

Peer reviewed paper

# proXimity: Walking the Link

Simon Harper, Carole Goble and Stephen Pettitt\*  
Information Management Group, Dept. of Computer Science,  
University of Manchester, Oxford Road, M13 9PL, UK

Email: [simon.harper@man.ac.uk](mailto:simon.harper@man.ac.uk)

\*Salford Electronic Consultants Ltd, Technology House, Lissadel Street, Salford, M6 6AP, UK

Key features: [References](#); [Figures 1, 2](#)

## Abstract

Our society is consistently told that the world is becoming increasingly connected, that the Internet can join physically disparate people by means of email, Web sites, and chatrooms, and that the one 'must have' is a personal domain name; in effect, that the virtual should be more respected than the physical. People are led to believe that computers, with the 'net' as their focus, are their portal to other worlds, their communication mechanism to remote peoples, 'blogging' their primary form of self expression. All this is in part true, but we think there are fundamental issues that are not addressed. The focus on only the virtual is skewing our perception to over-estimate the Web's importance. The increased complexity inherent in all large systems will become too great for many users as the Web develops and grows. The local environment, often most pertinent to the user, is currently completely ignored with regard to dynamic information giving. The Web's focus on information belies the fact that the world is also composed of physical artifacts. Therefore, we think that the next direction for the Web is the conjoining of the physical and virtual. We suggest that they must be connected because without a physical presence the virtual world cannot attain its full potential. To reduce the complexity and stress placed on the user, the Web should relate to the users' physical location and real-world artifacts encountered to make meaningful choices about what information is currently useful or required. In effect, the user acquires a real-world centric view of the Web in which the Web conforms to reality, not reality to the Web. The primary goal of our system, 'proXimity', is to augment realities by giving hypertext, and thus the Web, a physical presence in the real world.

**Keywords:** Mobile Hypermedia, Semantic Web, Mobility, Dynamic Ambient Networks, Universal Access and Control

## 1 Introduction

proXimity is a multi-disciplinary cross-platform system which seeks to widen the Web's scope from the virtual to the real. It is based on our previous work in hypermedia and real-world mobility ([Towel: Real World Mobility on the Web](#)). We analogise the real and virtual, and so aim to provide nodes (link targets), links and anchors in the real world so that hypertext information that describes a physical location also has a physical presence; and so that 'local' physical artifacts can be augmented by 'remote' hypertext and semantic information. Our purpose is to enhance the experience of individuals who would benefit from seamlessly delivered and interface-independent information with a temporal and spatial aspect (a traveller running to catch a flight can be directed to his gate based on artifacts encountered and knowledge of his final destination). We also wish to enable users to interact with aspects of their environment and, where appropriate, provide artifact control structures to enable this (a blind individual

## Contents

- 1 [Introduction](#)
  - [Motivation](#)
  - [The Story So Far](#)
  - [In This Paper...](#)
- 2 [A Vision of Common-Use Hypertext](#)
  - [Current Vision: proXimity](#)
  - [Future Vision: Walking the Link](#)
- 3 [Related Work](#)
  - [GeoNotes](#)
  - [HP Cooltown - websigns](#)
  - [Equator](#)
  - [Auto-ID](#)
  - [HyperTag](#)
- 4 [Conclusion](#)
- [Acknowledgements](#)
- [References](#)
- [Appendix: Ambient computing](#)

moves into the proximity of an ATM, the interface is appropriately recreated on his or her device to suit the individual interaction needs, he or she can now interact with the ATM to withdraw money). We suggest that institutions like public information services, museums, art galleries and digital libraries could use disparate hypertext information sources localised to a specific environment to augment the real-world experiences of their staff and clients. Further, individuals would also benefit, such as:

- clinicians, who could access patient-specific information harvested -- based on the proximity of the patient -- from many online medical resources like medical council journals;
- researchers, who may be able to pass experimental results harvested directly from experimental devices to online resources like the 'Go' database);
- travellers, who could use these devices for mobility information in complex and unfamiliar environments.

proXimity is all about linking the real and the virtual and so the conceptual design of the system is focused on providing this linking, *c.f. via a real-virtual SHIM (Davis et al. 1992, Davis et al. 1996)*. proXimity tries to extend the link metaphor from hypermedia into the real world. We try to link the real and virtual so that users can access complex travel scenarios and locate physical objects in the real world by links pointed to them in the accessible virtual world. To do this we use ambient devices in the role of a real-virtual SHIM to conjoin real and virtual worlds, and we use the physical traveller to 'walk the links'.

## 1.1 Motivation

This is a multi-disciplinary story (much like life) about people and hypertext. It also involves ambient computing, human computer interaction, disability, and augmented realities, but mainly it is about hypertext in the real world. We are writing it because we want you to share our 'blue sky' thinking before we get down to the technicalities.

Nice (on the French Riviera) can be a very artistic place. It has private galleries along its boulevards, state-run museums set in impressive grounds, and artwork - bought municipally - on display in public parks and squares. Strolling along the Boulevard de L'Anglais a promenader may encounter one such publicly displayed work in the shape of a 3 metre mirrored baseball player stooped, with bat extended, ready to strike. On further scrutiny they can find that it is created by a local artist, but the accompanying text is in French (unfortunate if the reader does not speak or read French) and so for that moment in that context it is possible for an individual to be linguistically handicapped by their surroundings.

Visit the Museum of Modern and Contemporary Art, and strangely enough there is an exhibit by an artist called Niki De Saint Phalle. All signs and descriptions are again in French with no translations (there are hundreds of languages so which one to translate to is an issue). But there, in a case, is a small model of the baseball player, Niki's initial sculpture, as a proof of concept for her main work now displayed on the boulevard.

Visitors have a lot of questions to ask about these artworks and the museum. What do other people think of the works? Did the museum (or any other museum in or around Nice) have more sculpture like this? What other information existed about Niki? How can they get back to the Boulevard de L'Anglais to have another look at the real sculpture?

Our principal research area is hypermedia and universal access, and so we started to draw comparisons between a museum visitor's inability to get answers to questions and that of other people in different situations who required information about the local environment and moveable artifacts within that environment. In an ideal world, solutions would exist to answer these questions. Our initial thought was that if we could access hypermedia resources and semantic XML descriptions in the real world based on our proximity to marked artifacts, then things would be easier for lots of people. Further, if we could add some sort of intelligent searching and universal access then we could benefit from the knowledge

contained in many disjoint and disparate resources, and those resources would be accessible by everyone. Finally, if we added the ability to create conjoined real and virtual resources and to expose those resources in a systematic way that included information about location and environment, then we could create maps of the environment which blur the boundary of real and virtual to augment a user's real-world experience.

In this connected world, when a user encounters the big baseball player their Personal Digital Assistant (PDA) reacts and presents information about the sculptor and other works. The user sees what other people think about the work (from virtual democratic annotation services), and are presented (from the intelligent search) with other museums that have work by the same artist. Visitors may choose the Museum of Modern and Contemporary Art because it has an exhibition by this sculptor, and are assisted to its location by the imbedded (invisible and embedded) devices in the environment reacting with their PDA and an XML map on the semantic Web.

On the way to the museum, the user meets others being directed to different locations. One of these is a blind woman on her way to the beach, but she doesn't need any help as she knows where she is, what obstacles are present, and a safe route to the beach. Her mobile device is using the same invisible ambient devices, but the combination of ambient device and user device gives her a completely egocentric experience as her mobile interface is tailored to her, just as another individual's interface is tailored to them.

Users who need some Euros to pay for their entry to the museum are directed to the nearest cash machine (ATM). Their device builds the interface to the cash machine based on descriptions it finds for this ambient device on the Web; all the questions asked by the user and functions accessed by that individual are through the mobile device and its semantic Web interface. Once the transaction is complete the mobile device transmits the final instruction to dispense 20 Euros and a receipt to the ATM.

## 1.2 The Story So Far

In previous work we asserted that lessons learned in the real-world mobility ([Green and Harper 2000](#)) of visually impaired individuals can also be used to solve their mobility problems when moving around hypermedia resources. Further, we likened hypermedia use to travelling in a virtual space, compared it to travelling in a physical space, and introduced the idea of mobility - the ease of travel - as opposed to travel opportunity - the chance to travel. Finally, we created a model of mobility that proposed a set of objects, techniques and principles useful in addressing the travel and mobility issues of visually impaired users ([Goble et al. 2000](#)).

In our continuing work we come full-circle by suggesting that similar concepts are present within the fields of both hypermedia and mobility, and that if the real can be applied to the virtual then the reverse also applies. Hypermedia techniques could be used to augment the real world just as real-world techniques have been used to augment hypermedia.

We think this is important because with the advent of the semantic Web and Tim Berners-Lee's desire to describe resources (many of them real-world resources, such as physical products and artifacts) more fully, the division between real and virtual will become more of a hindrance. Although users can currently search for Web resources from any location, the complexity of hits returned is increasing and a high number of unwanted hits are common (the precision versus recall tradeoff). When searching for physically local resources, the system could exclude hits and reduce complexity if the user's physical location was taken into account. We therefore reason that a system that takes into account physical locality could be useful for many user groups. For example, walking into proximity of Michelangelo's 'David' could display hypertext information from the local Web site, unique artifact histories, and other pre-searched information from the Web along with annotation services and the like.

## 1.3 In This Paper...

This paper has a vision, both current ([section 2.1](#)) and of the future ([section 2.2](#)) as to how we see common-use hypertext encroaching on the real world. We give a [multi-disciplinary background](#) (within [section 2.1](#)) to place the subject in a general context and then address similarities in other projects while highlighting differences with our project, and expand on this in related work ([section 3](#)). Finally, we wrap up with our conclusions ([section 4](#)) and a statement on the real-virtual synthesis of linking.

## 2 A Vision of Common-Use Hypertext

### 2.1 Current Vision: proXimity

We aim to give a real-world presence to hypermedia information by electronically marking physical artifacts and groups of artifacts with imbedded non-networked ambient devices. We aim to move to the virtual world (Web) descriptions of real-world node functionality, interaction protocols and interface specification, conceptual descriptions and groupings of nodes. We propose a connected mobile device (like a PDA or 3G mobile phone) be used as an egocentric user interface and computational hub to control devices and manipulate information structures stored as HTML and XHTML on the Web. In effect, we have devices that represent anchors and nodes, and devices that are interfaces - both types follow hypermedia and Web rhetoric.

We expect minimal computational and interface functions to be present on the imbedded device; rather, the device and the mobile interface will be used in combination. We aim to bring 'remote' information and control found on the Web to the 'local' physical environment, and we expect to use semantic descriptions, disjoint and disparate hypermedia information resources (like the Web), and enhanced searching based on accurate physical proximity to remotely described, physically local artifacts. As we connect our imbedded devices to the Internet using our mobile hub as the interface, we have small infrastructure costs, placement by domain experts is not required, and the user-interface is removed from the physical world and is therefore flexible and egocentric. Also, 'remote' hypermedia resources from the Web or other storage paradigms are constantly searched and presented to the user based on the real-world nodes and anchors, and therefore artifacts may also be moveable (like books in a library for example).


One specific research aim is to identify and resolve the issues in deploying proXimity to support large scale real-world hypertext presence. Although the approach seems straightforward there are significant research challenges to overcome, including:

- Specification of the imbedded device interface as both a mechanism for marking artifacts, nodes and anchor points for specific hypermedia information resources, and as a control interface to real-world resources / artifacts.
- Services to convert between real-world imbedded device identifications and virtual world DNS / IP and URI services.
- Services to convert each way so that by following a hyperlink in a hypertext document the real-world target point can be identified and contextualised so that the user is able to move to the real world presence.
- Resources to define the functionality, control specifications, physical context and location of real-world artifacts.
- Defining any explicitly linked hypermedia resources or search engines with physically relevant search criteria.

As we have stated, with the advent of the semantic Web the division between real and virtual will become more of a hindrance and therefore more of a concern ([Davis et al. 1992](#)). Other projects ([section 3](#)) go some way to address some of our concerns, but the projects most closely related to proXimity inadequately address the problems we envisage. Hewlett Packard Cooltown ([Cooltown 2004](#)) envisions a fully connected environment with Cooltown devices being expensive and internet enabled. Cooltown

websigns (Espinoza *et al.* 2001) only provide access to e-services on a coarsely spatial basis and do not locate artifacts. GeoNotes only provides a graffiti / annotation system based on Global Positioning Systems (GPS). The Equator Interdisciplinary Research Collaboration (Greenhalgh *et al.* 2001), developing projects like Narrative, City, Artequakt and AR, does not address our key issue of creating a universal real-virtual symbiosis through hypertext information and semantic knowledge. Auto-ID (2004) does not utilise semantic descriptions, does not take account of unique artifacts (as opposed to products), and does not take into account ambient device control mechanisms.

In practice our system works by distributing proXimity beacons (like that in Figure 1) into an area. These beacons can be placed such that the infrared (IR) 'cones' are at set angles to a maximum of 33 degrees. The footprint of each beacon, and therefore each ID ('unique' by atomic artifact) is very accurate, and if larger areas of cover are needed multiple beacons with the same ID can be used. In this way the same ID can be used over multiple beacons to denote one physically atomic area or artifact. We use a Sharp Zaurus SL-5500 (Figure 2) as our experimental mobile user device because it has infrared, Bluetooth, wireless networking, and an easily accessible development environment. Once the device comes into contact with IR from a beacon it sends a URL request via the normal browser while encoding the beacon ID to the server. The server translates the ID and sends back the correct information. The user devices keeps a record of transit and sends sequences of information to the server so that the correct direction information is returned. At present the system gives visual information only. This is obviously inappropriate for blind users and so audio will be added.



**Figure 1. Experimental proXimity Beacon (before inserting into housing). Notice the infrared emitter on the top left**





**Figure 2. Experimental user device: a Sharp Zaurus SL-5500 running Qtopia on Linux**

Services to convert between real-world imbedded device identifications and virtual-world DNS / IP and URI services are central to this type of system. These services need to be able to convert each way so that by following a hyperlink in a hypertext document the real-world target point can be identified and contextualised. These links are therefore bi-directional, and this means that the user is able to move either-way through the real and virtual worlds. The proXimity system relies on a SHIM (a thin, often tapered piece of material such as wood, metal or stone, used to fill space between things (Davis *et al.* 1996)) utility introduced between the Web server and the hypermedia resource repository to provide these services. The SHIM is activated when a request is sent to the Web server comprising a number followed by the 'Real Virtual Definition' (RVD) file name extension (345635.rvd) or with a MIME type of 'rvd'. The file does not actually exist as part of the Web server data repository but is present in a separate SHIM repository. The SHIM utility identifies items in the real and virtual worlds by XML descriptions of the real (sculpture) and virtual (hypertext descriptions of the sculpture) contained in the rvd file. These files also contain sets of explicit links to hypermedia resources, search criteria for certain search engines including keywords to search for, and links to annotation services for this particular real-world artifact.

Although still in its early experimental phase, proXimity is our answer to one vision of the semantic Web and hypermedia in that meaning is associated with links, and virtual structures are transferred to the real world. In this way the mobility of visually impaired travellers is increased and, in the future, access to hypertext structures based on their physical closeness to the requestee may be possible.

## 2.2 Future Vision: Walking the Link

We perceive inadequacies in the way the real and virtual worlds currently blend:

1. The presumption that ambient devices will always be statically (either by wire or by wireless means) networked.
2. Building the interface into a device (as is the normal convention) presumes the designer knows a user's interaction requirements.
3. Opportunities are being missed to leverage ambient devices to assist in other tasks beyond those originally intended.
4. Using GPS to gauge proximity to virtually marked out areas of interest is ineffectual because GPS fails to address: artifact mobility, complexity of environments, signal interference and inconsistency in internal and urban environments, and artifact uniqueness (Joest and Stille 2002).
5. Costs of current networked ambient devices are high.
6. Older artifacts that are not electronic or networked are not addressed.

We have a number of underlying beliefs and assertions that form our vision of the future. We think that information will become increasingly characterised by heterogeneous, evolving, disjoint and loosely structured data, of the type found on the Web. Therefore, we see distributed hypertext information as the logical repository for artifact knowledge. In this way individuals and organisations can create and maintain their own information stores for local artifacts. This has the additional benefits of reducing bureaucracy, decreasing infrastructure, and decentralising control. We do not believe that the world will

be fully networked with all real-world artifacts connected, but that ambient devices will be networked in an *ad hoc* manner. So we assert that ambient devices (used for augmenting artifacts and the environment) should not, by convention, be statically networked but should be dynamically networked when a user device is in the proximity of an ambient device. The user devices will therefore become communications hubs transferring information from the currently connected ambient devices to the Web. We believe spatial, temporal semantic descriptions of device functionality, control and artifact information are key to system interoperability, device independence and the creation of open standards. Description languages will be required to augment hypertext information with interface generation instructions and control sequences. We theorise that the interface to any ambient device should be with the user not in the environment and that this interface will be dynamically generated based on individual interface preferences. Therefore we enable universal access to device control and artifact information, and propose leveraging the placement of an ambient device to assist mobility within complex environments. Systems should be 'walk-by' in that the user device reacts to the environment. This means that both the user device and the ambient device must have the possibility of being mobile. As there will be a high number of imbedded devices, the costs of these devices (ambient devices) must be kept small so that they can be profuse.

The nub of our future vision is that the lattice of links that exist as hypertext on the Web will proliferate into the real world; and that in this new conjoined world users will 'walk the links' between Web instances in the real world using information from the virtual world.

We think combinations of technologies will be used. Existing items will be marked retrospectively and given 'identity' while new items will have their identity built in. Appropriate technologies - such as GPS/DGPS, radio markers (bespoke, Bluetooth, wireless LAN), RFID and IR tags - will be used to give different levels of granularity to the activity. We see devices being used in an *ad hoc* way based on user requirements so that the interface will conform to the user and be generated on-the-fly as each tag is encountered. Items can be quiet (passive) or loud (active), and users can virtually annotate real-world objects. Interaction will exist between both virtual and physical, movement of items will be registered and their location altered in the virtual world.

Our most important assertion, however, is that hypertext on the Web will have a semantic, a temporal, and a spatial component. These three components will enable mobility, mapping, direction, and way finding; attaching semantics to links enables portable devices with small screen sizes to order related items dynamically; and temporal components enable movement tracking of both user and real world artifact.

**In the current hypertext world we rely on technology to move from link to link, join the dots, traverse the arcs. In the future, hypertext and the real world will join so that the user can in effect 'walk the link'.**

## 3 Related Work

### 3.1 GeoNotes

Positioning technology is often associated with locating people in geographical space. The GeoNotes system, however, positions pieces of information. GeoNotes allows all users to annotate physical locations with virtual 'notes', which are then pushed to or accessed by other users when in the vicinity. GeoNotes uses GPS and DGPS, although this is only good for larger physical areas. Problems also exist with these technologies because they are not accurate enough to react to exact areas (measures suggest an accuracy level of only 6 - 18 m), they seldom work in buildings and internal environments, and are often inaccurate in complex built-up areas like cities. We also wish to group artifacts, which may have the ability of movement, not just geographic areas, so the GeoNotes paradigm is inappropriate ([Espinoza et al. 2001](#)).

## 3.2 HP Cooltown - websigns

HP Cooltown ([Cooltown 2004](#)) is a vision of a technology future where people, places and things are first-class citizens of the connected world, wired and wireless - a place where e-services meet the physical world, where humans are mobile, devices and services are federated and context-aware, and everything has a Web presence. HP researchers say:

In Cooltown, technology transforms human experience from consumer lifestyles to business processes by enabling mobility. Cooltown is infused with the energy of the online world, and web-based appliances and e-services give you what you need when and where you need it for work, play, life.

HP Cooltown proposes that all devices should be connected. We think this creates an infrastructure cost that is static and unmanageable, interfaces and functionality that are inflexible, and information access that is too specific and therefore negates the intention of hypermedia and the Web. We also propose that devices in the environment are second-class citizens (not first class) and that they cannot be networked without a user device being in range and used as a conduit for device to network communication.

Under the websign model (a sub-project of the main Cooltown effort), a user can point a websign-enabled device (e.g. PDA or phone) in a given direction and see a list of websigns related to the physical structures/objects that appear in the immediate vicinity. Physical structures that have a "websign" are bound to a URL, which contains information about, or a service related to, that structure. In essence, websign creates transparent links between the physical and the virtual and presents these links to the user as the user moves throughout the physical space. Websigns is limited in its scope, however, as it suffers from the problems associated with GeoNotes: it also uses coarse grained virtual signs to denote e-service portals on buildings with corresponding Web sites ([Brignone 2001](#)).

## 3.3 Equator

The central goal of the Equator Interdisciplinary Research Collaboration (IRC) is to promote the integration of the physical with the digital. In particular, it is concerned with uncovering and supporting the variety of possible relationships between physical and digital worlds ([Greenhalgh et al. 2001](#)). The objective in doing this is to improve the quality of everyday life by building and adapting technologies for a range of user groups and application domains. Our proposed project is relevant to the research agenda of the Equator IRC as it relates the digital and physical worlds. However, projects in the Equator stable like Narrative, City, Artequakt and AR do not address our key issue of creating a universal real - virtual symbiosis through hypertext information and semantic knowledge ([Michaelides et al. 2001](#)). This is because these projects rely heavily on deep infrastructure modification and do not have the profuse use of imbedded tags needed for truly ambient device use.

## 3.4 Auto-ID

Auto-ID is a project to give everything in the world a unique ID through the Physical Mark-up Language (PML).

Auto-ID technology will change the way we exchange information and products by merging bits (computers) and atoms (everyday life) in a way that dramatically enhances daily life in the all inclusive global arena of supply and demand - the supply chain.

Auto-ID does not utilise semantic descriptions, does not take account of unique artifacts (as opposed to products), and does not take into account ambient device control mechanisms. The main problem with Auto-ID is that the tags are moveable but the receiver is not. This is because a large amount of power is required to remotely charge a tag and read the returned information. The tag ID is 96-bits long, which also means there is a high possibility of the ID being lost if the receiver is not held next to the tag for a

period of seconds, obviously unacceptable in the course of normal daily activity. Finally, there are privacy issues associated with tags. The main concern centres around the fact that the user is not in control of the information being sent to the receiver. In the proXimity system this is not the case as the receiver is under the control of the user (i.e. the environment is tagged not the user, whereas Auto-ID tags the user's product not the environment).

The Auto-ID Project was disbanded on October 26th, 2003, partly due to privacy concerns. However, development continues commercially and academically ([Auto-ID Labs 2004](#)). The current focus has switched from tagging individual items to tagging grouped items (containers and pallets). As RFID tags are still used in a number of applications, and cataloguing systems to identify these tags are also used, it seems only a matter of time before commercial pressures will encourage companies to reassess the Auto-ID framework.

### 3.5 HyperTag

HyperTag is a system of infrared tags that can be detected by some mobile phones, and with the addition of some software to the phone it can be used to point the onboard browser to a specific Web address. It is only used for delivering single page and source hypertext content; it does not address the issues of interface independence, control of ambient devices, *ad hoc* networking, spatio-temporal searches over hypermedia resources, or semantic descriptions of unique real-world artifacts ([HyperTag 2003](#)). The system is not real time and is not imbedded or ambient because the user must know a tag is present, actively point and press to signal information is required, and then wait for that information to be delivered. It also has a high infrastructure cost and is mainly a form of advertising. Again there are privacy issues here because the supporting advertising server system can easily take a phone ID and issue unsolicited advertisements at will.

## 4 Conclusion

proXimity is directly relevant to the online sector's current interest in the semantic Web, intelligent searching and knowledge management. It is also relevant to mobile and ambient computing researchers and to organisations interested in universal access to real and virtual resources. proXimity is timely in that it uses new and evolving technologies across the knowledge, environment and user domains (semantic Web, ontologies, IR, Bluetooth, GPRS) to create a novel and unique system ([Various 2002](#), [McGinity 2003](#)).

We index physical objects using ambient devices and join the real to the virtual on an *ad hoc* basis, thereby creating large-scale physical networks for information and control. In this way we address many of the problems associated with the partiality and transient nature of current and future networks in semantic and physical space. We then deploy a real application over this combined space to demonstrate the feasibility of our ideas.

In the future we believe that the lattice (an  $n$ -dimensional geometrical structure of sites, connected by bonds) of links which form hypertext in the virtual world will start to encroach on the real world. Once this lattice starts to be extended to include real-world objects at physically discrete locations, truly hyperlinked real objects can be created. We assert that every item, real or virtual, will eventually be subsumed into a 'world' lattice and everything will be related to something else based on location, purpose, type, etc. In the real world these links can be used to view virtual information related to the physical location or object; in the virtual world we can be directed to related physical objects and 'walk the link' to investigate the reality. We can already see the beginnings of this encroachment with the advent of the proXimity project and others like HP Cooltown and Equator.

**The process has already begun, resistance is futile.**

# Acknowledgements

We would like to thank Julian Tomlin (Curator) and the Whitworth Art Gallery for their assistance in the evaluation of the proXimity project. We would also like to thank Bernard Horan and Sun Microsystems for their continued support.

# References

**Auto-ID Labs** (2004) <http://www.autoidlabs.org>

**Brignone, C., S. Pradham, J. Grundback, A. McReynolds and M. Smith** (2001) "Websign: A looking glass for e-services". In *Proceedings of the Tenth IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises (WET ICE)* (Los Alamitos, CA: IEEE Computer Society), pp. 311-312

**Cooltown** (2004) *Cooltown overview* <http://www.cooltown.com/cooltown/index.asp>

**Davis, H., W. Hall, I. Heath, G. Hill and R. Wilkins** (1992) "Towards an integrated information environment with open hypermedia systems". In *European Conference on Hypertext*, pp. 181-190

**Davis, H., A. Lewis and A. Rizk** (1996) "A draft proposal for a standard open hypermedia protocol" (levels 0 and 1: Revision 1.2 - 13th March 1996). In *2nd Workshop on Open Hypermedia Systems*, Washington, March

**Espinoza, F., P. Persson, A. Sandin, H. Nyström, E. Cacciatore and M. Bylund** (2001) "Geonotes: Social and navigational aspects of location-based information systems". In *UbiComp 2001: Ubiquitous Computing, International Conference*, edited by Abowd and Shafer (Berlin: Springer Verlag), pp. 2-17

**Goble, C., S. Harper and R. Stevens** (2000) "The travails of visually impaired web travellers". In *Proceedings of the 11th ACM Conference on Hypertext and Hypermedia* (New York: ACM Press), pp. 1-10

Towel Project <http://www.man.ac.uk/towel>

**Green, P. and S. Harper** (2000) "An integrating framework for electronic aids to support journeys by visually impaired people". In *International Conference on Computers Helping People with Special Needs*, pp. 281-288

**Greenhalgh, C., S. Benford, T. Rodden, R. Anastasi, I. Taylor, M. Flintham, S. Izadi, P. Chandler, B. Koleva and H. Schnadelbach** (2001) *Augmenting reality through coordinated use of diverse interfaces*, Technical Report Equator-01-002

**HyperTag** (2003) *Hypertag technical overview* <http://www.hypertag.co.uk/Technical/Home.view>

**Joest, M. and W. Stille** (2002) "A user-aware tour proposal framework using a hybrid optimization approach". In *Proceedings of the 10th ACM International Symposium on Advances in Geographic Information Systems*

**McGinity, M.** (2003) "Staying Connected: Hibernating During Telecom Winter". *Communications of the ACM*, 46(1):13-15, January

**Michaelides, D. T., D. E. Millard, M. J. Weal and D. C. De Roure** (2001) "Auld leaky: A contextual open hypermedia link server". In *Proceedings of the 7th Workshop on Open Hypermedia Systems, ACM Hypertext 2001 Conference*. Aarhus, Denmark, August (Springer Verlag), pp. 52-64

## **Appendix: Ambient computing**

At its simplest, ambient computing (also known as ubiquitous or pervasive computing) has the goal of activating the world by providing hundreds of wireless computing devices of all scales everywhere. While this concept of generalised computational devices situated through an environment is relatively recent, similar devices specifically to aid the mobility of visually impaired travellers have been in development since 1897. Since then more complex ambient-like devices (generically known as waypoint devices or beacon systems) have been proposed. These systems work by using infrared, radio and inductive or electrostatic technologies to transmit information between devices carried by the user and a device fixed within the environment. When the user moves into range either the beacon - within the environment - or the user device can give feedback. Beacons are often placed at strategic points - say on a platform or railway concourse - to augment implicit waypoints or create additional explicit ones, and the pools of mobility information around them are known as 'information islands'. Ambient systems stem from the belief that people live through their practices and tacit knowledge so that the most powerful things are those that are effectively invisible in use. Therefore, the aim is to make as many of these devices as possible 'invisible' to the user, where applicable. In effect making a system 'invisible' really necessitates the device being highly imbedded and fitted so completely into its surroundings that it is used without even thinking about the interaction.