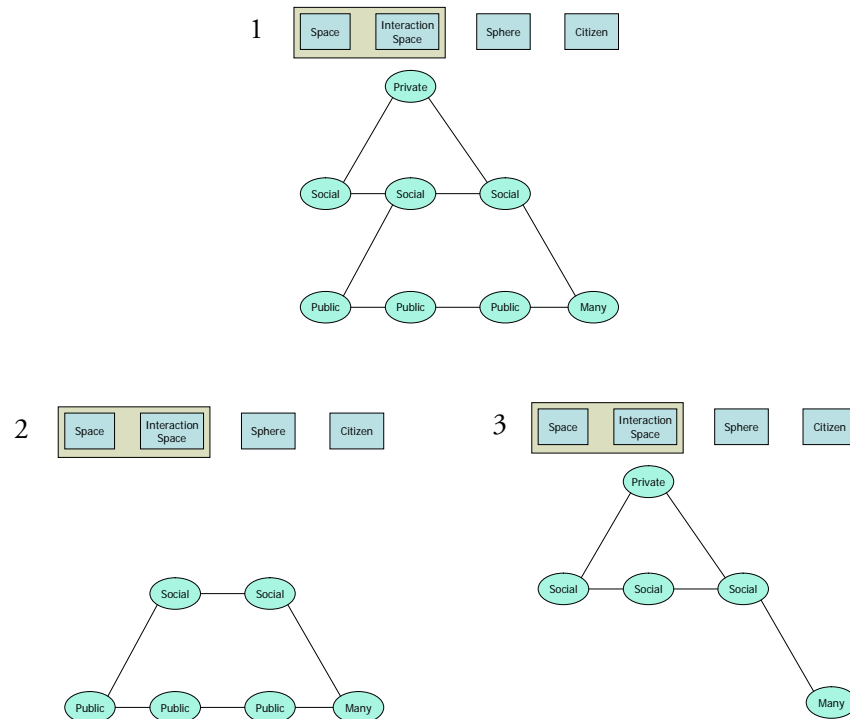
action technique we have described in Chapter 7, which would allow for the same interaction strokes to be used on the personal devices as well as the social interaction space.

The design decisions (❹) we have made in this example can be seen in Figure 8.11. In the first diagram we show a summary of our proposed system, which we have broken down into diagrams 2 and 3. Diagram 2 gives an overview of the system based on the public spaces involved in the system, while diagram 3 does the same for social spaces.

8.5   Summary

In this chapter we have shown how our ideas, design tool and method can be used to help us in evaluating, designing and representing pervasive systems. First, we

described some of the findings from an ethnographic study of a hospital, and explained that by providing information to the patients we could mitigate the violent incidents that often occur in the A&E departments of hospitals.

We demonstrated our design tool in action (post hoc), first with the simple example of the A4 sheets of phone lists posted in the waiting area, and explained why they result in a breach of privacy {I9}. We then went on to show how our method could guide us towards different design options and choices, and how we can use it to propose changes to our existing system. Specifically, we showed how the privacy issue could be addressed by using different technologies to deliver information.

We continued with the more complex example of the public healthcare information which is delivered using leaflets and posters. We described the existing setup using our design tool (post hoc), and then proposed a similar technology-driven system. We then used our design tool a priori to extend this system, and show how separating out subgroups of artefacts can help us focus on specific aspects of the system. We explored various technological alternatives, and tackled the problem we predicted in extending this system to deliver signage information to patients and visitors.

Our last example addressed directly one of the main issues that arose during our ethnographic study: providing personalised information to patients to help ease their wait, and to improve their treatment experience. Again we explored alternatives, and showed how various technologies could, or could not, be used for the kinds of situations we were designing for. As a secondary point, we also showed how our work on stroke recognition in Chapter 7 ties in with the kinds of designs we discussed in this chapter.

Throughout this case study we have addressed issues and ideas that we have discussed as part of our framework. We discussed the issue of providing access to information for free (at the point of consumption), which is a trend that characterises services of the public sector. From our analysis we also identified ways of representing and reasoning about conventional technology (such as posters and leaflets) in conjunction with digital technology. We also discussed how we can separate the information from the medium used to deliver it as a way of thinking about design alternatives. In terms of spaces, we explained how making changes in the physical environment can work to our advantage, and should therefore also be considered as design alternatives.

There are many more issues and ideas that stem from our framework. We feel that the large number of ideas we discussed in Chapters 3, 4 and 5 should not be seen as a barrier to a simple, straightforward design approach. Rather, our designs can become enriched if we take account of all the ideas we have presented. Design options and approaches become available, or apparent, when one considers the design of a pervasive system from the different viewpoints (i.e. by focusing on one of the three elements of our framework).

The A&E department we studied in this chapter was a rich, yet simplistic, environment where we could apply our ideas. This simplicity has helped us in explicating our framework, design tool and method. For this case study we drew on the findings of our ethnographic study.

What remains to be demonstrated after this case study is a way to generate design recommendations which can lead to design. We have done just that in the following chapter, where we show how we can study a location, apply our framework at a higher level and generate recommendations for a pervasive system.

CASE STUDY: THE CITY OF BATH

So far we have discussed our ideas and framework in the context of a relatively small environment, such as the A&E department of a hospital. In this chapter, we apply our ideas and framework a priori to a large-scale setting. To achieve this we use the city of Bath in the UK as a case study. In this setting we make predictions and explore the possibilities for public pervasive systems within the city of Bath.

We remind the reader that our previous case study, and all the analyses of various design alternatives, was based on two simple recommendations: first, that providing digital posters, leaflets and directions could save staff time (see Section 8.3), and secondly that by delivering personalised treatment information we could alleviate some of the problems we reported (see Section 8.4). We then followed through

each of these recommendations by using our design tool and method. Here, we present many more recommendations and predictions that could similarly lead to a long and fruitful exploration of design alternatives and solutions. Following this through to completion is beyond the scope of this chapter. However, the previous case study serves as a model of how this can be done and demonstrates the results of such a process.

There is a potentially large number of issues that could be addressed when looking at a city as part of a case study. Even when confined to the design of pervasive systems, there are still many issues, problems and findings that could potentially be of interest to us. In an attempt to avoid overstretching our case study, we have tried to limit the issues which we felt to be of direct interest to the rest of this thesis. We have therefore approached this case study with the following enquiries in mind:

If a pervasive system were to be designed for, and installed in the city of Bath,

- What types of services would it offer? Domestic pervasive systems are most appropriate for providing in-house services and functionality to the visitors of shops, tourist sites and restaurants. How could a public pervasive system make our lives better if it does not provide such tailor-made services for each site, but rather offers services to the city as a whole?

- What physical form and shape should it take? How can we apply our analysis of space to the design of pervasive systems?

- How can a pervasive system strike a balance between civic needs and commercial drives? Public pervasive systems are expensive to design, install and maintain. Who is going to pay for them? In the case of Bath, how could the local government and commerce find a mutually beneficial recipe for this system? To address this, we need to analyse the civic needs that could potentially be served by such a system.

- How would the different locations of the city affect and be affected by the pervasive system? As we have said in Section 5.5, a pervasive system can be regarded as a set of digital artefacts, or as a part and extension of the physical environment. How can this duality help us analyse how our system affects the physical environment and architecture, and vice versa? Furthermore, how can we know if these effects are wanted or not?

These are the questions that underlie this case study. Although Bath is a (geographically) small city, the richness and complexity of this city, and almost certainly of any other city, prohibits an holistic analysis within the scope of a single thesis chapter. Thus, we have decided to analyse in detail a small central part of the city rather than remain vague in the analysis of the whole city.

In Section 9.1 we focus on the central part of Bath which is at the heart of our case study. In this section we show photographs and maps of the specific locations in which we are interested, and describe in detail the various sites and situations that we have identified as being of importance to our case study. Specifically, we pinpoint on the map all those locations that are of interest to us, give some general information on these locations, along with photographs, and also discuss some general requirements that cannot be pinpointed on the map but rather apply to the whole of the city of Bath.

In Section 9.2 we analyse the important sites and situations we have identified in light of the main lines of enquiry of our case study. For each location and situation we analyse what potential services could be used in either the location or situation, what effect they would have on the form of the pervasive system, how they would impact and be affected by architecture, and how the element of civil needs balances against the consumer and market interests for the specific location or situation.

We are interested in the design of a *public* pervasive system. Therefore, what we are *not* proposing in this case study is various systems to be used for specific locations such as a restaurant. In the beginning of this thesis we have emphasised that specific locations and situations (like shops, tourist sites, etc) can and should implement a small-scale pervasive system to suit their specific needs and aspirations. What we *are* interested in is how the various sites we examine contribute to the design requirements of a public pervasive system. If all of the sites we examine are to be covered by and included in our public pervasive system, then the variety, richness and complexity of those sites should be reflected in our design rationale.

9.1   Our region of focus - central Bath

For the purposes of this case study we focus on one of the central locations of Bath. As we show, there is a rich diversity of locations and events within this region. In Figure 9.1 you can see an aerial picture of the location we study. Our area of focus is around a pedestrian street, which is crossed by a number of roads. This street is indicated by the arrow in the figure. A number of tourist sites are in this area, along with a large number of shops, restaurants, pubs, banks, bus stops and open spaces where street performances take place. Figure 9.2 contains a simplified map of this area. The numbered locations in this map correspond to the following locations:

• Location 1: The *Roman Baths*, which is a tourist site. This is one of the main attractions of Bath. Visitors must buy a ticket to enter this site of ancient ruins, while local residents are entitled to free entry. The main entry to the Roman Baths is from the west side of the block.

• Location 2: A *public space with benches* where sometimes performances take place. This open space is often used by street artists to host their performances. Sometimes musicians perform here, while people are sitting on the benches. This is

FIGURE 9.1: The area of our focus.

This aerial photograph shows the area of our focus. The arrow shows the main pedestrian walkway that is our focus, and also happens to be the rough South-North direction. The numbered locations correspond to the same locations in Figure 9.2.

also a location from where tourists take photographs of the Abbey which is to the east of this location.

- Location 3: A *clothes shop*. This is a shop that sells clothes. There is nothing special about this location in terms of technology or anything else that might make it attract attention. This shop has a shop window which is used to display items on sale and promotions. The doors of the shop are kept open so that the inside can be seen from the street. People may enter freely into this shop, but the shop manager reserves the right to deny entry to any customer.

- Location 4: A *public space* which is sometimes used for performances. This is another open space where street performances take place, but there are no

FIGURE 9.2: A simplified map showing the area of our focus.

This map shows the area that was depicted in Figure 9.1. The top of this map points towards the north.
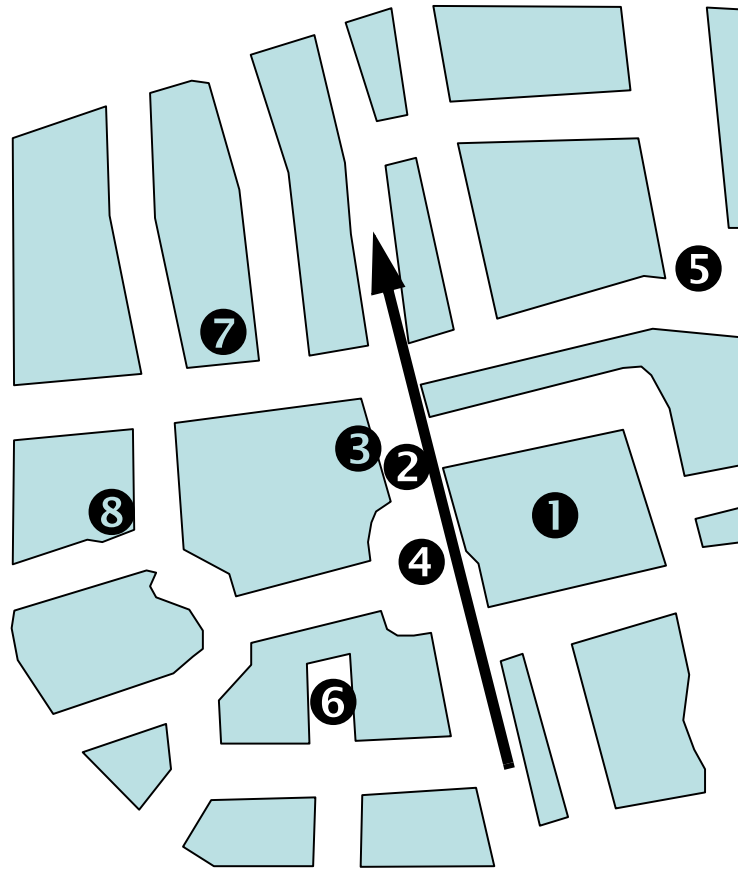
benches here. Under normal circumstances, people simply walk past this space, since there is nothing to attract their attention. The entry to the Roman Baths is located here, but it contributes very little to attracting the attention of those who pass by.

- Location 5: The *bus stop*. In fact, there is a number of bus stops all located along a 40 metre sidewalk. A number of buses make a stop here, and a very wide range of destinations is served from these bus stops.

- Location 6: Local residents' *parking area*. Cars with a special parking permit on show may park here. To obtain a permit, one must live in the area and also pay a relatively small annual fee. No other car is allowed to park, not even for an hourly fee.

- Location 7: A *pub*. Like with most other pubs, entry is free, but customers must be over 18 years of age to enter this pub. Furthermore, the pub reserves the right to deny entry to any customer, and this is usually done with customers who have previously caused trouble. The presence of security staff at the entrance reinforces this.

- Location 8: A *cinema*. People of all ages may enter the cinema (depending on the film), but only upon paying a ticket. Like almost every privately-owned shop and company, the managers of the cinema may deny entry to customers.

The above sites were selected for the diversity of their uses, which would enrich our case study findings. In addition to the specific locations that we have discussed so far, there are some more general "situations" or needs which could be seen as applicable to the whole city of Bath. The ones we are interested in and have analysed are:

- *Providing civic information*. A huge amount of civic information is currently made available via the Bath city council website.[35] The information on this website is related to the environment, learning, council announcements, housing, tourism and local events such as concerts, gallery exhibitions and special museum collections. This information could also be made available by utilising a city-wide pervasive system. This service follows directly from our discussion in Chapter 3 of citizens and citizenship rights. It is therefore quite appropriate for a pervasive system to address these issues by providing civic information to citizens.

- *Providing navigational information*. This category of information would allow people to find their way to various locations by asking the pervasive system for "directions". This would also enable people to enquire about specific types of locations such as restaurants, cinemas and parks. We believe that such services

---

35. http://www.bathnes.gov.uk

can best be addressed as part of an holistic approach to the design of pervasive systems, and not as isolated services. Specifically, we have described the duality of views (see Section 5.5) that we can take on pervasive computing. Thus, we may consider the system as giving directions and information about parts of the city, or giving directions and information about itself. Furthermore, navigational cues have been explored in architecture, and this work should be utilised in the design of pervasive systems. In fact, all the lessons we have learned from our discussion in Chapter 5 may be drawn on here.

- *Helping people meet each other.* Building on the provision of navigational information, we would like to add a social element to our system. This can best be explored by attempting to help individuals coordinate their activities by assisting them in meeting each other, forming groups and socialising. Although the provision of navigational information could be an integral part of this, it is not essential. For instance, local residents may not need to get directions to a pub where their friends are. Rather, they may simply require the name of the pub. A number of different possibilities could be explored with such a system, both for local residents and visitors in Bath. In both cases, however, we are primarily interested in the element of sociability and supporting the formation of groups, which we have seen to play an essential role in the success of urban environments and cities (see page 105).

In the next section we discuss in detail each of the eight locations and three more general needs, and outline the important points that we have considered for each of them. We relate this discussion to our original points of enquiry stated on page 196, but also relate our findings from each of these points to every other one. In the process of doing so, we see emerging an overall picture of our envisioned public pervasive system.

9.2   Detailed analysis

We now proceed with a detailed analysis of each of the interesting sites and situations that we have listed above. Throughout this analysis we reference our previous case study (Chapter 8). This is because a number of situations that we describe here are similar to situations we examined previously. Rather than replicate our work by presenting even more design diagrams, we reference our previous work and let the reader decide whether to go back and look at our previous discussions.

In our analysis here we raise interesting points relating to our four main lines of enquiry. In Section 9.3 we use our analysis in this section to make a full recommendation for a public pervasive system. Therefore, for now we simply raise important design considerations without necessarily offering a direct solution to every one. We then offer such solutions once we have analysed all the locations and situations. In this way, we can develop a more coherent approach to the design of a public pervasive system for Bath, rather than offer sporadic suggestions and recommendations which would be difficult to put together in one system.

Location 1: The Roman Baths

This location may be seen in Figure 9.3. What is shown in this photograph is the main entrance to the site, which is on the West side of the block. This site has restricted access, in the sense that visitors must buy a ticket to enter it. Local residents are allowed to enter for free, so long as they can prove that they are indeed residents of Bath. This can usually be done by showing a driving licence, an electricity bill, a bank statement, a university card, or any other "official" document that proves that one lives in Bath. Because of these restrictions on access, we classify this as a social space.

When considering the space element, we must be aware that this is a protected site. Therefore, any pervasive system that is deployed here should have an absolutely

FIGURE 9.3: The Roman Baths (location 1).

The Roman Baths (location 1) as seen from outside the main entrance to the site. The entrance is located on the West side of the building.

minimum impact on the artefacts of this site, since they are of archaeological importance.

In terms of information, we believe that a public pervasive system could provide the opening times, ticket prices and historical information for this site, all of which belong to the public sphere. Because of the public nature of the information we wish to deliver, a public interaction space is appropriate {I2}{I7}{I9}.

Delivering public information to public spaces is best achieved with public interaction spaces.

There is a wealth of additional information and services that could be provided in the Roman Baths. A domestic pervasive system employed inside the site could provide the services and functionality of traditional tourist guides and museum guides. This could entail providing information for every artefact, allowing visitors to record their visit and put it on the web, or leave messages and have conversations. This could be done in an engaging, content-rich interactive way. All these services could be delivered by a domestic pervasive system, and so are outside the scope of the system we are designing.

Note that by visiting and seeing the site, in principle we are accessing information from the social sphere. (Here, by information we mean the physical representation of the statues, baths and artefacts, as well as all the knowledge that has gone into developing the site). In this case, one has to be inside the building to access this information, and such an interaction space within the building is either social or private. Because we proposed to use a public interaction space, it would have to physically be in a public space {I1}, and thus outside the boundaries of the building.

If we examine Figure 9.3 closely, we can see that there are information boards on each side of the main entrance which serve the purpose of creating a public interaction space. These posters contain information about the opening times of the site, as well as an overview of what the site has to offer. This is very similar to the "poster technology" we examined in our previous case study (Section 8.3). The diagrams and design rationale that we have previously employed in order to explore just how to design digital counterparts of posters can be replicated here.

As we have noted, local residents may enter for free because this site belongs to the local community. Therefore, the public pervasive system we are designing could be used to validate if a person is indeed a resident of Bath and thus is entitled to free entrance. This type of information could belong to the private sphere depending on what proof of residence is used (i.e. one's home address could be regarded as private information, while one's University of Bath student card belongs to the social sphere). Therefore, we need to pay attention to how this information is given for validation to the system, and what happens to this information afterwards: is it stored or discarded?

In summary, let us try to relate our discussion of this site to our original points of interest regarding our public pervasive system (see page 196). The Roman Baths are a social space with restricted access and visitors need to buy a ticket to enter.

Citizenship rights allow local residents to enter for free, upon proof of their residency. Therefore, private information provided by visitors (such as proof of residency) allows them to enter the site for free. The physical space itself is of archaeological importance and should not be altered by the pervasive system we are designing. Furthermore, the interior of this site would have minimal effects on our system, since we are addressing the public aspect of this site, not the interior aspect of it. The interior would be best served by a specialised domestic system. Finally, the information that we wish our system to provide in relation to this site is of a public nature (ticket prices, opening times, historical facts). Thus, it should ideally be delivered by a public interaction space. As we noted, such spaces are already created by traditional "poster technology".

Location 2: Benches

This is part of a pedestrian street where a number of benches have been installed. These benches naturally form a gathering place for visitors, and this in turn attracts merchants who bring stands to do their business. During quiet times (such as late night hours and early morning hours) this street is empty. Conversely, busy hours see a flood of pedestrians and merchants engaging in all sorts of activities as we can see in Figure 9.4.

This is a completely free open space, and anyone may "enter" or walk through this location. Furthermore, anyone can use the benches. This is therefore a public space. We note that merchants and visitors create social and private spaces in the sense that a merchant's till or a visitor's purse is a private space. However, at the level of our analysis it is more useful to consider this location as a public space.

To understand what information relating to this location may be useful, we point out the different activities that are going on here. In photograph 1 we see a newspaper stand which sells a local newspaper. In photograph 2 we see a number of benches, aligned back to back, which are occupied by people who are resting, chat-

FIGURE 9.4: The bench area (location 2).

This is a location where a number of benches have been installed (location 2). The photographs here depict a number of stands which are brought by merchants and some billboard maps which show the city of Bath.

ting, and having a snack. We also see a poster stand, with a logo of the Bath city council. This stand provides information about local attractions on one side (picture 2) and a map of central Bath on the other side (picture 4). We also see the attempts of the council to make this a more enjoyable location by adding hanging flower baskets (pictures 2 and 3). Finally, in picture 3 we see a stand that sells handbags and other small clothing items (left) and a stand that sells frames and pictures by a local artist (right).

This location, with the set of activities we have just described, is a rich and vibrant environment that provides a lot of ideas for our pervasive system. For instance, the presence of the newspaper stand informs us that there is an apparent commercial opportunity, with enough demand and supply for accessing local news. The presence of other stands denotes that this location may act like a small market, although

this can change: one moment the stands are there, the next moment the stands are gone. The adaptability of this location is something we have noted on page 106, and we need our pervasive system to keep up with the changes that are taking place.

In addition, we note the efforts of the council to provide residents and visitors with information about local events and attractions, as well as navigational information. This information belongs to the public sphere, and the posters that are delivering it create a public interaction space. This information should be provided by our system, and most probably this should be done using public interaction spaces {I2}{I7}{I9}. In addition to this effort, the council has also made attempts to make this location aesthetically pleasing. We need to consider just how we can make our system aesthetically pleasing, both in terms of direct physical presence as well as the experience it provides. In relation to this experience is the fact that this location is a natural gathering place. People tend to gather here, have a chat and enjoy a snack or an ice cream from the nearby shops and stands. Our system should provide such services and functionality that would assist people in this tendency to gather, meet, and enjoy a conversation.

Once more we encounter public interaction spaces. We shall refer to this setup as "poster technology".

Location 3: A clothes shop

This is one of the purely commercial locations in our case study. This shop is located on the pedestrian street we examined previously and its main function is to sell clothes to customers. In Figure 9.5 we can see what this shop looks like from the outside. The shop window is used to display items on sale and promotions. The entrance doors of the shop are kept open so that the inside can be seen from the street. This may perhaps cause customers to be inclined to enter the shop, a pull that is enhanced by the absence of brick walls facing the outside, such that the internal activity of the shop is quite visible to the outside.

People may freely enter the shop, but the shop manager reserves the right to deny entry or service to any customer. Therefore, we classify this location as a social

space. Reinforcing this classification is the fact that the manager or shop owner has complete control over what happens in this shop, how it is internally arranged, what colour the lights are, and so on. Furthermore, the fact that this location may be "closed" indicates that it does not offer unbarred access to its customers.

We now turn to the informational aspects of this location. From our analysis so far, we can deduce certain categories of information in relation to this location that could be provided by our public pervasive system. It is obvious from the arrangement of the shop that the owner wishes to make certain bits of information public. Clearly, the shop promotions and prices are something that the shop advertises in public. Another publicised aspect is the internal activity which we have seen to be clearly visible from outside the shop due to the absence of brick walls and the fact that the doors are always kept open. Further relevant information includes the opening times of the shop, which are advertised by a small poster on the entrance of

FIGURE 9.5: The clothes shop (location 3).

This is a primarily commercial location. Notice the absence of brick walls on the side of the shop facing outside. Additionally, the doors are kept open at all times. These two features allow people outside to see what is happening inside, the products on sale, and the current promotions.

Shop-owned interaction spaces.
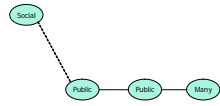
Poster technology.

the shop. Finally, general announcements regarding the operation of this shop may be of interest to the general public. Such announcements could for instance inform that the shop will be closing in a few weeks, changing management, or even changing business. We therefore {I2}{I7}{I9} propose that public interaction spaces be used to deliver this information. However, we need to take into account that this is a social space, and just like our earlier example of the hospital reception offering a public interaction space (see Figure 8.4) {I3}, issues of direct physical interaction with the system may arise. To address this, we could create these interaction spaces in public spaces.

Location 4: Street performances

This location is a prime example of how spaces may change and adapt the function they offer. As we can see in Figure 9.6, this location is an open space. It is located outside the main entrance to the Roman Baths (location 1). What is interesting about this location is that although most of the time people simply walk past, it is also a gathering point for performances. A number of street performers use this

FIGURE 9.6: The public performances area (location 4).

This is the paved area outside the Roman Baths where street performances take place.

location as their venue. The circular shape of this location allows street artists to gather attention by urging those who are passing by to form a circle around them.

As we saw in Chapter 5, activities and performances contribute to an enjoyable experience of a public space. An interesting feature street performances is that they can be for free.   As one street artist has said,[36]

> "Street performances are the most democratic type of art. Nowhere else can you watch a performance and then decide if you are going to pay for it. You could say it is like try before you buy".

The same person claimed that city councils offer him money to perform on the streets of their city and in events like festivals. We argue, therefore, that street performances, which are a great source of entertainment, contribute to the image of a city and quality of life and are free, should be supported by a public pervasive system.

The most common characteristic of all street performances is that they happen in a public space. This ensures that whoever wants to attend can do so {I7}, and also helps in attracting a greater audience. A second characteristic which can be found in many street performances is that the audience participates {I3}. Some members of the audience may be asked to help by clapping, holding an object, or being subjected to the wit of the artist. Conversely, in some cases the audience simply enjoys the performance of a violin player or a mime without having to take a more active role. The third characteristic found in almost all street performances is the act of donation on behalf of the audience. The end of the performance usually initiates the donation process, whereby those who watched the show may give money to the artist. This sometimes happens during the performance, such as in the case of the

---

36. Personal communication on July 20, 2003, with a Canadian street artist who called himself Bill.

violin player or the mime. The final common characteristic, which we have already noted, is that street performances usually take place in non-designated public spaces. By this we mean that the space which hosts the performances serves different purposes when the performance finishes and everyone leaves.

We have now introduced a new set of requirements for our public pervasive system. In street performances we have found a situation where public spaces are used to host activities where the audience can participate along with the main artist, and where a donation mechanism is used to compensate for the artist's work. Finally, this can happen in non-designated areas or in areas that serve additional purposes. For our pervasive system, therefore, we need public interaction spaces {I2}{I7}{I9} (interaction spaces that are in public spaces), we need to support two-way communication and enable participants to physically contribute (thus other types of interaction spaces do not qualify), and we also need a mechanism whereby participants may financially contribute {I3}{I12}. Finally, we need to take into account that any public space which serves another primary function may be used for a street performance.

Poster technology.

Location 5: The bus stop

In addition to the bus station, located to the south of the city centre, this location serves a large number of passengers. As we can see in Figure 9.7, there is a forty metre pavement where signposts have been distributed. Each signpost indicates which services make a stop at that particular point. On some signposts, but not all, there is a concise schedule for the services that stop there.

This location has many common elements with other similar locations in other parts of the city or even other cities. The purpose of a bus stop in most cases is well-defined and understood. Because they offer a specified set of requirements and needs, numerous projects have attempted to create "smart" bus stops. The results have been deployed in real world instances, such as in Bristol, UK. Typically,

FIGURE 9.7: The bus stop (location 5).

This is where bus services make a stop to drop passengers and take on new ones. There are a number of posts along the side of the road which indicate where each service stops. On some, but not all, signposts there is a concise schedule for the respective bus services.

"smart" bus stops allow the passengers to see what buses are imminent and when they will arrive. The functionality they offer is similar to the functionality of the London Underground digital signs and the National Rail digital signs placed in train stations. In a sense, bus stops have and are being addressed by the research community. This is exactly why we are *not* proposing specifically how to design "smart" bus stops.

Since Chapter 1 we have noted that individual projects have been implemented to address some of the issues we are examining. In this instance, the needs and requirements for bus stops have been well studied. However, we have also stressed that a top-down approach has *not yet* been developed in the design of these projects, and that this causes many of the problems within pervasive computing. In this case, the design of "smart" bus stops has had very little impact on the design of other small-scale pervasive systems. As testbeds, such bus stops have helped in assessing and evaluating certain technologies in real-world situations where they have been installed. As we have argued in Chapter 1, little more has been gained.

Unfortunately, the understanding of designing pervasive systems has not noticeably been advanced by the design of enhanced bus stops.

Coming from such a perspective, how can the study of this location be of use to us? We believe that this location, just like every other one in this case study, has a lot to offer to the *whole* system and to our conception of the requirements for a public pervasive system in Bath. Just as we have done with the locations we have already analysed in this case study, we indicate those aspects of this location that are of interest to us and could potentially impact our design of a larger system. To be precise, we are interested in the four main themes we set out on page 196 and how they are affected by each of the locations we study.

In terms of space therefore, we find that the bus stop has no designated area. The only aspect of space that makes this a bus stop is the fact that buses actually stop here, as well as the presence of posts which give information telling us that this is a bus stop. We note that other bus stops offer shelter from the weather in the form of a small roof, which doubles as an information board. Therefore, if the information gets delivered in a different manner, all that would be left of the bus stop is the physical shelter. Taking this one step further, we could argue that locations for other purposes can offer shelter (such as coffee shops, parks and other public spaces), and therefore a bus stop completely loses its physical element. For the time being, however, the bus stop is placed in a public space.

In terms of information, the passengers usually retrieve information from the time-tables, quite often after a specific enquiry (i.e. a two-way interaction rather than a one-way feed). The timetable information is of a public nature, since anyone is granted access to it, and therefore it belongs to the public sphere. This type of information can therefore be delivered in a public interaction space without privacy concerns. Furthermore, timetable information does not always reflect the actual

times that the buses arrive. Therefore, there is a distinction between delivering the information of the timetable (which is static) and delivering information about when the next bus is actually coming (which is live). The latter is preferable for the passengers, although the former can be helpful in planning one's journey.

Delivering bus information to where the passengers are.
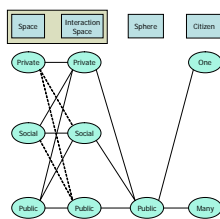
In addition, we need to consider our argument that the physical element of a bus stop is due to the sheltering requirements and not to delivering the information. This implies that social and private interaction spaces could be used to deliver the bus schedule information in locations where the *passengers are, not where the bus stops are.*

Finally, we wish to ensure that the information remains intact despite its public nature. Therefore we wish to limit the amount of possible interactivity between the passengers and the information, so that enquiries can be made but not changes {I9}. This issue is currently addressed by enclosing the posted timetables in a protective sleeve, such that people cannot change the information they provide (although they can still obscure it).

Location 6: Local residents' parking area

This is another location where, just as at the Roman Baths (location 1), the citizens and residents of Bath have been given special privileges. In Figure 9.8 we can see what form this takes. This area is signposted with notices indicating that only local residents may park here. Each car which is parked in this area needs to bear proof of residency, which takes the form of a special permit. This permit, which is made available by the city council only to local residents, must be placed on the windscreen or other clearly visible part of the vehicle. Any car which is found parked in this area without having such a permit is removed and the owner is charged a fine for retrieving their vehicle.

As we have discussed earlier (see page 205) a public pervasive system could be used to validate the credentials of a car or of the person driving the car. If a car bears some form of permit, it is allowed to park here. The difference from our earlier example is that now we wish our system to be restrictive. Therefore, validating authorised cars is only part of what we wish to do. Ideally, we would like the system to identify unauthorised vehicles and perhaps inform the parking wardens of the presence of such cars.

In our previous discussion on producing proof of residence, we said that some of the information involved could be private, and this would cause us to think twice about how this information is produced and validated. In this case, however, this issue is not relevant because the only form of proof that can be used is the special permit. This existing mechanism for authenticating cars serves two purposes: First, local residents must pay a certain amount to obtain the permit (therefore by enforcing the use of permits the council has found a way to generate income). Secondly, it now becomes easier for both a parking warden as well as a driver to resolve the issue of whether a car is legally parked or not: there can be a large number of docu-

FIGURE 9.8: The local residents' parking area (location 6).	This location offers special privileges to local residents. This is a parking area where only local residents may park their vehicles. A special permit which drivers place in their car is used as proof of authorisation.
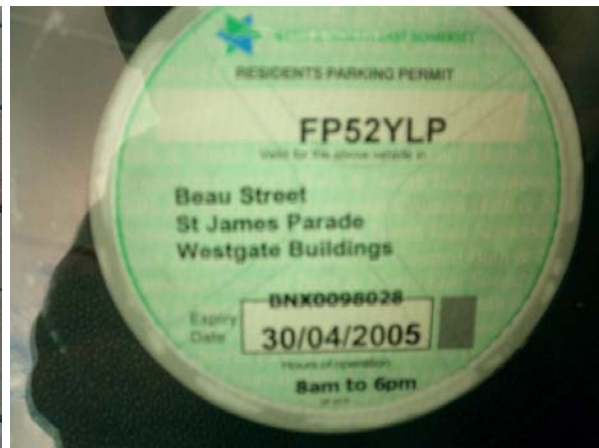
ments which could potentially prove residence, some could be ambiguous for the parking warden, some could compromise private information of the car owner. This leads us to the issue of interaction spaces.

Currently, the credentials used to park a car in this area (i.e. the special permit) are produced in a public interaction space: a car's windshield is a public interaction space so long as the car is in a public space {I1}. Therefore, we can assume that the same can hold true in the case of our pervasive system. As we have said, the existing mechanism does not compromise personal information, and so we could digitally extend it while allowing it to retain this characteristic.

Analysing this location in terms of space reveals some very interesting details. This area is a public space. Yet, the whole point of our analysis so far has been about how we can restrict people parking here. Why, then, is this a public and not a social space? Informally, we could say that because anyone is ultimately allowed to *walk* here, this is a public space. More formally, we can look back at our original discussion on public spaces (Section 5.2). We have said that public spaces belong to the community and are filled with norms and expectations. In this case, we could regard the act of parking by "outsiders" as something unacceptable as well as illegal. On the contrary, a social space (such as the clothes shop we discussed earlier) would ultimately be controlled by someone who can dictate the rules and norms within the space itself.

Displaying a parking permit in a car. The lack of direct physical interaction is appropriate.

Location 7: A pub

We now turn to another commercial location, similar in some ways to the clothes shop we studied earlier. Here, the location we are focusing on is a pub. The exterior and main entrance of this pub can be seen in Figure 9.9. As with most other pubs, entry is free, but customers must be over 18 years of age to enter this pub. Furthermore, the pub reserves the right to deny entry to any customer, and this is often done with customers who have previously caused trouble. The presence of

FIGURE 9.9: A pub (location 7).

This pub allows entry to customers aged 18 or more. The posters and boards on the walls deliver varying types of information, while the three large windows and the open door allows passers-by to glance at the internal activity of the pub.

security staff at the entrance reinforces this. Therefore, we categorise this location as a social space.

Looking at Figure 9.9 we can see that posters and information boards have been placed outside the pub to inform about special promotions within the pub. This information is intended by the pub owners to reach as many potential customers as possible, and is therefore delivered in a public interaction space.

An interesting point we can make, which is of relation to quite a few of the previous locations we have studied, is that although this information is intended to be "injected" into the public sphere, it is also intended to be left intact. We therefore need to be cautious in making sure that such information is not changeable; the fact that it belongs to the public sphere means that others have access to it, can comment on it, or can reference it. The two-way communication with the public sphere (see page 79) is intended as a means of exchanging information and ideas. As we will discuss on page 231, however, this begs the question of how we can

avoid flooding by those who have more resources, skills or motivation than others. We address this issue in the last part of this chapter.

Pub-owned interaction spaces, and delivering the information to where the customers are.
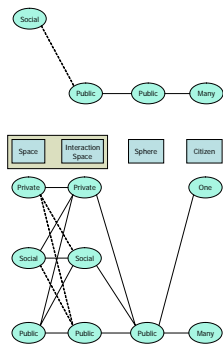
Besides the fact that aspects of one's citizenship (i.e. age, criminal record, etc) affect entry to this location, the rest of the informational aspects of this location are similar to the clothes shop (location 3). Again we have here some information that the pub may wish to advertise: opening times, special promotions, the state of internal activity, and perhaps general announcements. The interaction spaces most suitable {I2}{I7}{I9} for delivering this information in this location are public, although we need to be aware that if they are within a social space then issues of limited direct physical interaction arise {I3}. Social and private interaction spaces could also be used to deliver such information, and their use could be explored for situations where potential customers are elsewhere {I6}{I8}{I10}{I12}.

Location 8: A cinema

The last location we examine is yet another commercial location. This cinema, shown in Figure 9.10, offers only one screen. During a day more than one films is

FIGURE 9.10: A cinema (location 8).

This is the view of the cinema from outside. A number of posters attached to the walls are visible. These provide information about screenings and ticket prices.

screened, and typically each day the cinema has four to five screenings. This cinema has very similar requirements and restrictions to any other cinema. Customers must have a valid ticket for the specific screening, and be old enough (depending on the movie) in order to enter the screening area. Customers may buy tickets in advance, something which is not available in all cinemas. In this case we have a social space due to economic constraints (as opposed to citizenship constraints in the pub).

Poster technology.

Regarding informational needs, we can see that a number of posters and boards outside the cinema deliver advertising and screening information. These create public interaction spaces since the area outside is a public space. In this case, however, the internal activity of the main screen room is not visible from outside. What can be known, upon an enquiry, is the number of free seats available in the screening room (some customers may not wish to view a movie in an empty or almost full cinema).

Like the pub (location 7), the cinema would wish to advertise certain information using our public pervasive system. This would include opening times and screening times, information about the films being shown, ticket prices and special promotions. Public interaction spaces could be used for delivering this information in public spaces {I2}{I7}{I9}. Social and private interaction spaces could be explored for delivering this information directly to customers {I6}{I8}{I10}{I12}.

Civic information

Besides looking at specific locations for our case study, it can be worthwhile to examine issues that cannot be discussed in relation to any particular location. The first of such issues we discuss is the provision of civic information to the citizens of Bath. As we described earlier, the city of Bath council has commissioned a web site to provide civic information aimed primarily at the residents of Bath.

The information available on this website touches upon many different aspects of daily life in Bath. It relates to the environment, housing and tourism, general council announcements such as debates and decisions made, events taking place in the area such as gallery exhibitions and festivals. By providing this information on a publicly accessible web site, the council of Bath has deemed this information as being of a public nature. This website is a way of publicising such information, but also of getting citizens involved.

This theme of encouraging citizens to participate is being researched quite vigorously. Specifically, e-democracy and e-participation are major research areas within HCI [Carroll & Rosson, 2001] as well as within Computer Science [Watson & Mundy, 2001], Sociology and Politics. Being of such importance and relevance to everyday life, democracy and participation should have an impact on the design of a public pervasive system in Bath. Although this approach would appear to be the opposite of what other researchers are pursuing (essentially the investigation of how technology can impact democracy and participation), we argue that pervasive technologies are not developed to stand and exist for their own sake; rather they are envisioned as becoming assimilated with everyday life and everyday activities. As such, they are susceptible to influence by whatever is relevant to everyday life. Therefore, since civic information is of importance in relation to everyday life, then our public pervasive system should be designed accordingly.

What, then, can we say about the provision of civic information in relation to the main lines of enquiry of this case study? First, we have already identified that civic information belongs to the public sphere. In terms of services, therefore, these depend on which specific information is being delivered, whether it requires two-way communication between citizens and the public sphere or amongst citizens themselves. The range of possible services in relation to civic information is vast; a

glance at the list of topics on the council's website gives a good indication of this. All we can say at a high level in relation to services delivering civic information is that they are of a public nature.

Next, we can consider what we discussed in our previous case study in relation to placing (or removing) obstacles in the way of getting to the information (see page 173). This is of direct relation to the form and shape of our system. Ideally, the civic information under discussion should be accessible by all citizens {I7}. Currently, it is accessible only by those who have access to the Internet and the website where the information currently resides. In deciding upon the form and shape of our public pervasive system, we need to consider the obstacles in the way of people who wish to access the information. First and foremost, have we placed economic obstacles in the way? Are people required to own and operate a mobile phone, PDA or other type of device to take advantage of our public pervasive system? These obstacles can be straightforward to remove, in comparison to other more complex ones, such as accessibility: are all social groups able to access and use our public pervasive system?

We should bear in mind that these obstacles need not necessarily be the result of technology itself. In merging the physical spaces of architecture with the interaction spaces of technology, we need to take account of the obstacles that each of these two components has introduced or removed. The physical form of our public pervasive system, the locations where people come in physical contact with it, the way it integrates with the built environment, are all issues that could result in either creating or removing obstacles.

Finally, we need to keep in mind the issue of balancing civic and commercial needs and interests. A public pervasive system flooded by civic information is no better than an equivalent system flooded by advertisements and promotions. Both civic

and commercial needs are part of our lives. This decision and these values really depend on the local people who will actually benefit from the system. But the need to find a balance between the two remains.

Navigational information

Another situation that is not directly related to any specific location in a city, but rather poses an holistic requirement is the provision of navigational information. By navigational information we refer to information that helps people find their way to various locations in Bath. People could enquire about specific locations (by giving a specific address), or about types of locations (by giving specific requirements such as "Italian restaurants within half a mile"). Rather than duplicate the efforts of many research projects looking into the design of tourist guides and digital map applications which explore such functionality and services [Brown & Perry, 2002], we wish to provide an alternative approach. Our approach is based on the duality of views that can be adopted when considering a public pervasive system (see Section 5.5). In terms of providing navigational information, this duality of views leads us to view the system as either giving directions and information about parts of the city (thus distinguishing and distancing itself from "the city" and promoting the view that "the system" is just a set of digital and physical artefacts), or giving directions and information about parts of the system itself (thus enforcing the view that it is embedded in "the city"). Certain aspects of the system could be designed using the former approach, whilst other parts could adopt the latter.

Delivering the information in private and social spaces.
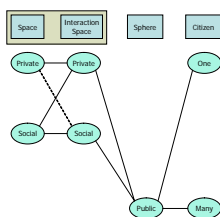
Let us give an example of how these two different approaches may be adopted according the needs of a situation. In private and social spaces, we might want to promote the feeling of ownership and control of the space by those present within it. To do so, it would be appropriate for the system to distance itself from "the city", and act as a set of artefacts. For example, within a car it would be appropriate to provide navigational information without necessarily bringing in other aspects or

services of the pervasive system. On the other hand, getting directions within a public space could potentially be a situation where both the physical and digital spaces are susceptible to navigation, and thus along with the exploration of physical space one could explore digital space as well.

Furthermore, parts of our discussion in the latter sections of Chapter 5 have relevance to supporting navigation within a city. Landmarks, both physical and digital are of potential benefit to those who wish to navigate physical and digital space. In terms of the form and shape of our public pervasive system, we should ask just how it eases and supports navigation. The built environment could make the navigation of digital space difficult (i.e. noisy and busy streets make it hard to use the system, or a specific location disorients a person who then cannot make use of navigational information provided by the system). The opposite also hold true, i.e. that the digital spaces can make it hard to navigate physical space (trying to navigate the system requires too much focus and attention to be able to walk around at the same time, or the navigational information provided is confusing and people make wrong turns while driving).

Finally, once again commercial interests have to be balanced against civic needs. In terms of navigation, businesses want people to find them easily and thus have motive to make full use of this system. On the other hand, people need to be able to find parks and streets without necessarily falling victim to spam advertising or other unwanted intrusions.

Helping people meet each other

In the final part of our analysis we extend our previous discussion of navigational information with a social element. A good example for achieving this is to explore how our system could help people in forming groups, friends in meeting each other, and as a result to help sustain social activities within the city. The provision of navigational information is needed to pinpoint locations of interest and get

directions to them. Additionally, however, we now need to convey information about the status and activities of a location, and the whereabouts of specific individuals (i.e. "friends").

The effects and impact of architecture have been suggested throughout our discussion of specific locations in this case study. We have described how shops and pubs wish to publicise their internal status and activities by keeping doors open and having big windows through which people may look inside. Should our public pervasive system make use of such principles, then it could also assist in getting people to meet others and organise into social groups.

Locating friends is researched within Computer Science [e.g. Benford et al., 2003]. In our case, we are interested in how this functionality blends in with the rest of the activities that go on in a city, as well as the rest of the public pervasive system that supports these activities: locating friends, along with having access to information about specific locations and navigational directions to them, coupled with a focus on supporting locations that host activities, is a recipe for successful public spaces according to the projects we surveyed in Chapter 5.

Furthermore, we have surveyed locations which offer themselves as gathering places within central Bath. Locations such as 2 (the benches, see page 206) and 4 (street performances, see page 210), and the respective functionality and coverage by our public pervasive system, can assist in creating and sustaining social activities.

9.3   The Bath public pervasive system

We have now analysed various aspects of central Bath, and are ready to make our recommendations for a public pervasive system. Before we proceed, we would like briefly to recall the variety and range of issues we have discussed so far in this case study.

In this case study we have looked at locations that have a specific purpose, locations whose purpose and activities may change, and locations whose purpose is replicated across a number of similar locations in the city or other cities. We studied public and social spaces, also over-arching needs which are not confined to one location. Of the public spaces we studied, some are static in the activities they support, others adapt, whilst others support multiple activities simultaneously. In respect to the social spaces we looked at, some had social limits as to who could access them, others posed economic constraints, and others posed citizenship constraints. On this last point, we also examined locations where citizenship would actually remove economic obstacles (by giving privileges to local residents) or conversely place obstacles to those who did not meet certain criteria. Finally, the over-arching needs we discussed dealt with the provision of civic information targeted mainly at the citizens of Bath and the provision of navigational information targeted at anyone who might need it. In an attempt to socially extend the latter, we discussed the possibility of our public pervasive system supporting the creation of social groups and activities by helping people meet each other.

In this last part of our case study, we present various recommendations, each of which may be followed through to derive concrete design solutions. We now discuss what we have learned, and make suggestions in respect of the main enquiries (see page 196) that were of interest to us in approaching this case study. The recommendations we present here are not meant to be complete: still many issues remain untouched on the road to actually building and installing a public pervasive system in Bath. However, what we present here is a result of our analysis of central Bath, based on our design framework. Therefore, these ideas and issues are not arbitrary or random. They are underlined by our general approach embedded in our design framework. Furthermore, as we showed in the previous case study, the ideas that we discuss here can be followed through from specification to design and

implementation by adopting the more low-level approach of our design tool and method.

The available services

A theme that we have revisited more than once is the provision of credentials on behalf of the people. It would therefore be reasonable to conclude that such a service would be useful since it addresses requirements. This service would ideally allow people to prove that they are who they claim to be. This would involve proving their residency, their age, and perhaps any characteristic they might wish to prove. To be able to do so, a mechanism for acquiring credentials must be implemented. We would also expect third parties to react to the absence of required credentials (such as in the parking area). Furthermore, in implementing such a service, we need to make sure that privacy is not compromised either by means of spillovers or by other security breaches.

A second service we propose deals with the provision of civic information. This would take the form of delivering local daily news as well as information regarding the environment, jobs, events, and the rest of the information that currently resides on the Bath city council website. This information forms part of the public sphere, so design decisions should be made accordingly.

The next service we propose deals with supporting commerce, and underlies many of the locations we studied. This service would allow someone to make a payment to another person. This functionality is already available with conventional systems in the form of on-line banking. Furthermore, it has been studied in the context of specific locations such as restaurants [Kindberg et al., 2004]. From our case study, however, we saw various examples of locations that could use this service, such as the pub and the cinema. Additionally, this service would support locations that may be temporarily transformed into a commercial venue, such as the merchant stands or the comedian's hat. A public pervasive system, carrying the name and

authority of the local council or government, could make use of this authority to encourage and support these types of transactions.

Another service we propose deals with accessing public information which may be dependent or independent of location. By this, we mean information such as bus times, navigational information, the location of friends or the promotions of a pub. Although this sounds very similar to having a PDA with Internet access, there are subtle differences in what we are proposing. We are not proposing that because all the information is available it should be immediately accessible to the user. Furthermore, we are not proposing that the system gives a list of "available services" or "available information" to the user. Standing in the middle of central Bath, the available services and information, even in the form of a list, would still be overwhelming. We are proposing that interaction spaces be used to filter out the information that gets delivered to the user. This means that if the user is within the interaction space created by the pub, the cinema, or the street artist, then the corresponding information is delivered. This may still be overwhelming, but it is a first step, and when coupled with context awareness or direct user input, it could prove helpful. Of course, the user may still wish to have access to information which is location independent (such as the news, bus times, or even the promotions of a pub on the other side of town). Even in this case, interaction spaces could be used to identify which information should not be delivered. Finally, this distinction could also be helpful for those providing the information. The pub, for example, is aware that the information it provides may (or may not) allow for direct physical interaction - something that our design tool can predict based on spaces and interaction spaces {I3}{I10} - and could thus deliver the information in an appropriate manner.

Our next recommendation does not take the form of a service, but rather gives us ideas about potentially numerous services. In our discussions on successful public

spaces (Section 5.4) we highlighted the need to support activities and socialisation in public spaces. We also explored this issue in this case study, and examined how a relatively asocial navigational system could be extended to provide support for social groups. There are still many potential ways of providing support for social groups with a public pervasive system. For instance, we identified the location with benches as a natural gathering point. We could create "digital" gathering points, or even extend physical gathering points with functionality to support what people do there. This could involve some of the services we have already proposed, such as digitally saying who you are, or buying stuff to eat from the merchant.

A final, but very important recommendation has to do with the ability of the public pervasive system to keep up with the changing functionality that a location might host. This example is highlighted by the bench area as well as the street performance area. Those locations change the functionality they support, and our pervasive systems needs to keep up with these changes. Of course, this can also be the case in social and private spaces. What we suggest is that locations are provided with a form of context awareness, such that the services that are available are those that best support, or even describe what is happening within the location. Therefore, we introduce a second filtering layer of available services within a location. We do not imply that the services that are filtered become blocked; we are simply proposing that priority should be given to the services that pass through the filters.

The form of the pervasive system

At the beginning of this case study we raised questions about the potential physical form and shape of our pervasive system. We are now in a position to comment on this issue.

The first location we studied, the Roman Baths, is a location where absolutely no physical change may take place. In this sense, therefore, our public pervasive system is to remain invisible within the Roman Baths. This is also to be expected of

social and private spaces (such as people's homes), where physical change by the government would be inappropriate (imagine the local council trying to install a camera or a screen in every home!). Furthermore, the city as a whole is peculiar in the sense that very few alterations can be made to buildings (as we saw in the case of the clothes shop). At first, these findings might suggest that a public pervasive system needs to remain completely invisible. This, however, is not necessarily the correct approach for a number of reasons.

First, we have noted the need for providing navigational clues throughout the city. These could be either physical or digital. In the case of physical clues, we could use existing landmarks, or we could deploy new landmarks such as signposts. In either case, we need to provide a design that efficiently takes advantage of physical landmarks, and is able to merge effectively with them.

Another reason for externalising and physically representing our public pervasive system has to do with the use of public interaction spaces. By definition, a public interaction space needs to be accessible from public spaces. To create public interaction spaces, we need technology that does this, and which is currently quite visible (such as screens and projectors). Even if the technology for creating such interaction spaces were miniaturised and made invisible, still the public interaction spaces, which are part of the system, would persist. Therefore, the presence of public interaction spaces means that the public pervasive system cannot remain invisible.

A final argument is that the system needs to place, as well as to remove, obstacles to accessing the information. Picking up from our discussion on page 222, we have said that our system needs to address the issue of placing or removing economic, social or physical obstacles. In the case of physical obstacles, this could involve locking doors, opening doors, raising or lowering bars in the parking space. There-

fore, this is another way in which our pervasive system becomes represented and visible in the physical world.

Most of the ideas we presented in Chapter 5 in fact discuss how to make the presence of our system "felt" in such a way that our lives can be improved. Navigation, gathering points, safety and accessibility, all of which are themes that we see being repeated in suggestions for urban space design, result in physical representation.

Poster technology and delivering information where the citizens are.
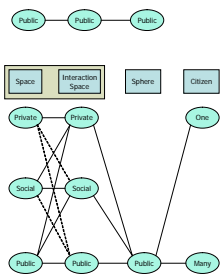
The final point we would like to make on physical representation is that in our case study we have seen many times some design requirements being repeatedly addressed with the same solutions. Such patterns are for instance the poster solution for delivering information to the public, or the private interaction space for accessing remote information. These have implications for the physical representation of the system. More important, however, is the fact that they have similar representation in terms of our design tool, and this could be a first step towards a common representation language for various systems as we describe in the final chapter.

Civic needs vs commercial interests

It would be simplistic to assume a public pervasive system could provide information for all shops in the city without giving the perception that the system has been flooded by commercial interests. What could be potentially worse, however, is something we have mentioned in the opening chapter. The perception of commercial flooding could undermine people's confidence in the integrity of the system or information. Why should this public pervasive system be yet another advertising opportunity for businesses? Yet businesses would claim that they form a core part of our way of life, our culture and our behaviour.

The key to resolving this argument is understanding what it is fundamentally about. We believe that it is the result of tension between civic interests and com-

mercial interests. This friction between the commercial and the civic is very vivid at the physical border or boundary that separates the shop (any shop) from the public realm. For example, shop managers are aware of rules and regulations, set by the council or government, regarding the types of items that shop owners may put out on the street. The fact that the shop we have just analysed kept the doors open is no accident. It is an attempt to overcome the limitations set by the council, which dictate that shops cannot flood out onto the street and take over the public realm. Another interesting example, especially in Bath, is the fact that many buildings are deemed protected, and thus must conform to the traditional golden stone look of Bath. Shop owners cannot be creative and paint the outside of their shop any colour they wish. We see that the way most cities have chosen to address the issue of buildings' appearance by first declaring that the public realm is a very important part of everyday life, and then assuming that commercial drives, if left unchecked, could overtake the public realm. The result is this model of controlling shops and merchants by adopting rules and regulations governing the looks of a shop and the types of items it brings to the public.

The friction between the conceptual boundaries of the civic and commercial within our public pervasive system could also be resolved by adopting similar laws and regulations. When translated into the terms of our framework, we are essentially proposing that the information spheres that are accessed by our system, and the spaces and interaction spaces involved, could be regulated so that the public sphere and public interaction spaces are not overrun by eager-to-contribute commercial agents. This would of course involve discussions with the local community, as well as an understanding of the values of the wider community.

Finally, we have already made clear our preference of a public service model for the design of public pervasive systems. We have said that by adopting this model we

could promote the use and familiarity of the system. Furthermore, it could prove to be a catalyst in getting people to participate (aspects of e-democracy and e-participation which we discussed on page 221). Finally, this model is appropriate for introducing and enforcing the regulation and control we discussed above.

Effects between city and system

This last part of our analysis deal with how we can assess the impact of our system. There are a number of questions we can ask of the designers of the system, the local community, and the commercial enterprises.

To assess the impact on public space, we could ask how we have made public spaces more enjoyable. Has our system provided security, comfort and support for social activities in public spaces? Is our system itself secure and accessible, and does it directly support social activities? Is it able to adapt to the dynamics of locations? Does it improve on the variety and diversity of choices that people have? Overall, we could refer to the issues we discussed in Section 5.4.

On the commercial front, we could enquire as to how local businesses have benefited. Do they disseminate information to more people? Are they easier to find and locate? Has commerce in general benefited by, say, increased transactions? Or is it the case that our system has had a debilitating effect?

To balance the commercial issues, we also need to consider the civic issues. Have citizens become more involved and aware of local issues? Can they find the information they are looking for? Do they feel that rules and regulation are being adhered to, or is the system constantly abused? Has the community become more active, and are public events supported? What about the presence or absence of obstacles? Do people have ready access to the system we have designed?

Thinking about and answering the above questions is not an easy task, especially once our pervasive system is in place and being used. However, it highlights that a

number of potential problems lurk about our system besides the obvious technical problems. It can be easy to concentrate on the day-to-day running of the system, making sure the software and hardware operates smoothly. However, the issues we have discussed here are also of importance, and could ultimately lead to the failure of our pervasive system.

9.4   Summary

In this chapter we have presented a case study that examined at the city of Bath as a potential setting for a public pervasive system. We analysed a priori various locations in terms of our design framework as well as certain specific enquiries we had in mind. This case study should be seen as complementary to the study we presented in Chapter 8. In the hospital case study of Chapter 8 we showed how design ideas and recommendations could be translated into design alternatives that can be tested and explored. What remained to be demonstrated after the previous case study was a way to generate these design recommendations. We have done just that in this chapter. Here, we showed how we can study a location, apply our framework a priori and generate recommendations and predictions for a pervasive system. These recommendations may then be translated into design options which can be explored as we showed in the previous case study.

CONCLUSION AND FUTURE WORK

10.1 Summary

In our work, we have sought to contribute to the sadly sparse body of theory within pervasive computing, but in a way that can provide tangible results as well as assistance to those who will actually design and implement pervasive systems. Because of the nature of pervasive computing, many different domains had to be drawn upon for a comprehensive and coherent theoretical understanding of such systems. From the start of this work we claimed that pervasive computing is still in its infancy. We offered our own definition and vision of what pervasive systems could offer us in the future.

Our starting point has been established theory and practice in HCI, which we have attempted to extend. The results have been the creation of a design framework, the implementation of an interaction technique, and the development of a design method and tool for pervasive systems design. The concepts we presented as part of our framework have been applied on three different scales: systems level (Chapter 7), design exploration and evaluation (Chapter 8), and design recommendations (Chapter 9).

10.2   A framework for pervasive computing

The main thrust of the work presented here evolves around the framework we have developed from previous work in HCI. In Chapters 1 and 2 we explained that the domain of pervasive systems requires a new approach to and understanding of design. Our survey of related research and practice showed that currently the notions of user, task and domain are not adequate for pervasive systems. We went on to describe our own alternative concepts, citizen, sphere and space respectively, as a way of sharpening our focus on the important issues.

We explained each of our terms in separate chapters. In Chapter 3 we proposed the notion of citizen as a new way of thinking about the people who will be using our system. We reviewed some of the relevant work in sociology and political science in order to elicit those aspects of citizenship that are of relevance to our work. In our analysis we found that free access to information, a widely supported freedom associated with citizenship, is a key right which transforms pervasive systems into public services. In line with our overall vision for pervasive systems, we discussed public services and pointed out important characteristics that can offer design ideas within pervasive computing. Our discussion of citizenship was also reflected in Chapter 9, where we saw a number of services being allowed (or disallowed) based on various aspects of citizenship. Pervasive systems can affect us in many ways

throughout our daily lives. We therefore need a concept such as citizenship that can embrace the variety in our society, the range of rights and obligations, and the justification by which certain obstacles are removed or placed in our daily activities.

In Chapter 4 we discussed the second aspect of our framework: spheres. Using spheres (which are pools of information) we addressed issues of task and privacy. As we noted, an increasingly pervasive computing environment has the potential for fragmenting the information that is required to carry out a task. By reasoning in terms of information spheres, and designing accordingly, we can assist people in their tasks by grouping the information that is of relevance. At the same time, this grouping offers support for our approach to privacy (using public, social or private spheres). Our top-down approach to privacy addresses the privacy issues that are raised in a pervasive environment in a way that is related to the tasks being carried out as well as everyday life. Our reference to the public sphere also carries with it a number of social, political, and ethical considerations that could affect the use and acceptance of pervasive systems. We have embraced such non-technical issues because we wish to make pervasive systems relevant to everyday life.

The third aspect of our proposed framework was space, discussed in Chapter 5. Here, we argued that architecture is a pervasive system. As such, we can learn a lot from how architecture deals with issues such as designing parks for public access, the effects of their designs on people, and so on. Furthermore, we have argued that architecture should not be in conflict with pervasive computing systems: they both are designed to be part of everyday life, and should thus coexist in harmony. We claimed that this harmony can be achieved if our designs are based on the notion that architecture manipulates physical space whilst pervasive systems manipulate interaction spaces. Finally, we surveyed various approaches to design within archi-tecture, and provided a number of suggestions for the design of pervasive systems.

Having provided a detailed discussion of all three of our framework's elements, we proceeded in Chapter 6 to discuss how these three elements come together to form a design tool. The ideas of our framework play a central role when applying our design tool, and we operationalised this by developing a method for using our design tool. The strengths of our design tool lie in the fact that we now have the ability visually to represent, compare and evaluate our designs, as well as convey information in our framework in a way that is accessible to designers.

Having reached the stage where we had a framework and a design tool available to us, we demonstrated their usefulness on three different levels. First we showed how to inform design on an interface level by the development of an interaction technique (Chapter 7). We showed how interaction can be separated from the physical form of the system, thus minimising the constraints it sets. Furthermore, we showed how our discussion of privacy and interaction spaces can be reflected in the way we interact with the system. We focused on situations where the interaction itself creates inappropriate interaction spaces, and showed how an interaction technique can generate varying interaction spaces, thus giving the opportunity to users to select an appropriate interaction space. Finally, we discussed non-technical obstacles, such as economic issues impeding the use of a pervasive system, which may be taken into account when designing at the interface level.

We then moved our focus to design solutions. We showed how we take design recommendations and generate design solutions which we can explore and evaluate, as well as applying our design tool a post-hoc to an existing setup. Specifically, in the context of a hospital A&E department (Chapter 8) we showed how we can turn design suggestions into design solutions. Furthermore, we showed how to explore design alternatives, evaluate them for potential problems, as well as reason about existing and potential technologies. We discussed the issue of providing access to

information for free (at the point of consumption), which is a trend that characterises services of the public sector. From our analysis we also identified ways of representing and reasoning about conventional technology (such as posters and leaflets) in conjunction with digital technology. We also discussed issues of information, and how we can separate the information from the medium used to deliver it as a way of thinking about design alternatives. In terms of spaces, we explained how making changes in the physical environment can work to our advantage, and should therefore also be considered as design alternatives or solutions.

At the highest design level, we showed how we can generate overall design recommendations by making use of our framework. Specifically we looked at the city of Bath (Chapter 9), and applied our framework a priori to the design of a hypothetical pervasive system. We also analysed issues in relation to civic and navigational information, commercial interests, and the mutual effects between the city and our system. We applied our concepts and analysed how the various locations we studied can help us in understanding the needs and requirements for a public pervasive system in the city of Bath. These recommendations can be translated into lower-level design solutions and alternatives as explained in Chapter 8, which in turn can make use of our interface level solutions (Chapter 7).

10.3   Future work

Many problems still confront the development of theory within pervasive computing. For instance, we have highlighted the need to identify requirements and research areas in an appropriate way, not based on technological issues. The three dimensions we have proposed as part of our framework can be a starting point, and can be used as three general directions of research that are relevant to pervasive computing. Of course, there is room for technological issues within all three of the dimensions we have proposed.

We also identify the lack of a common language or a common frame of reference as a problem for the field of pervasive computing. We believe that it would be possible to define different categories of pervasive systems based on their diagram instances. For instance, we could speak of a "Class 2" system, as a way to refer to a pervasive systems in public spaces using private and social interaction spaces to deliver public information to many citizens.

In relation to our design tool diagram, we should note that based on a rough calculation there are approximately 240,000 possible diagram instances. It would be interesting to investigate how many instances are problematic, which instances tend to be the most popular ones, which instances are popular for specific problems or situations, and finally investigate ways of deriving the optimum path which would lead us from one (existing) instance to another (preferable) instance.

Also interesting would be the development of an interactive system which algorithmically uses our method to explore design instances. Perhaps visual and multimedia annotations could be useful to convey more information about the specific diagram instance. Although the design process is largely creative and innovative, there is still room for an "intelligent system" approach which would work out design alternatives and propose optimum solutions based on our framework here as well as on the specific requirements and needs of the designer.

Implementation   Many existing software architectures are related to the ideas we have developed in this thesis. Such implementations can be related back to the design process by means of our framework. To a certain extent our notions of information spheres are similar to the notions of web presence (see page 27) and InfoSpaces [Hong & Landay, 2004]. Further implementations of interest are described in [Fitzpatrick et al., 2002; Greenberg & Roseman, 1999; Patterson et al., 1996], while access mechanisms to information have been developed by Langheinrich [2002]. In terms of

our framework, such work addresses issues of implementing mechanisms for accessing the public, private and social spheres, transferring information between those spheres and enabling interaction with other participants of a social or public sphere. In trying to implement the ideas we have presented, the above work offers appropriate starting points, although extensions would need to be made so that our concepts and the relations we have described between them are reflected in the software architecture.

Awareness

Our work on interaction spaces has not directly touched on awareness issues. We have presented work in how to design appropriate interaction spaces in respect of space and spheres, and also a way to interact with interaction spaces. Having done so, the next step in the development process is to instantiate those interaction spaces using appropriate artefacts and interfaces. Our work on gestural interaction forms the basis for addressing interaction with the interaction spaces themselves. Yet, a number of issues relating to users' awareness of those interaction spaces needs to be addressed in order to instantiate and implement those interaction spaces. There is a body of research and literature on understanding the critical role of awareness in information interaction [Kraut et al., 1988] and how people track awareness information within their physical environment [Gutwin & Greenberg, 2002]. Furthermore, social issues such as concerns about distraction [Bellotti, 1996; Boyle et al., 2000] and how it is used within virtual communities [Turkle, 1995] have been researched. Awareness of information in the environment [Fitzpatrick et al., 2002; Grudin, 2001b; Gutwin & Greenberg, 2002] as well as awareness of multimedia interaction spaces [Finn et al., 1997] play an important role in understanding which interaction space is appropriate for a particular situation. Research in this area may be drawn upon to help the designer understand which interaction spaces should be incorporated in a pervasive system (i.e. while designing a system), but also to allow for the run-time manipulation of the system's interac-

tion spaces (i.e. once the system is being used). Once an interaction space has been chosen (at design time) or instantiated (at run time) our method could be used to indicate potential issues. Our work complements work on awareness in addressing the issue of designing appropriate interaction spaces, and offers the appropriate concepts of spaces and information spheres to bridge work on awareness with pervasive information access.

Privacy

On page 81 we explained that our work on information spheres is not an attempt at providing a theory of privacy. Rather, our notions of spheres, spaces and interaction spaces provide an important link from existing work in privacy to the design process for pervasive information access. Previous research has shown that privacy is affected by how the information is received in addition to the characteristics of the information itself [Davies, 1997]. A major part in the privacy equation is the receiver of the information [Adams & Sasse, 1999a; Adams, 1999; Bellotti & Sellen, 1993; Bellotti, 1997] as well as how the received information is used [Adams & Sasse, 1999b]. Our concepts of interaction spaces can greatly improve the quality of feedback that the sender of the information is getting, something which has been shown to be of importance [Adams & Sasse, 1999b]. The system could inform the user that, for example, "a public interaction space is being used at the other end of the line". Currently, however, users have to make a trade-off between how sensitive the information is, who will be receiving it, and what it will be used for [Adams, 2003]. Based on these assumptions, there exist models of user's perceptions of privacy, which we can be used to help a system classify information in the appropriate sphere. Once this happens, the appropriate interaction spaces can by dynamically generated depending on the available spaces and vice versa. Finally, it has been argued [Bellotti & Sellen, 1993] that relying merely on social controls for safeguarding privacy is dangerous, and that unobtrusive technology increases the risk of privacy invasion. Ultimately, privacy will have to be man-

aged through a combination of technology, legislation, corporate policy and social norms [Lessig, 1998]. Our work brings under the umbrella of pervasive systems issues from technology, legislation and social norms.

10.4 Conclusion

We are now in a position to assess how well we have addressed our original research question on page 12: "*How can we design for pervasive access to information?*". We have shown how we can represent pervasive systems, and thus relate them to each other. By doing this, we have also shown how we can compare systems and evaluate them against specific requirements. Furthermore, we have provided grounds on which design decisions can be based. Also, by observing certain repeated patterns, either in systems themselves or in their diagrammatic representations, we can learn to avoid problems that have been addressed before and utilise existing work. Finally, as far as prescribing technology goes, we have provided the language (namely interaction spaces) that can be used to describe characteristics of an arbitrary new technology, without this technology existing. Overall, therefore, we have addressed all of the issues that we set ourselves as benchmarks for enabling us to *design for pervasive access to information.*

We addressed the design of pervasive systems on three discrete levels, each of which has different requirements. We did this by applying our design ideas to interface level design (with the design of an interaction technique), design exploration and evaluation (in the A&E case study), and finally overall design recommendations and guidance (in the city of Bath case study). The fact that our ideas were applicable to varying degrees of analysis is an encouraging indicator for us.

Finally, it is worth noting that the ideas presented in this work relate primarily to the built environment and to industrialised societies with a high penetration of computing resources and wealth. To design truly pervasive systems in the purest

global sense will require major changes in political, economic and technological development. A framework and design tools for such a global system remains a challenge.

The challenges facing pervasive computing are great, but its goal is even greater. We strongly believe that society can benefit in many different ways from this kind of technology, and we also believe that progress within this area should be based on strong theoretical grounds.

APPENDIX

A worked example - the case of the phone lists (see Section 8.2).

(Note that here we do not enter into the discussion of evaluating the possible alternatives. This discussion is presented in Section 8.2).

The artefacts of interest are (❶):

Artefact
Phone directory
A4 Telephone lists
Waiting room
Reception

We now classify each of the artefacts (❷):

Artefact	Space			Interact. Space			Sphere			Citizen	
	P	S	Pr	P	S	Pr	P	S	Pr	M	1
Phone directory								✓			
A4 Telephone lists				✓						✓	
Waiting room	✓									✓	
Reception		✓								✓	

The A4 sheets are placed in the waiting room. The artefacts that form a self-bound group are (❸):

Artefact	Space			Interact. Space			Sphere			Citizen	
	P	S	Pr	P	S	Pr	P	S	Pr	M	1
Phone directory								✓			
A4 Telephone lists				✓						✓	
Waiting room	✓									✓	

The setup created by our artefacts looks like this (❸):

In the above diagram we identify issue {I9}(❹). This is a potential privacy breach. The solution we choose is to relocate the a A4 lists, and place them in the reception area(❹). Now the A4 lists create a social interaction space {I1}. Our complete list of artefacts now is (❷):

Artefact	Space			Interact. Space			Sphere			Citizen	
	P	S	Pr	P	S	Pr	P	S	Pr	M	1
Phone directory								✓			
A4 Telephone lists					✓					✓	
Waiting room	✓									✓	
Reception		✓								✓	

Since the A4 sheets are in the reception, the artefacts that form a self-bound group are (❸):

Artefact	Space			Interact. Space			Sphere			Citizen	
	P	S	Pr	P	S	Pr	P	S	Pr	M	1
Phone directory								✓			
A4 Telephone lists					✓					✓	
Reception		✓								✓	

The diagrammatic representation of the artefacts now becomes (❸):
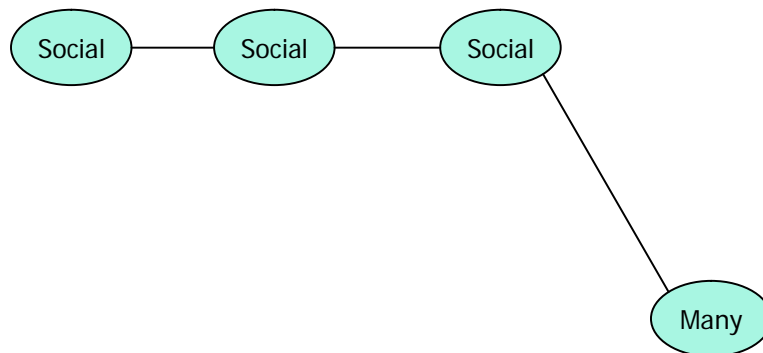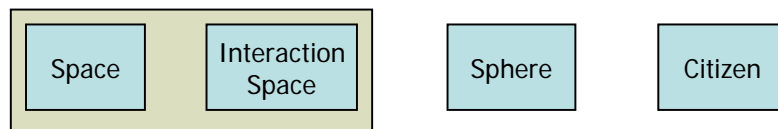
We now backtrack. Instead of relocating the A4 sheets to the reception area, we choose another solution.

We decide to use new technology - personal devices (❹). The artefacts of interest now are (❶):

Artefact
Phone directory
Personal device
Waiting room
Reception

We now classify each of the artefacts (❷)

Artefact	Space			Interact. Space			Sphere			Citizen	
	P	S	Pr	P	S	Pr	P	S	Pr	M	1
Phone directory								✓			
Personal device						✓					✓
Waiting room	✓									✓	
Reception		✓								✓	

All the artefacts form one group (❸). The diagram representing the setup now is (❸):
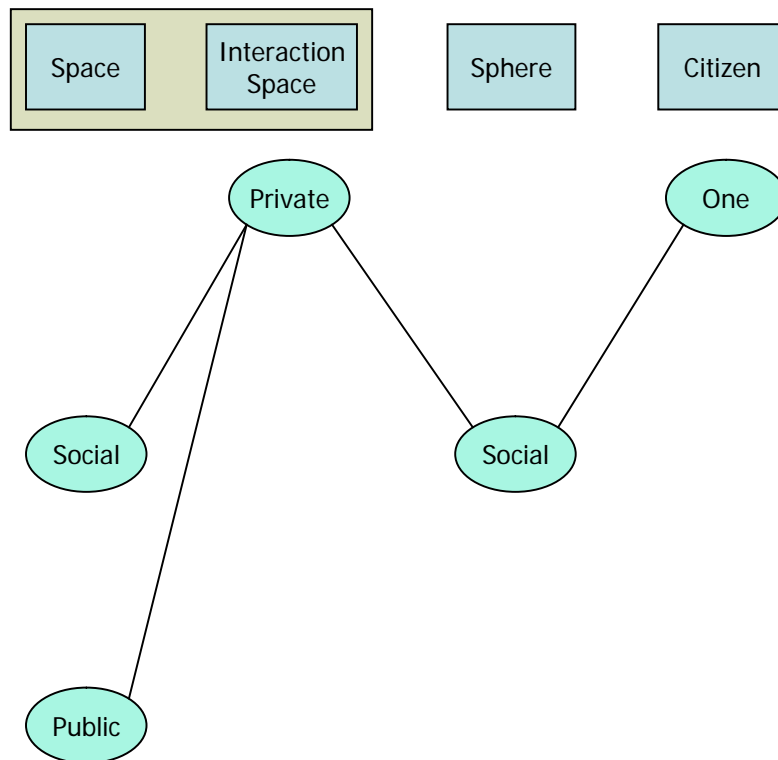
In the above diagram we note that the technology needs to be insulating {I4}.

The process we have just been through is summarised in the following diagram:

Initial setup

Relocate the A4 lists

Change the technology
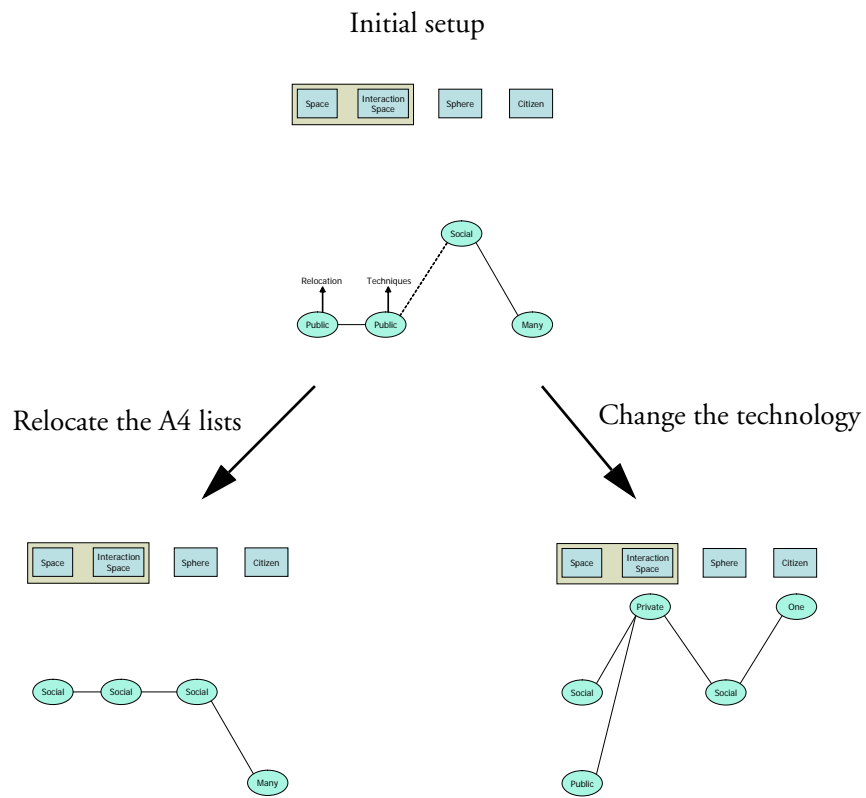
The end.

REFERENCES

Ackerman, M. S., Cranor, L. F. and Reagle, J. (1999). Privacy in e-commerce: examining user scenarios and privacy preferences. In proceedings of the *1st ACM conference on Electronic commerce*, ACM Press, New York, NY. p. 1-8.

Ackerman, M.S. (2000). The intellectual challenge of CSCW: The gap between social requirements and technical feasibility. *Human-Computer Interaction*, 15, 179-203.

Adams, A. (1999). Users' perception of privacy in multimedia communication. In proceedings Vol 2 of *the ACM Conference on Human factors in computing systems*, ACM Press. p. 53-54.

Adams, A. (2003). A whole picture is worth a thousand words. SIGCHI Bulleting (Supplement to Interactions). 2003: 12-12.

Adams, A. and Sasse, M. A. (1999a). Taming the wolf in sheep's clothing: privacy in multimedia communications. In proceedings of *the ACM international conference on Multimedia*, Orlando, Florida, ACM Press. p. 101-107.

Adams, A. and Sasse, M. A. (1999b). Privacy Issues in Ubiquitous Multimedia Environments: Wake Sleeping Dogs, or Let Them Lie? In proceedings of *INTERACT 1999.* p. 214-221.

Agre, P. (1997). Computing as a social practice. In Reinventing Technology, Rediscovering Community, P. Agre and D. Schuler, (eds). 1997, Ablex Publishing Corporation: London.

Alexander, C. (1975). The Oregon experiment. New York, Oxford University Press.

Alexander, C., Ishikawa, S. and Silverstein, M. (1977). A pattern language: Towns, buildings, construction. New York, Oxford University Press.

Annett, J. and Duncan, K. D. (1967). "Task analysis and training design." *Occupational Psychology* 41: 211-222.

Aoki, P. M., Romaine, M., Szymanski, M. H., Thornton, J. D., Wilson, D. and Woodruff, A. (2003). The Mad Hatter's Cocktail Party: A Social Mobile Audio Space Supporting Multiple Simultaneous Conversations. In proceedings of the conference on *Human factors in computing systems*, ACM Press, New York, NY. p. 425-432.

Baillie, L. (2002). The home workshop: a method for investigating the home. Phd Thesis. School of Computing. Edinburgh, Napier University.

Banavar, G., Beck, J., Gluzberg, E., Munson, J., Sussman, J., and Zukowski, D. (2000). Challenges: An Application Model for Pervasive Computing. In proceedings of the 6th Annual ACM/IEEE International Conference on Mobile Computing and Networking.

Barker, R. (1968). Ecological psychology: concepts and methods for studying the. Stanford, California, Stanford University Press.

Barnard, P. (1985). Interaction cognitive subsystems: A psycholinguistic approach to short-term memory. In Progress in the Psychology of Language, Vol 2(6), A. Ellis, (ed). 1985, Lawrence Erlbaum Associates: Hove.

Barnard, P. (1991). Bridging between basic theories and the artifacts of Human-Computer Interaction. In Designing interaction: psychology at the human computer interface, J. M. Carroll, (ed). 1991, Cambridge University Press: Cambridge. p. 103-127.

Bederson, B. B., Lee, B., Sherman, R. M., Herrnson, P. S. and Niemi, R. G. (2003). Electronic Voting System Usability Issues. In proceedings of the conference on *Human factors in computing systems*, Fort Lauderdale, FL, ACM Press, New York, NY. p. 145-152.

Bellotti, V. (1996). What you don't know can hurt you: Privacy in collaborative computing. In proceedings of *People and Computers XI, HCI 1996*, Springer. p. 241-261.

Bellotti, V. (1997). Design for privacy in multimedia computing and communication environments. In Technology and privacy in the new landscape, P. Agre and M. Rotenberg, (eds). 1997, MIT Press: Massachusetts.

Bellotti, V. and Sellen, A. (1993). Design for Privacy in Ubiquitous Computing Environments. In proceedings of *the European Conference on computer supported cooperative work*. p. 77-92.

Bellotti, V. and Edwards, K. (2001). "Intelligibility and Accountability: Human Considerations in Context-Aware Systems." *Human Computer Interaction* 16: 193-212.

Benford, S., Anastasi, R., Flintham, M., Drozd, A., Crabtree, A., Greenhalgh, C., Tandavanitj, N., Adams, M. and Row-Farr, J. (2003). Coping with uncertainty in a location-based game. IEEE Pervasive Computing, July-September. 2: 34-41.

Bentley, I. (1985). Responsive environments: A manual for designers. London, Architectural Press.

Biddulph, M. (1993). Consuming the sign value of urban form. In Making better places: urban design now, R. Hayward and S. McGlynn, (eds). 1993, Butterworth Architecture: Oxford.

Borning, A. (1981). "The programming language aspects of ThingLab, a constraint-oriented simulation laboratory." *ACM Transactions on Programming Languages and Systems*, 3(4): 353-387.

Borovoy, R., Martin, F., Resnick, M. and Silverman, B. (1998). GroupWear: Nametags that Tell about Relationships. In proceedings of the conference on *Human factors in computing systems*, Los Angeles, CA, Addison-Wesley. p. 329-330.

Boyle, M., Edwards, C. and Greenberg, S. (2000). The effects of filtered video on awareness and privacy. In proceedings of *the ACM Conference on Computer supported cooperative work*, ACM Press. p. 1-10.

Brebner, J. M. T. (1982). Environmental psychology in building design. London, Applied Science.

Brown, B. and Perry, M. (2002). Of maps and guidebooks: designing geographical technologies. In proceedings of the conference on *Designing Interactive Systems*, ACM Press, New York, NY. p. 246-254.

Buchecker, M. (2003). Public place as a resource of social interaction. In proceedings of the workshop on *Space, Spatiality and Technology*, Napier University, Edinburgh. p. 57-61.

Buyya, R., Abramson, D., Giddy, J. and Stockinger, H. (2002). "Economic models for resource management and scheduling in Grid computing." *Concurrency and Computation* 14(13-15): 1507-1542.

Callahan, J., Hopkins, D., Weiser, M. and Shneiderman, B. (1998), An empirical comparison of pie vs. linear menus. In proceedings of the conference on *Human Factors in Computing Systems*, ACM Press, New York, NY. p. 95-100.

Card, S. K., Moran, T. P. and Newell, A. (1983). The Psychology of Human Computer Interaction. Hillsdale, New Jersey, Lawrence Erlbaum Associates.

Carroll, J. M. and Rosson, M. B., (2001). Better home shopping or new democracy?: evaluating community network outcomes. In proceedings of the conference on *Human factors in computing systems*, ACM Press, New York, NY. p. 372-379.

Castells, M. (1995). The Information City. Oxford, Blackwell.

Catarci, T., Matarazzo, G. and Raiss, G. (2000). Usability and Public Administration: Experiences of a Difficult Marriage. In proceedings of the conference on *Universal usability*, Washington, DC, New York. p. 24-31.

Chapman, D. (1996). Creating neighbourhoods and places in the built environment. London, E & FN Spon.

Checkland, P. B. (1981). Systems Thinking, Systems Practice. Chichester, John Wiley.

Collins, C. C., Collins, G. R. and Sitte, C. (1965). City planning according to artistic principles. London, Phaidon Press.

Crampton, S. G. (1995). The Hand that Rocks the Cradle. I.D. May/June 1994. p. 60-65.

Dansky, K. H. and Miles, J. A. (1997). "Patient satisfaction with ambulatory healthcare services: waiting time and filling time." *Hospital and Health Services Administration* 42: 165-177.

Davies, N., Friday, A., Blair, G. S., and Cheverst, K. (1996). "Distributed Systems Support for Adaptive Mobile Applications." *Mobile Networks and Applications* 1(4): 399-408.

Davies, S. (1997). Re-engineering the right to privacy. In Technology and privacy in the new landscape, P. Agre and M. Rotenberg, (eds). 1997, MIT Press: Massachusetts.

Davis, F. D. (1993). "User acceptance of information technology: system characteristics, user perceptions and behavioral impacts." *International Journal of Man Machine Studies* 38(3): 475.

Dearden, A. and Walker, S. (2003). Designing for Civil Society. In Proceedings Vol 2 of *People and Computers XVII, HCI 2003: Designing for Society*, Bath, UK, Springer.

Dey, A. K., Abowd, G. D. and Salber, D. (2001). "A Conceptual Framework and a Toolkit for Supporting the Rapid Prototyping of Context-Aware Applications." *Human Computer Interaction* 16(2-4): 97-166.

Dix, A., (2000). Welsh mathematician walks in cyberspace. In proceedings of the *3rd international conference on collaborative virtual environments*, ACM Press, New York, NY. p. 3-7.

Dix, A. (2003). Managing Multiple Spaces. In proceedings of the workshop on *Space, Spatiality and Technology*, Napier University, Edinburgh.

Dix, A., Finlay, J., Abowd, G. D. and Beale, R. (1998). Human-computer interaction. 2nd Edition, London, Prentice Hall Europe.

Dix, A., Rodden, T., Davies, N., Trevor, J., Friday, A. and Palfreyman, K. (2000). "Exploiting Space and Location as a Design Framework for Interactive Mobile Systems." *ACM Transactions on Computer Human Interaction* 7(3): 285-321.

Donnelly, V. and Merrick, R. (2003). Community portals through communitization. In proceedings of the conference on *Universal usability*, ACM Press, New York, NY. p. 9-14.

Edwards, W. K., Mynatt, E. and Stockton, K. (1995). "Access to Graphical Interfaces for Blind Users." *Interactions* 2(1): 54-67.

Esler, M., Hightower, J., Anderson, T. and Borriello, G. (1999). Next century challenges: data-centric networking for invisible computing: the Portolano project at the University of Washington. In proceedings of the conference on *Mobile computing and networking*, ACM Press, New York, NY. p. 256-262.

Eysenck, M. W. and Keane, M. T. (2000). Cognitive psychology : a student's handbook. Hove, Psychology Press.

Finn, K. E., Sellen, A. J. and Wilbur, S., Eds. (1997). Video-mediated communication. Manwah, NJ, Lawrence Erlbaum.

Fitzpatrick, G., Kaplan, S. and Mansfield, T. (1998). Applying the Locales Framework to Understanding and Designing. In proceedings of the *Australasian Computer-human interaction conference*, Adelaide; Australia, Los Alamitos CA. p. 122-129.

Fitzpatrick, G., Kaplan, S., Mansfield, T., David, A. and Segall, B. (2002). "Supporting Public Availability and Accessibility with Elvin: Experiences and Reflections." *Computer Supported Cooperative Work* 11(3): 447-474.

Fitzpatrick, G., Mansfield, T. and Kaplan, S. (1996). Locales Framework: Exploring Foundations for Collaboration Support. In proceedings of the *Australasian Computer-human interaction conference*, Hamilton; New Zealand, IEEE Computer Society Press. p. 34-41.

Foster, I. and Kesselman, C. (1999). The grid: blueprint for a new computing infrastructure. San Francisco, Morgan Kaufmann Publishers.

Foster, I., Kesselman, C. and Tuecke, S. (2001). "The Anatomy of the Grid: Enabling Scalable Virtual Organizations." *International Journal of High Performance Computing Applications* 15: 200-222.

Fraser, N. (1995). Politics culture and the public sphere: toward a postmodern conception. In Social postmodernism: beyond identity politics, L. J. Nicholson and S. Seidman, (eds). 1995, Cambridge University Press: Cambridge, New York.

Garzonis, S., O'Neill, E., Kostakos, V., Kaenampornpan, M. and Warr, A. (2004). A Novel Approach for Identification and Authentication of Users in a Pervasive Environment. In proceedings of the *2nd UK-UbiNet Workshop*, University of Cambridge, UK.

Gaver, W. W., Beaver, J. and Benford, S. (2003). Ambiguity as a Resource for Design. In proceedings of the conference on *Human factors in computing systems*, Fort Lauderdale, FL, ACM Press, New York, NY. p. 233-240.

Giddens, A. (1982). Profiles and critiques in social theory. London, Macmillan.

Gifford, R. (1987). Environmental psychology: principles and practice. Boston, Allyn and Bacon.

Green, L. (2002). Communication, technology and society. London, SAGE.

Greenberg, S. and Roseman, M. (1999). Groupware Toolkits for Synchronous Work. In Computer-Supported Cooperative Work (Trends in Software 7), M. Beaudouin-Lafon, (ed). 1999, John Wiley & Sons Ltd. p. 135-168.

Grigni, M., Papadias, D. and Papadimitriou, C. (1995). "Topological Inference." *International Joint Conference on Artificial Intelligence* 14(1): 901-907.

Grosz, B. J. and Kraus, S. (1996). "Collaborative plans for complex group action." *Artificial Intelligence* 86(2): 269-358.

Grudin, J. (2001a). "Desituating Action: Digital represenation of Context." *Human Computer Interaction* 16(2-4): 269-286.

Grudin, J. (2001b). Partitioning digital worlds: focal and peripheral awareness in multiple monitor use. In *the ACM conference on Human factors in computing systems*, ACM Press. p. 458-465.

Guimbretiere, F., Stone, M. and Winograd, T. (2001). Fluid interaction with high-resolution wall-size displays. In proceedings of the symposium on *User Interface Software and Technology*. p. 21-30.

Guimbretiere, F. and Winograd, T. (2000). FlowMenu: combining command, text and data entry. In proceedings of the symposium on *User Interface Software and Technology*. p. 213-217.

Guttman, E., Perkins, C., Veizades, J. and Day, M., Eds. (1999). Service Location Protocol, Version 2, RFC Editor.

Gutwin, C. and Greenberg, S. (2002). "A Descriptive Framework of Workspace Awareness for Real-Time Groupware." *Computer Supported Cooperative Work* 11(3): 411-446.

Habermas, J. (1962). Strukturwandel der èOffentlichkeit; Untersuchungen zu einer Kategorie der bèurgerlichen Gesellscahft. Neuwied,, H. Luchterhand. Translated as The structural transformation of the public sphere, Cambridge, Polity, 1989.

Hall, E. T. (1969). The hidden dimension. Man's use of space in public and private. Bodley Head, London.

Hanson, V. L. (2001). Web access for elderly citizens. In proceedings of the *2001 EC/NSF Workshop on Universal Accessibility of Ubiquitous Computing: Providing for the elderly.* p. 14-18.

Harrison, J. and Woods, L. M. (2001). "Defining European Public Service Broadcasting." *European Journal of Communication* 16(4): 477-504.

Harrison, S. and Dourish, P. (1996). Re-place-ing space: the roles of place and space in collaborative systems. In proceedings of the conference on *Computer supported cooperative work*, ACM Press, New York, NY. p. 67-76.

Hayes, P. J., Szekely, P. A., and Lerner, R. A. 1985. Design alternatives for user interface management systems based on experience with COUSIN. In proceedings of the conference on *Human Factors in Computing Systems*, ACM Press, New York, NY. p. 169-175.

Held, D. (1989). Political theory and the modern state: essays on state and power. Cambridge, Polity Press.

Hillier, B. and Hanson, J. (1984). The social logic of space. Cambridge, Cambridge University Press.

Hodes, T. D., Czerwinski, E., Zhao, Y., Joseph, D. and Katz, H. (2002). "An Architecture for Secure Wide-Area Service Discovery." *Wireless Networks* 8(2-3): 213-230.

Hodes, T. D., and Katz, R. H. (1998). Enabling "Smart Spaces": Entity Description and User Interface Generation for a Heterogeneous Component-Based Distributed System. In proceedings of the *DARPA/NIST Smart Spaces Workshop*, Gaithersburg, Maryland, 1998.

Hong, J. I. and Landay, J. A. (2004). An architecture for privacy-sensitive ubiquitous computing. In proceedings of *the 2nd international conference on Mobile systems, applications, and services*, ACM Press. p. 177-189.

Ishii, H. and Ullmer, B. (1997). Tangible Bits: Towards Seamless Interfaces between People, Bits and Atoms. In proceedings of the conference on *Human factors in computing systems*, Atlanta; GA, Addison-Wesley. p. 234-241.

Johnson, H. and Hyde, J. (2003). "Towards modeling individual and collaborative construction of jigsaws using task knowledge structures (TKS)." *Acm Transactions on Computer Human Interaction* 10(4): 339-387.

Johnson, P. and Johnson, H. (1991). "Task Knowledge Structures: Psychological basis and integration into system design." *Acta Psychologica* 78: 3-26.

Jolly, E. J. and Reardon, R. (1985). "Cognitive differentiations, automaticity, and interruptions of automatized behaviours." *Personality and Social Psychology Bulletin* 11(3): 301-314.

Kaenampornpan, M. and O'Neill, E. (2004). Modelling context: an Activity Theory approach. In proceedings of the *2nd European Symposium on Ambient Intelligence, EUSAI 2004*, Eindhoven, The Netherlands. Springer LNCS 3295, p.367-374

Kaenampornpan, M., O'Neill, E., Kostakos, V. and Warr, A. (2004). Classifying Context Classifications: an Activity Theory Perspective. In proceedings of the *2nd UK-UbiNet Workshop*, University of Cambridge, UK.

Kasik, D. J. (1982). A user interface management system. *Computer Graphics* 16(3): 99-106.

Kaya, N. and Erkip, F. (1999). "Invasion of Personal Space Under the Condition of Short-Term Crowding: An Automatic Teller Machine." *Journal of Environmental Psychology* 19(2): 183-190.

Kayden, J. S., New York Department of City Planning and Municipal Art Society of New, Y. (2000). Privately owned public space: the New York City experience. New York ; Chichester, John Wiley.

Kieras, D. E. and Polson, P. G. (1985). "An approach to the formal analysis of user complexity." *International Journal of Man Machine Studies* 22: 365-394.

Kindberg, T., Barton, J., Morgan, J., Becker, G., Caswell, D., Debaty, P., Gopal, G., Frid, M., Krishnan, V., Morris, H., Schettino, J., Serra, B. and Spasojevic, M. (2002). "People, places, things: web presence for the real world." *Mobile Networks and Applications* 7(5): 365-376.

Kindberg, T., Sellen, A. and Geelhoed, E. (2004). Security and Trust in Mobile Interactions: A Study of Users' Perceptions and Reasoning. In proceedings of the *2nd UK-UbiNet Workshop*, University of Cambridge, UK.

Koleva, B., Benford, S. and Greenhalgh, C. (1999). The properties of mixed reality boundaries. In proceedings of the *European Conference on Computer Supported Cooperative Work*, Copenhagen, Dordrecht. p. 119-138.

Kostakos, V. and O'Neill, E. (2003). A directional Stroke Recognition Technique for Mobile Interaction in a Pervasive Computing World. In proceedings of the conference on *People and Computers XVII, HCI 2003: Designing for Society*, Bath, UK, Springer-Verlag. p. 197-206.

Kostakos, V. and O'Neill, E. (2004a). Designing Pervasive Systems for Society. In proceedings of the *Second International Conference on Pervasive Computing, Volume 2, First International Workshop on Sustainable Pervasive Computing*, Vienna, Austria.

Kostakos, V. and O'Neill, E. (2004b). Extending Traditional Design approaches for Pervasive Computing. In proceedings of *PREP2004*, University of Hertfordshire, UK.

Kostakos, V. and O'Neill, E. (2004c). Pervasive computing in emergency situations. In proceedings of the *Thirty-Seventh Annual Hawaii International Conference on System Sciences*, Hawaii, Computer Society Press. p. 81-90.

Krajewski, M. (2001). Public Services and the Scope of the General Agreement on Trade in Services. Geneva, May 2001 Center for International Environmental Law.

Kraut, R., Egido, C. and Galegher, J. (1988). Patterns of contact and communication in scientific research collaboration. In proceedings of *the ACM conference on Computer supported cooperative work*, ACM Press. p. 1-12.

Krier, R., Czechowski, C. and Black, G. (1979). Urban space. London, Academy Editions.

Langheinrich, M. (2002). A privacy awareness system for ubiquitous computing environments. In proceedings of *Ubicomp*, Goteberg, Sweden. p. 237-245.

Laski, H. J. (1962). The Rise of European Liberalism. London, Allen and Unwin.

Lees, L. H. (1994). "Urban public space and imagined communities in the 1980s and 1990s." *Journal of Urban History* 20(4): 443-465.

Lessig, L. (1998). The architecture of privacy. In proceedings of *the Taiwan NET 98 Conference*, Taipei, Taiwan. See online version http://cyber.law.harvard.edu/works/lessig/architecture_priv.pdf.

Lister, R. and Child Poverty Action, G. (1990). The exclusive society: citizenship and the poor. London, Child Poverty Action Group.

Little, L. (2003). An activity theory approach to technology use in public areas: the case of the ATM. In Proceedings Vol 2 of *HCI 2003: Designing for Society*, University of Bath, UK.

Logie, G. C. (1954). The Urban Scene. London, Faber & Faber.

Long, A. C., Landay, J.A., Rowe, L.A. and Michiels, J. (2000). Visual Similarity of Pen Gestures. In proceedings of the conference on *Human Factors in Computing Systems*, ACM Press, New York, NY. p. 360-367.

MacCaulay, L., Fowler, C., Kirby, M. and Hutt, A. (1990). "A new approach to requirements specification." *Interacting With Computers* 2(1): 92-108.

Maister, D. H. (1984). The psychology of waiting lines. Boston, Mass., Harvard Business School.

Malina, A. (1999). Perspectives on citizen democratisation and alienation in the virtual public sphere. In Digital Democracy: Discourse and Decision Making in the Information Age, B. N. Hague and B. D. Loader, (eds). 1999, Routledge: London.

Maniatis, P., Roussopoulos, M., Swierk, E., Lai, K., Appenzeller, G., Zhao, X. and Baker, M. (1999). "The mobile people architecture." *Mobile Computing and Communications Review* 3(3): 36-42.

Mann, M. (1987). "Ruling class strategies and citizenship." *Sociology* 21: 339-354.

Marshall, T. H. (1950). Citizenship and social class, and other essays. Cambridge, University Press.

McCall, R. (2003). The Place of Space in User Interface Design. In proceedings of the workshop on *Space, Spatiality and Technology*, Napier University, Edinburgh.

McGinity, M. (2004). "Staying Connected RFID: Is This Game of Tag Fair Play?" *Communications of the ACM* 47(1): 15-18.

Miller, Franz (2001). "Wired and Smart: from the Fridge to the Bathtub." *ERCIM News* 47: 17-18.

Mumford, E. (1993). The participation of users in systems design: an account of the origin, evolution, and use of the ETHICS method. In Participatory design: principles and practices, D. Schuler and A. Namioka, (eds). 1993, Lawrence Erlbaum Associates: Hillsdale, N.J.

Myers, B. A., Giuse, D. A., Dannenberg, R. B., Kosbie, D. S., Pervin, E., Mickish, A., Zanden, B. V., and Marchal, P. (1990). "Garnet: Comprehensive support for graphical, highly interactive user interfaces." *IEEE Computer* 23,(11): 71-85.

Naumann, S. and Miles, J. A. (2001). "Managing waiting patients' perceptions: the role of process control." *Journal of Management in Medicine* 15(4-5): 376-386.

Newman, W. M. (1968). A system for interactive graphical programming. In proceedings of the *Spring Joint Computer Conference*, AFIPS Press, Arlington, VA. p. 47-54.

Norberg-Schulz, C. (1971). Existence, space & architecture. London, Studio Vista.

O'Neill, E., Kaenampornpan, M., Kostakos, V., Warr, A. and Woodgate, D. "Can we do without GUIs? Gesture and speech interaction with a patient information system." *Personal and Ubiquitous Computing*, under Review.

O'Neill, E., Johnson, P. and Johnson, H. (1999). "Representations and User-Developer Interaction in Cooperative Analysis and Design." *Human Computer Interaction* 14(1-2): 43-92.

O'Neill, E. (2000). User-developer cooperation in software development: building common ground and usable systems. London, Springer Verlag.

O'Neill, E., Woodgate, D. and Kostakos, V. (2004). Easing the Wait in the Emergency Room: Building a Theory of Public Information Systems. In proceedings of the conference on *Designing Interactive Systems*, Cambridge, Massachusetts. p. 17-25.

Oldenburg, R. (1999). The great good place: cafés, coffee shops, bookstores, bars, hair. New York, Marlowe.

Olsen, D. R. (1986). "MIKE: The menu interaction kontrol environment." *ACM Transactions on Graphics* 5(4): 318-344.

Olsen, D. R., Jr., and Dempsey, E. P. (1983). Syngraph: A graphical user interface generator. In proceedings of the *ACM Conference on Computer Graphics*, ACM Press, New York, NY. p. 43-50.

Papadias, D. and Egenhofer, M. J. (1997). "Algorithms for Hierarchical Spatial Reasoning." *Geoinformatica* 1: 251-274.

Parsons, T. and Platt, G. (1973). The American University, Cambridge. Cambridge, MA, Harvard University Press.

Pascall, G. (1986). Social policy: a feminist analysis. London, Tavistock.

Patterson, J. F., Day, M. and Kucan, J. (1996). Notification servers for synchronous groupware. In proceedings of *the ACM conference on Computer supported cooperative work*, ACM Press. p. 122-129.

Patterson, W. (1999). "Transforming Electricity." *Journal of Energy Literature* 5: 89-91.

Pedersen, E. R., Sokoler, T. and Nelson, L., (2000). PaperButtons: expanding a tangible user interface. In proceedings of the conference on *Designing interactive systems*. 2000, ACM Press, New York, NY. p. 216-223.

Perkins, C. and Harjono, H. (1996). Resource discovery protocol for mobile computing. In proceedings of the conference on *Mobile communications*, Canberra, Chapman & Hall. p. 219-238.

Pier, K. and Landay, J. (1992). Issues for location-independent interfaces. Technical Report ISTL92-4. Xerox PARC, Palo Alto, CA. Available at http://www.cs.berkeley.edu/~landay/research/publications/LII.ps.

Pirhonen, A., Brewster, S. and Holguin, C. (2002). Gestural and Audio Metaphors as a Means of Control for Mobile Devices. In proceedings of the conference on *Human Factors in Computing Systems*, ACM Press, New York, NY. p.291-298.

Potel, M., and Cotter, S. (1995). Inside Taligent Technology. Addison - Wesley, 1995.

Project for Public Spaces (2000). How to turn a place around: A handbook for creating successful public spaces. New York, NY, Project for Public Spaces.

Preece, J., Rogers, Y. and Sharp, H. (2002). Interaction design: beyond human-computer interaction, John Wiley & Sons.

Putnam, R. D. (2000). Bowling alone: the collapse and revival of American community. New York ; London, Simon & Schuster.

Reid, J., Hull, R., Melamed, T. and Speakman, D. (2003). "Schminky: The design of a cafe based digital experience." *Personal and Ubiquitous Computing* 7(3-4): 197-202.

Rekimoto, J. (1997). Pick-and-drop: a direct manipulation technique for multiple computer environments. In proceedings of the *10th annual ACM symposium on User interface software and technology*, ACM Press, New York, NY. p. 31-39.

Rekimoto, J. (2001). GestureWrist and GesturePad: Unobtrusive Wearable Interaction Devices. In proceedings of the conference on *Wearable computers*, Zurich, Switzerland, IEEE. p. 21-30.

Rekimoto, J., Ullmer, B. and Oba, H. (2001). DataTiles: a modular platform for mixed physical and graphical interactions. In proceedings of the conference on *Human factors in computing systems*, ACM Press, New York, NY. p. 269-276.

Rezgui, A., Ouzzani, M., Bouguettaya, A. and Medjahed, B. (2002). Preserving privacy in web services. In Proceedings of the *4th international workshop on Web information and data management*, ACM Press, New York, NY. p. 56-62.

Rodden, T. and Benford, S. (2003). The evolution of buildings and implications for the design of ubiquitous domestic environments. In proceedings of the conference on *Human factors in computing systems*, Fort Lauderdale, FL, ACM Press, New York, NY. p. 9-16.

Roseman, M, and Greenberg, S. (1996). "Building real-time groupware with GroupKit, a groupware toolkit." *ACM Transactions on Computer Human Interction* 3(1): 66-106.

Ruback, R. B., Pape, K. D. and Doriot. p. (1989). "Waiting for a phone: intrusion on callers leads to territorial defense." *Social Psychology Quarterly* 53(3): 232-241.

Rubine, D. (1991). Specifying gestures by example.   In proceedings of the conference on *Computer Graphics and Interactive Techniques*. p. 329-337.

Rui, Y., He, L., Gupta, A. and Liu, Q., (2001). Building an intelligent camera management system. In proceedings of the *9th ACM international conference on Multimedia*, ACM Press, New York, NY. p. 2-11.

Schmidt, A., Beigl, M. and Gellersen, H. W. (1999). "There is more to context than location." *Computers and Graphics* 23(6): 893-901.

Schuler, D. (2001). "Cultivating society's civic intelligence: patterns for a new 'world brain'." *Information Communication and Society* 4: 157-181.

Schulert, A. J., Rogers, G. T., and Hamilton, J. A. (1985). ADM-A dialogue manager. In proceedings of the conference on *Human Factors in Computing Systems*, ACM Press, New York, NY. p. 177-183.

Shepherd, A. (1989). Analysis and training in information technology tasks. In Task analysis for Human-Computer interaction, D. Diaper, (ed). 1989, Ellis Horwood: Chichester. p. 15-55.

Shirky, C. (2003). A Group Is Its Own Worst Enemy: Social Structure in Social Software. Keynote in the *O'Reilly Emerging Technology conference*, Santa Clara.

Shneiderman, B. (1983). Direct manipulation: A step beyond programming languages. *IEEE Computer* 16(8): 57-69.

Shneiderman, B. (2002). Leonardo's laptop: human needs and the new computing technologies. Cambridge, Mass., MIT Press.

Silverstone, R. and Hirsch, E. (1992). Consuming technologies: media and information in domestic spaces. London, Routledge.

Stephanidis, C. (2001). "Adaptive Techniques for Universal Access." *User Modeling and User Adapted Interaction* 11(1-2): 159-179.

Stewart, D. (2002). "Taking the Stress from the Queue." *Logistics and Transport Focus* 4(10): 70-76.

Stoker, G. and Williams, J. (2001). Memorandum by Professor Gerry Stoker, Chair and John Williams, Executive Director, New Local Government Network (PSR 11). Select Commitee on Public Administration, 29 November 2001.

Streitz, N. A., Geißler, J., Holmer, T., Konomi, S. I., Müller-Tomfelde, C., Reischl, W., Rexroth, P., Seitz, P., and Steinmetz, R. (1999). i-Land: An interactive landscape for creativity and innovation. In proceedings of the conference on *Human Factors in Computing Systems*, ACM Press, New York, NY. p. 120-127.

Suchman, L. (1987). Plans and situated actions: The problem of human-machine communication. New York, Cambridge University Press.

Sukaviriya, P. N., Foley, J. D., and Griffith, T. (1993). A second generation user interface design environment: The model and the runtime architecture. In proceedings of the conference on *Human Factors in Computing*, ACM Press, New York, NY. p. 375-382.

Sutherland, I. E. (1963). Sketchpad: A man-machine graphical communication system. In proceedings of the *AFIPS Conference*. p. 329-346.

Svanaes, D. and Verplank, W., (2000). In search of metaphors for tangible user intefaces. In proceedings of the conference on *Designing augmented reality environments*, ACM Press, New York, NY. p. 121-129.

Szekely, P., Luo, P., and Neches, R. (1993). Beyond interface builders: Model-based interface tools. In proceedings of the conference on *Human Factors in Computing*, ACM Press, New York, NY. p. 383-390.

Tilley, S. (2003). Computer documentation for senior citizens. In Proceedings of the *21st annual international conference on Documentation*, ACM Press, New York, NY. p. 143-146.

Tuan, Y. F. (1971). "Geography, phenomenology and the study of human nature." *The Canadian Geographer* 15: 181-192.

Tunstall, R. (2001). "Devolution and User Participation in Public Services: How They Work and What They Do." *Urban Studies* 38(13): 2495-2514.

Turkle, S. (1995). Life on the screen: identity in the age of the Internet. New York ; London, Simon & Schuster.

Turner, B. (1990). "Outline of a theory of citizenship." *Sociology* 24(2): 189-217.

Turner, B. (1994). Contemporary Problems in the Theory of Citizenship. In Citizenship and Social Theory, B. Turner, (ed). 1994, Sage: London.

Turner, P. and Turner, S. (2003). Two Phenomenological Studies of Place. In proceedings of the conference on *Human computer interaction; People and computers XVII: designing for society,* Bath, UK, London. p. 21-36.

Vanhala, L. (2001). "A flood of Intelligence - the Living Room Project." *ERCIM News* 47: 14-15.

Vazhkudai, S. and Von Laszewski, G. (2001). "A Greedy Grid - The Grid Economic Engine Directive." *International Parallel and Distributed Processing Symposium*: 173.

Venkatesh, A. (1996). "Computers and Other Interactive Technologies for the Home." *Communications of the ACM* 39(12): 47-54.

Vicente, K. J. (2003). The human factor: revolutionizing the way people live with technology. Toronto, A.A. Knopf Canada.

Want, R., Hopper, A., Falc, V. and Gibbons, J. (1992). "The active badge location system." *ACM Transactions on Information Systems* 10(1): 91-102.

Wasserman, A. I., and Shewmake, D. T. (1982). Rapid prototyping of interactive information systems. *Software Engineering Notes* 7(5): 171-180.

Weiser, M. (1991). "The Computer for the Twenty-First Century." *Scientific American* 265(3): 94-104.

Westin, A. F. (1970). Privacy and Freedom. New York, NY, Atheneum.

Watson, R. T. and Mundy, B. (2001). "A strategic perspective of electronic democracy." *Communications of the ACM* 44(1): 27-30.

Wiecha, C., Bennett, W., Boies, S., Gould, J., and Greene, S. (1990). "ITS: A tool for rapidly developing interactive applications." *ACM Transactions on Information Systems* 8(3): 204-236.

Wyckoff, P., McLaughry, S. W., Lehman, T. J. and Ford, D. A. (1998). "T Spaces." *IBM Systems Journal* 37(3): 454-474.

Yamato, M., Inoue, K., Monden, A., Torii, K. and Matsumoto, K.-i., (2000). Button selection for general GUIs using eye and hand together. In proceedings of the working conference on *Advanced visual interfaces*. ACM Press, New York, NY. p. 270-273.

Zachary, W. W., Ryder, J. M. and Hicinbothom, J. H. (2000). Building cognitive task analysis and models of decision-making team in a complex real-time environment. In Cognitive Task Analysis, J. M. Schraagen, S. F. Chipman and V. L. Shalin, (eds). 2000, Lawrence Erlbaum and Associates: Mahwah, N.J. p. 365–383.

Zhang, L., Deering, L., Estring, D., Shenker, S., and Zappala, D. (1993). RSVP: A New Resource Reservation Protocol. *IEEE-Network*, 7(5): 8-18.

Zhao, R. (1993), Incremental recognition in gesture-based and syntax-directed diagram editors. In proceedings of the conference on *Human factors in computing systems*, ACM Press, New York, NY. p. 95-100.