

The Reality Gap: Pragmatic Boundaries of Context Awareness

Recent developments in wireless communication and mobile computing have prompted a new vision of the world in which we live. Augmented reality has been proposed as the integration of physical and digital space. This approach is built upon the idea that it is possible to blur the distinctions that currently exist between computers and other forms of artifacts, ranging from household appliances to pieces of furniture. The exchange of digital information both about and between these artifacts creates powerful opportunities for context-aware computing (Ishii et al., 1997). Inspired by this vision, we decided to engineer a practical and extensible system for augmented reality using ‘off the shelf’ wireless-enabled computing devices. The first demonstration of this technology was to be the development of electronic notes that could be attached to inanimate objects and read using a PDA. Wireless links were used to trigger the presentation of information when a user approached a ‘tagged’ object. They were also used to update information about the notes in real-time from any location. Rather than extending this metaphor of augmented reality or proposing further developments to the abstract ideas of context aware computing, this paper describes the practical challenges of engineering such an interface. The hope is that others might learn from our mistakes.

Keywords: wireless communication, mobile computing, augmented reality.

1 Introduction

Mobile phones and personal digital assistants (PDA) help people to traverse the divide between physical and digital space. These devices enable users to access information from remote servers about the physical location that they are currently in. Such developments form part of a more general trend that is pushing physical and digital environments closer and closer together (Casewell et al., 2000, Oppermann et al., 1999). For instance, chat rooms and web-based shopping malls duplicate the physical experience within a digital format. It is a matter of personal experience as to whether these digital developments enhance or parody their physical counterparts. Their success cannot be denied, at least in the short term. Many of these developments have, however, been ad hoc. In consequence, a number of HCI researchers have searched for a more systematic approach to the integration of physical and digital space (Ishii et al., 1997, Want et al., 1999). This work has developed sophisticated theories and a number of case studies that point

to the benefits of tangible bits and augmented reality. A smaller group of studies have pointed to the ethical problems and social exclusion that might result from such development.

Rather than focus on the mass of theoretical and methodological research, this paper examines the engineering challenges that HCI developers must address if they wish to achieve the benefits proposed for augmented reality. Our focus is on systems that will work in an indoor environment. This is appropriate because it is the context in which the greatest benefits have been claimed for augmented reality. It also poses the greatest engineering challenges. Physical barriers prevent GPS, the lingua franca of location detection, from supporting the location detection that is a basic building block for context awareness.

1.1 Previous Work

This section reviews a number of different engineering approaches that have supported the development of augmented reality systems. Brevity prevents a complete overview of the many techniques that have been proposed. We, therefore, focus on the pioneering work of the Active Badge developers. We then go on to consider the Cyberguide approach. It is important to introduce this application because it provided a blueprint for our initial prototypes. However, Cyberguide relies upon specialist hardware. A primary motivation behind this paper is to show how augmented reality systems can be constructed using widely available, off the shelf components. Finally, we describe Hewlett Packard's Cooltown project. This vision of context aware computing is important for the rest of our paper because it forms a contrast with the more prosaic problems that must be addressed before HCI developers can fully achieve the proposed benefits.

The Active Badge system is arguably the first context aware system (Harter et al., 1994). In particular, it provided early demonstrations of 'follow me' and security control applications in the office domain. Using this system, users are required to wear badges. Each badge is able to transmit a periodic infrared signal, which contains a unique code as its identifier. These periodic infrared signals are detected and picked up by networked badge sensors. Badge sensors are placed at fixed locations around the environment within the host building and are wired into a local area network. The badges' locations can be obtained once the sensors receive the infrared signal. A 'sighting' event is then sent to a server, which hosts the location-dependent information. The server can then provide relevant information to the user based on their current location. Although the Active Badge system can tailor information based on the users' current location, it did suffer from a number of problems. Significant installation costs were associated with the sensor network this inevitably involved the use of custom-made hardware. There was also a myriad of social and cultural issues associated with the involuntary disclosure of location information. Users responded by adopting a range of informal protocols to determine when and where a badge would actually be worn.

The Active Badge system exploited sensors that actively detected signals from the users' badges. In contrast, passive location detection relies upon the users' equipment detecting of signals that are generate from beacons in their environment. This is an important distinction because these beacons can be low cost signal generators, such as remote control devices, rather than the active sensors of the pioneering systems. For instance, the Indoor Cyberguide system utilized passive infrared sensing technology to demonstrate a context awarene tour guide application (Abowd et al., 1997). It was one of the first attempts to use PDA technology to bridge the divide between physical and digital environments. Unlike the Active Badge system, Cyberguide used infrared transmitters to communicate a unique identifier for each location. It did not use a networked system of base sensors. Each user carries a PDA, which is equipped with an additional specialist infrared sensor using a Motorola 68332 microcontroller. Once the user enters the signal of the infrared transmitter, otherwise known as its cell, the PDA can detect and decode the transmission into a unique identifier. In the original implementation, each PDA cached all of the location dependent information. The infrared location transmission was used to index into this stored data. However, this approach again relied upon a custom -de infrared unit. This is a significant issue for developers who lack access to the specialist hardware or who lack the engineering skills that are necessary to incorporate the signals from these units into existing PDAs. Cyberguide also relied upon cached information. There was no means of dynamically updating information as the user moved around their environment and so the context was essentially frozen once the PDA was disconnected from its network infrastructure.

Hewlett Packard's Cooltown project is based on the idea that every user, every object and every location in physical space will have a web-based representation (Caswell et al., 2000, Hohl et al., 1999, Oppermann et al., 1999). For instance, the visitors to Cooltown Museum will be able to access information about the exhibits from their PDA. These devices will automatically download this information from the web using a URL that is associated with an infrared beacon placed close to each exhibit. Visitors implicitly select the web reference by moving within the transmission range of an infrared beacon. This project, therefore, exploits the same passive infrared sensing technology as the Cyberguide system. However, it is augmented with web-enabled devices, such as printers and projectors. Users can exploit their PDA's to print documents by sending their URLs to an appropriate printer. Similarly, documents might be sent to web-enabled projectors during meetings. This illustrates an important difference between the Cooltown project and the other approaches in this section. It is a pioneering vision rather than an engineering infrastructure. It looks beyond what is currently available to look at what might be feasible in the short to medium term. There are, however, a number of limitations. For instance, it does not consider the engineering of an appropriate network topology to support the more visionary applications. For instance, what happens when users are out of transmission range? Similarly, although some of the papers assume an infra-red architecture, the implementation details are not sufficient for HCI developers to implement a functioning system without recourse to the specialist hardware of previous approaches.

Some of the objections that we have raised will be addressed by technological developments. In particular, Bluetooth devices will enable the transmission and reception of short-range location information without additional hardware. These short-range signals enable users to identify objects that are in their current vicinity, for instance between 2-5 meters away. In contrast, existing commercial radio LANs provide signals that extend from 100-200 meters away. This resolution is not fine enough for most location detection systems. Alternatively, Microsoft are working on triangulation techniques that exploit differing signal strengths from multiple LAN radio beacons to perform fine-grained location detection. Unfortunately, both of these approaches face a number of technological challenges. Differential signal detection is subject to distortions from the movement of objects in the environment. The Bluetooth consortium has been working for more than three years to agree the details of their technical specification and commercial devices are only just beginning to appear on the market. In the meantime, there is a requirement for HCI researchers and developers to have some means of validating the claims for context awareness. Ideally, such a system should be modular, easily extensible and should exploit off the shelf technology.

1.2 The Glasgow Context Server (GCS)

A number of issues exacerbate the development of context aware applications. These can be summarized as follows:

- (1) They have not been integrated with existing mobile computing technology and have relied upon additional hardware;
- (2) The supporting infrastructure has been too costly to deploy and maintain;
- (3) There is no efficient network topology to support disconnected users;
- (4) They have failed to address the social and personnel concerns of their users.

The Glasgow Context Server (GCS) system has been developed to address these concerns (Johnson and Cheng, 2000). GCS uses the existing infrared port on most PDAs. Software has been developed so that the existing PDA hardware can detect signals from most commercial infrared remote controllers. As a result, it is possible to add a new beacon simply by buying a new remote controller from any one of a number of commercial suppliers at minimal cost. It is possible to cache information on the PDAs, in the manner described for the Cyberguide system. However, the full GCS system exploits commercial wireless local area networks. Current implementations use off the shelf systems produced by Apple and Lucent. As a result, PDAs detect their position and the presence of other objects using infrared signals from commercial remote controllers. They can then use the radio network to either communicate their location to a central server or to request information about the objects in their environment. Rather than adopt the active sensing techniques adopted by the Active Badge projects, GCS adopts a passive approach. In other words, GCS will never disclose the users' position unless this is specifically permitted.

2 The Virtual Notelet Case Study

As mentioned, the GCS environment has been developed to enable HCI designers to construct context aware applications using off the shelf technology. In order to demonstrate the application of our techniques, we focussed the initial development work on a proposal made by Baldonado et al. (2000), Pascoe (1997) and others. The idea is to implement a system that enables its users to stick virtual notes to physical objects, such as doors or physical objects. These annotations can then be read by holding a PDA close to that object. Alternatively, users might read the notes on their door from any location that offered Internet connectivity, following the model proposed by the Cooltown project. Figure 1 illustrates the general architecture for this application.




Figure 1: Traditional Notelet (left) and Virtual Notelet (right).

2.1 Design Overview

The development proceeded by identifying a number of scenarios. Observing how people used these physical notes in many different ways did this. An augmented reality system cannot hope to match the flexibility of the physical medium. However, we hoped to offer a range of additional functionality that might blend with and support the conventional uses of these physical annotations. For instance, we designed a note that users could update with their location as they moved around a building. Another note simply displayed the individuals published diary from a time management system that was already being used within our target organization. A further template enabled people to link notes to other forms of digital media.

2.2 Pragmatic Concerns and Usability Issues

Many of the design concerns that arose in the design of this system had little to do with the underlying communications technology, that is often a focus for work in

this area. In contrast, it has to do with social and personal concerns. For instance, a current problem with existing notes is that they are visible to anyone who might pass a particular door. Users were concerned to associate password protection or other permission mechanisms to the notes that they left. Other concerns related to the longevity of a note. Users wanted to specify when a communication should be removed and so on. More importantly, the initial design activities were dogged by a form of skepticism that is not recognized by many of the more euphoric attitudes expressed by the proponents of augmented reality systems. The earlier proponents of this technology had stressed the benefits of virtual notelets. However, they often failed to recognize the pragmatic barriers. Many people did not possess PDAs. Others forgot to carry them with them or simply did not attempt to access the digital information. Very importantly, an infrared transmitter lacks the affordances provided by a physical note! Touch sensitive screens in public areas can address some of these problems. However, this again raised security concerns and is far from an ideal solution. An interim conclusion from this work is, therefore, that visionary projects such as CoolTown not only depend on a revolution in the underlying technology. They also require fundamental changes in the way in which people access information as they move around their environment. Our observations revealed a tendency to 'graze' information that was close at hand. Electronic resources, even if they are wireless and Internet enabled, seem to be a last resort for many people.

It should be emphasized that many of these insights only emerged as the implementation was developed. This emphasizes the key point behind this paper; unless HCI developers have access to off the shelf implementation techniques then systems such as the Virtual Notelet will continue to be proposed with little or no supporting evidence for its usefulness. Section 2, therefore, describes the technical problems that arose during this project. about technical factors and finance for implementing Virtual Notelet application and also the human factors will be described and discussed. The conclusion will be discussed in section 3.

3. The Technical Challenges for Off the Shelf Systems

This section reviews the technical issues that arose during the development of the GCS. It also explains how these contributed to the engineering of the Virtual Notelet application. Two primary reasons motivate this presentation. The first is that it illustrates the sorts of issues that developers must consider if they want to validate HCI design ideas for context aware computing. The second is that our techniques might act as a blue-print for other HCI developers.

3.1 Infrared

As mentioned, GCS exploits infrared beacons. These are based on domestic remote controls (i.e. TV, CD player remote control). However, there is an important technical obstacle to communication between these infrared devices and mobile devices' existing transceivers. The protocols and standards used by the

PDA's built-in IR transceivers are different from the physical and data link layers used by domestic remote controls.

The manufacturer of domestic appliances have exploited three different means of encoding infrared signals in consumer devices. They are pulse coded, space coded (AKA REC-80), and shift coded (AKA RC-5) (DuBios III 1991). These coding methods are different from that employed by the infrared transceivers on existing mobile devices. This explains why the Cybergudie system was forced to exploit additional hardware to enable its users to detect infrared signals. The majority of infrared transceivers on mobile devices conform to the Infrared Data Association (IrDA) standard. These standards were developed to support applications whose users needed to frequently transfer significant amounts of data from portable to fixed, portable to portable, and fixed to fixed equipment (i.e. Palm PC to Desktop, Hand-Held PC to Printer, PC to PC). According to the serial infrared physical layer specification from IrDA, It utilizes a RZI (Return to zero-Inverted) coding scheme (IrDa 1998). These technical issues might seem to be a very long way indeed from the claimed benefits of context aware computing but they help to illustrate the pressing need for reliable, off the shelf development environments. In passing it is worth noting that we began this project as a study of context awareness and not of infrared transmission standards!

Domestic remotes send commands in a low-speed burst for distances of up to 30 feet in one direction (BuBios III 1991). They use directed IR with Light Emitting Diodes that have a moderate cone angle. In other words, the signal spreads out from the transmitted in a cone shape. This explains why you can change the channel on your television even if the remote control is not directly pointing at the device. In contrast, PDAs rely upon IrDA to transfer point-to-point data. In other words, it is deliberately intended to support line of sight communication. Data transfer is bi-directional. It operates in high-speed bursts over short distances. It uses directed infrared with LEDs having a narrow cone angle (~30 degree cone) (IrDa 1998). As before, all of these technical considerations might seem to have little to do with mainstream HCI. However, we were again forced to recognize the consequences that they implied for the development of context aware systems. IrDA transmissions require relatively careful aiming; to use this protocol the users of a notelet would have to take careful aim with their device so that it lined up with a transmitter. If any other object comes between the PDA and the transmitter then the signal will fail. These characteristics make it impossible to use IrDA as an architecture for context aware computing, especially given the previous observations about the way in which users require rapid access and low set-up overheads for such information resources.

It is, however, possible to overcome the limitations of the IrDA protocol for our purposes. We have developed software that temporarily disables the existing protocols for infrared reception through the existing ports on most PDAs. Instead, they are programmed to operate in what is termed 'raw infrared mode'. In raw infrared mode, the infrared port can be accessed as a normal serial port with an infrared transceiver attached. The port is no longer IrDA compliant, however, it is then possible to receive signals directly from a domestic control. This supports

the robust transmission of infrared location information and does not force users to carefully position their PDA within the line of sight to an IrDA device.

3.2 Radio frequency Wireless LAN

It should be emphasized that the GCS has a modular design that is consistent with both the Bluetooth standard and with differential radio signal detection. These more elegant approaches can simply be substituted for the infra-red beacons that we currently use as a pragmatic alternative.

In the late 1980s, radio frequency based wireless LANs provided a bandwidth of approximately 250kbps. More recent standards support up to 2Mbps (IEEE 802.11). Nowadays, IEEE 802.11b makes a provision of high communication bandwidth up to 11Mbps (Lough et al., 1997). This raises the potential for wireless LANs to replace more traditional alternatives, especially in buildings that are difficult or expensive to rewire such as schools, hospitals, museums and domestic homes. GCS uses a commercial wireless LAN. The wireless LAN adapters that are installed on the users' mobile devices are compatible with most standard platforms. Infrastructure costs depend on the size of the required coverage and the number of users to be served. However, wireless LANs are marketed as a low-cost solution to conventional networks. The user simply inserts a PC card into their mobile device and installs the relevant driver. The system can provide a transparent network connection from then on. This means that users see no difference between wired and wireless based network interaction. In addition, radio frequency base stations can have a range of approximately up to 1500 feet depending on the construction of the building.

3.3 Network topology for off-line support

The radio frequency local area network (LAN) provides a mobile computing environment for GCS users. However, even with maximal coverage there will be areas in which connection cannot be guaranteed. The physical construction of most buildings creates problems of shielding and interference that cannot easily be resolved. This raises a vast range of usability issues; many of which have not been considered by the more extreme proponents of context aware systems (Ebling et al., 1998). In particular, it seems unlikely that we will be able to solve the problems associated with partial network coverage for many years to come. This comment applies both to indoor and outdoor systems, to cellular and satellite systems.

The communication technologies in GCS are based on radio frequency and infrared. When the users walk around inside the radio frequency and infrared network coverage, they can access data from the location server without any problem. Unfortunately, however, we initially chose to install our system in what can be described as a 'pathological' environment. We ended up demonstrating the technology in a number of Victorian and Edwardian buildings. The walls were extremely thick and these were numerous sources of electromagnetic

interference. This severely reduced the range of the radio transceivers that formed the backbone of our system. Such factors again might seem to have little relationship to human computer interaction. However, it is precisely these issues that developers must address if they are to realize the more visionary objectives of context aware computing.

Fortunately, a number of features in the architecture of the GCS system enabled us to profit from adversity. The problems of achieving uniform radio coverage forced us to consider the usability issues that arise when people move outside our system. This led to an unsurprising but critically important observation: radio transmitters are expensive but infrared transmitters are cheap. As a result, the GCS system continues to provide context aware services even when users move beyond radio coverage. We can deploy dozens of low cost infrared transmitters that index into pre-cached information on each PDA. The cache is then periodically refreshed when the user moves into radio coverage. Figure 2 shows an example of the distribution of radio transmitters and infrared beacons. It also illustrates how a user can walk from radio coverage in cell 1 through an uncovered area to reach cell 2. On their way, they pass by infrared beacons D, B, and A. In order for the user to access information about D, B and A, it must be preloaded before the user leaves cell 1. In other words, GCS must anticipate the users' likely trajectory through the uncovered region. This trajectory is at the heart of the caching mechanisms that support the development of interactive, context aware systems using the GCS.

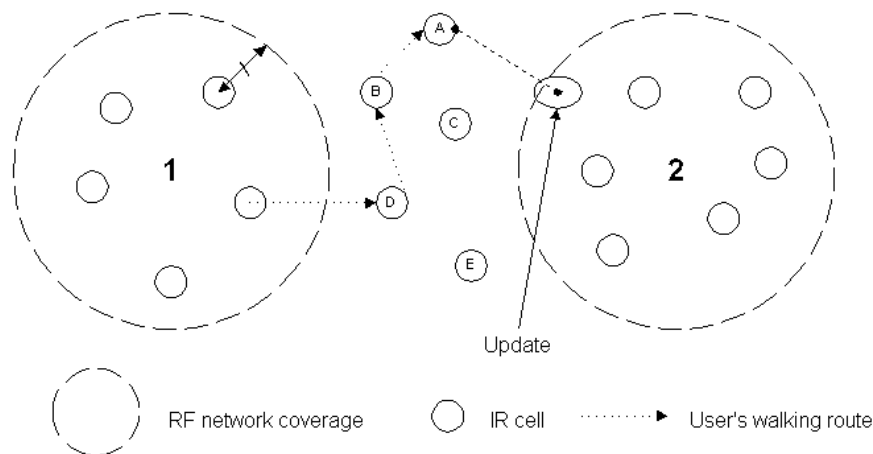


Figure 2: Network topology for off-line support.

As mentioned, the users' cache is updated when the system detects a trajectory that will take them beyond radio coverage. This is done by ensuring that infrared beacons are positioned a short distance inside the perimeter of radio coverage. This distance helps to determine how long GCS has to finish refreshing the information in the users' cache. It is possible to derive a number of equations that support the installation of the GCS network. It should be noted that the following formulae represent gross simplifications. The time taken to transfer location

information from the final beacon and to establish wireless communication is factored into the effective data transfer rate. They do, however, illustrate important properties of the general architecture:

$$\text{Time required to refresh the cache} = \frac{\text{Quantity of data to be transferred}}{\text{Effective data transfer rate}}$$

$$\text{Distance of last beacon from perimeter} = \text{Time required to refresh the cache} * \text{Users average walking speed}$$

For many applications, this predictive caching relies upon a mixture of task analysis and user modeling. It can also be made more adaptive if the system automatically records information about the users previous path when they re-enters radio coverage. The practical development of context aware applications, including the Virtual Notelet case study, has also increased our interest in many issues that have not conventionally been part of human computer interaction. For example, we have become acutely interested in the speed at which people walk. As we have seen, this determines how long we have to refresh the cache before a user leaves cellular coverage. We have also become more interested in architectural modeling. It is far easier to cache information if users must follow a single path along a corridor than it is to predict information needs in a less 'confined' architecture.

4 Conclusion

This paper starts from the premise that context awareness offers many potential benefits to future users of interactive systems. It is, however, difficult to validate the many claims that are made about these applications. Very few interface designers possess the necessary resources to build context aware systems and so there is a lack of empirical data or even of direct practical experience with these applications. Previous approaches have relied upon specialist hardware or have involved re-wiring entire buildings. Future technology offers means of avoiding these implementation issues, for instance through the emerging Bluetooth standard or through differential radio signal processing. Unfortunately, technological and commercial difficulties have delayed the delivery of these potential solutions in a form that can easily be used by HCI developers.

We eventually got tired of waiting for these technologies to mature. This paper, therefore, describes how we set out to engineer an infrastructure for the development of context aware systems. It is worth noting, however, that the GCS has a modular design that is consistent with both the Bluetooth standard and with differential radio signal detection. These more elegant approaches can simply be substituted for the infra-red beacons that we currently use as a pragmatic alternative. The design principles behind this architecture were that it should be low cost, should be modular and easily extensible and that, above all, it should exploit 'off the shelf' hardware. The Glasgow Context Server (GCS) satisfies these requirements. However, previous sections have described the many

technical difficulties that arose during the implementation of this system. We have also described the implications that these engineering issues have for the development of context aware interactive systems. For example, our use of passive infra-red transmitters helps to address some of the security concerns that have been expressed about previous context aware systems.

We have briefly introduced the Virtual Notelet example to illustrate the argument that is presented in this paper. This system is not the central focus for our work. In contrast, the intention has been to demonstrate just how difficult it is to build the context aware applications that many other authors have proposed. These difficulties not only relate to the technical challenges that are addressed by the GCS system. They also relate to the very real human factors issues that often ignored by the proponents of this technology. In particular, two of these issues stand out. Firstly, users seem reluctant to stand in front of an inanimate object consulting their PDA. There is an inclination to search for immediately accessible information sources, physical notes, notices and displays and other people. The electronic resource is a last resort. Secondly, the use of passive infra-red technology does not address all of the security concerns that users have about context aware systems. Tagging inanimate objects with paper notes is a form of social disclosure that is governed by numerous conventions. We have yet to fully understand what those conventions might be like when people from a different organization or country might read those notes. In closing, it is worth reiterating the key point that such key observations will not emerge from the theoretical analysis of potential technologies. They will, however, become readily apparent if more HCI designers had access to the technological infrastructure that is necessary to support context aware systems.

References

- Abowd, G. D., Atkeson, C., Hong, J., Long, S., Kooper, R., and Pinkerton, M. (1997), "Cyberguide: A mobile context-aware tour guide", *ACM Wireless Networks* Vol. 3, Issue 5, pp. 421-433.
- Baldonado, M., Cousins, S., Gwizdka, J., and Paepcke, A. (2000), Notable: At the Intersection of Annotations and Handheld Technology, *in* Thomas, P., and Gellersen, H. W. (Eds.), Springer-Verlag Berlin Heidelberg: pp.100-113. Proceeding of Second International Symposium, HUC 2000, 25th-27th September 2000, Bristol, UK.
- Caswell, D., and Debatty, P. (2000), Creating Web Representations for Places, *in* Thomas, P., and Gellersen, H. W. (Eds.), Springer-Verlag Berlin Heidelberg: pp.114-126. Proceeding of Second International Symposium, HUC 2000, 25th-27th September 2000, Bristol, UK.
- DuBois III J.H. (1991), A Serial-Driven Infrared Remote Controller, <http://www.armory.com/~spcedt/remote/Irremote.html>
- Ebling, M. R., John, B. E., and Satyanarayanan, M. (1998), The Importance of Translucence in Mobile Computing Systems. Proceedings of the 15th ACM Symposium on Operating Systems Principles, May 1998.
- Harter, A., and Hopper, A. (1994), "A Distributed Location System for the Active Office", *IEEE Network*, Vol. 8, No. 1.
- Hohl, F., Kubach, U., Leonhardi, A., Rothermel, K., and Schwehm, M. (1999), Next Century Challenges: Nexus – An Open Global Infrastructure for Spatial-Aware Applications, *ACM Mobicom*: pp. 249-255. Proceedings of the fifth annual ACM/IEEE international conference on Mobile computing and networking, 15th-19th August 1999, Seattle, WA USA.
- IrDA (1998), Infrared Data Association Serial Infrared Physical Layer Specification V. 1.3.
- Ishii, H., Ullmer, B. (1997), Tangible bits: towards seamless interfaces between people, bits and atoms, *in* Pemberton, S (Eds.), *ACM SIGCHI*: pp.234-241. Proceeding of Conference on Human Factors and Computing Systems, 22nd-27th March 1997, Atlanta, GA USA.
- Johnson, C., and Cheng, K. (2000), The Glasgow Context Server: A Wireless System for Location Awareness in Mobile Computing.
- Lough, D., T. Blankenship, T. K., and Krizman, K. J. (1997), A Short Tutorial on Wireless LANs and IEEE 802.11, The Bradley Department of Electrical and

Computer Engineering, Virginia Polytechnic Institute and State University,
Blacksburg, Virginia USA.

Lucent Technologies (1997), WaveLan Wireless LAN, Technology and Market
Background.

Oppermann, R., Specht, M. (1999), Anomadic Information System for Adaptive
Exhibition Guidance, GMD – German National Research Centre for Information
Technology Institute for Applied Information Technology.

Pascoe, J. (1997), The stick-e note architecture: extending the interface beyond the
user, ACM IUI: pp. 261-264. Proceedings of the 1997 international conference on
Intelligent user interfaces 6th-9th January, 1997, Orlando, FL USA

Want, R., Fishkin, K. P., Gujar, A., and Harrison, B. L. (1999), Bridging Physical
and Virtual Worlds with Electronic Tags, ACM SIGCHI: pp. 370-377. Proceeding
of the CHI 99 conference on Human factors in computing systems: the CHI is the
limit, 15th-20th May 1999, Pittsburgh, PA, USA.