

Proximity-Sensitive Messaging

Mark Perry and Caroline Umashangar

Brunel University, UK

Abstract: As an approach to the delivery of context-appropriate services, proximity offers a rich and different set of features compared to those of location. We illustrate this by demonstrating how proximity-sensitive services could be utilised in a prototype communications application that supports interpersonal media messaging. Three wireless ‘sensing’ technologies are used to support proximity-sensitive messaging at different levels of distance, loosely characterised as measuring over the short (<50cm), medium (<10 metre) and long (> 10 metre) ranges, allowing message attachment to places and entities at different granularities of scale. This paper also shows how we can utilise the distinctive properties of the various proximity technologies employed to achieve different communicative affordances, as well as allowing access to messages at these different levels of spatial proximity. This type of proximity-sensitive technology offers a far more *relational* form of interaction than simple location, and physical environments that are already becoming saturated with sensor technologies offer a ready-made infrastructure for such proximity-sensitive services.

Keywords: Media messaging; Mobile and collaborative computing; Proximity-dependent and sensitive services.

Proximity-sensitive communication

We often need access to, or find useful, information that relates to people, places or objects that are nearby, and which are not possible to pin down to an absolute grid reference. Indeed, people often leave personal messages and reminders to others in specific places and on objects, the post-it note being a prime example of such a paper-based technology. Yet not all informationally interesting entities reside in a fixed physical location – people and small objects, in particular are more commonly not located in a static space, but dynamically move around a variety of locations, all of which have different configurations and resources associated with them. Moreover, we often wish to mark information at various levels of scale, from making a note on a small object (e.g. a book, a desk, a bag), to marking entities or places at a larger range (e.g. a room, a person, a vehicle), or to post messages over a wide area (e.g. a building, a shopping centre, car park). We may therefore wish to mark messages onto objects that are both mobile and which can only meaningfully be interpreted when the message annotation is scaled to an appropriate distance from the thing that it is appended to. Using location-based technologies to electronically tag these objects and places with information is therefore problematic, and we have examined the notion of proximity, rather than location, as a means of resolving this.

The mobility of computing devices has created the potential for locating mobile devices in close spatial proximity to deliver proximity-sensitive messages to mobile technology users. In this respect, a number of technologies allow us to offer proximity-sensitive services, most notably, Bluetooth, RFID, Barcoding and Infra-red. When used in concert with wireless data communications systems such as GPRS, WiFi, 3G and other forms of network access, these can offer users access to remotely-held information about the electronically enhanced objects that they may *temporarily* become co-located with. However, existing messaging systems

have not evolved to support such messaging among mobile device users, although the potential for this exists, and has been utilised in a limited variety of ways in a number of research projects.

Recent advances in sensing technologies have made it possible to embed and relocate various sensors on people and objects, and within places (echoing Hewlett-Packard's Cooltown slogan 'people, places and things'; Kindberg *et al.*, 2000) in the environment. Such device discovery has the potential to support a wide variety of applications on mobile devices. At its simplest level it may just be used to determine the presence of another device (for example the Lovegetty, see Iwatani, 1998), to supporting far more complex interrelationships: commercial systems already allow mobile devices with integrated sensors to offer communication, data transfer and resource sharing (for example, printers and fax machines) facilities while research developments have further explored the possibilities of using integrated sensors in providing location-based applications on mobile devices. But whilst numerous location-based systems have been implemented using integrated sensors, none of these have so far attempted to use pre-existing and diverse sensor types that allow operation across different environments and sensor infrastructures, nor have they explored the various options supporting the marking of people, places and objects across these environments and infrastructures.

Although associating real-world places and entities with messages is a relatively recent area of interest, much of the 'tagging' technology that enables that is not new: bar-coding products for example has long allowed users to access remote information about the objects that they are stuck to, and this is something that RFID technologies are expected to enhance further. Similarly, we have already begun to see proximity-sensitive signalling, such as dating profiles held on mobile devices that trigger when apparently compatible prospective mates pass nearby systems (known as 'proxidating' or 'bluedating'). A related service called Serendipity (Eagle and Pentland, 2004) makes use of users' online profiles to instigate serendipitous interactions between co-located Bluetooth-enabled devices. Yet these technologies are relatively simple in terms of the services offered, and few offer any great degree of user-configurability or dynamism within the interaction or whilst mobile. They are typically used as electronic initiators for face-to-face communication, or communication in another media (e.g. initiating a subsequent web-based interaction), rather than tools for conducting electronic communication through.

This lack of support for mobile communicative interactivity utilising proximity is not entirely absent: the Nokia 'sensor' application (Nokia, 2007) allows a more sophisticated form of proximal interaction than these other short-range sensor-based applications. The sensor application runs on a mobile telephone, and passively scans the area for other Bluetooth devices running the same application. Users can initiate connections with other Sensor users in the immediate vicinity, and is intended as a social networking tool. The application supports information sharing between users, so that once a connection between devices has been established, users can browse user profiles, and messages can be sent between devices. Notably, the Sensor application is limited to co-present proximal interaction (connections are not persistent over time), and when devices move out of range, connections are lost. It is not intended that users be able to leave messages in places or on people, just to use a sensed proximal connection as an initiator of an electronic communication. Nevertheless, this is an interesting development, and offers an insight into the value and potential utility that commercial developers are beginning to place on proximity-based systems.

Proximity sensitive systems have had some notable developments and research findings in the area of gaming, such as Pirates! (Björk *et al.*, 2001; Falk *et al.*, 2001) and Feeding Yoshi (Bell *et al.*, 2006). It appears that

proximity has a particular role to play in game-playing on mobile devices as developers attempt to build systems that bring mobile participants into physical contact with one another and with the material environment as a means of enriching the gaming experience (which itself is often relatively impoverished on a mobile gaming platform). Pirates!, for example, has a thing-to-thing, and a thing-to-place model of interaction, although proximity-driven events are fixed as to their meaning, and cue interaction with the game rather than enriching other forms of connectedness. Thus in sensed player-to-place events, players are informed that they are near to an island, whilst player-to-player relationships show that they are close to other players (Falk, 2001). As with our own work on the Marked Space system, Pirates! employs a ‘relative’ model of interaction – rather than an absolute model of where people are located in real space.

There is one notable proximity-sensitive communication system: Hummingbird (Redström *et al.*, 1999; Holmquist, Falk and Wigström, 1999) – the first and one of the only systems that we have seen use to the term ‘proximity-sensitive’ in the way that we have applied it here. Hummingbird devices provided users with an awareness of other Hummingbird users in their proximity that had a predetermined wish to communicate. Hummingbird is a particularly interesting research project to us in that it has made explicit its interest in ‘local interactions’ based on proximity as a device for social interaction, albeit on with a strong notion of synchronous interaction to support interpersonal awareness (they call their device an IPAD, or interpersonal awareness device). Notably, the Hummingbird researchers were also interested in integrating people and places into their system, although technologies to support place were relatively poorly developed. The Hummingbird system had a fixed model of proximity – with little or no notion of scale – this is an understanding of proximity as that which could be measured by the computer (in this case, roughly 100 meters), rather than its more human counterpart of varying levels of space and scale. Like our own system developed in the rest of the paper, the Hummingbirds did not depend on an installed infrastructure, making them flexible and open to use in a wide variety of settings. Yet, although proximity was used by Hummingbird as a mechanism to enable interpersonal awareness, and in the research papers that describe it, little is made of the notion of proximity other than its effects on awareness, which is a very limited lens with which to examine this complex notion through.

Alongside this previous research, the widespread use of mobile devices with integrated sensors inspired us to explore the possibilities to use integrated sensors to leave an electronic ‘mark’ on people, place and objects. The work presented in this paper exploits the diverse characteristics of existing wireless sensing technologies in a proximity-sensitive messaging system that incorporates people, places and objects. We present details of our system architecture that is designed to deliver messages with different levels of proximity in both indoor and outdoor environments. This architecture comprises of a single software application running on a mobile device that detects signals from existing wireless devices that our physical environment is already richly populated with. Here, we briefly discuss a generic architecture that utilises embedded sensors to offer proximity-sensitive messaging to mobile device users. We explain how the design exploits the distinctive characteristics of different types of commonly available wireless sensors to resolve the environmental and proximity issues faced in offering messaging across various environments with different levels of proximity. Through this, we explore the properties of the different sensor technologies in terms of the communicative roles and affordances that they offer.

Location-based messaging systems

Whilst traditional, non-computational, technologies such as paper and pins, sticky tack and tape, and recently, PostIt notes, have evolved to support communication that enables proximal or location-based relationships, most electronic communication systems do not have such a sense of place, or indeed, offer the interactional potential to *be* placed. Information and communications are received and sent from device to device, irrespective of the relationships between those devices and the other entities that surround them. Although mobile technology allows us to communicate ‘anytime, anywhere’, this is not necessarily *appropriate* to its location (see, for e.g. Perry, *et al.* 2001) because of the very *placelessness* of that communication and its temporal decontextualisation. Such communications technologies do not, for example, allow a user to send a message to someone reminding them to post a letter on their way home that will only be received when that person leaves their office. This is a serious failing in the promise of computing in enabling richer interaction with the world, and enabling a symbolic layering of content over the things that we encounter in the world around us.

In practice, most existing location-based technologies have focused on *locating*, but not *marking*. However, researchers have begun developing a number of outdoor and indoor location-based systems that allow users to locate devices and to offer variety of services such as tracking, messaging and data collection. One of the best-known messaging systems is ComMotion (Schmandt, *et al.*, 2000); this is a typical example of a location-based service that supports personal messaging on mobile devices. ComMotion allowed its users to leave reminders and messages in frequently visited locations. These locations are not user-defined, but automatically generated by the device.

Following developments from the ComMotion project, GeoNotes (Espinoza *et al.* 2001) has attempted to resolve some of the volume problems that would be expected occur as the number of messages scales up by looking at interaction issues in filtering notes and messages. Both of the ComMotion and GeoNotes projects have utilised fixed location-based technologies (although sensor types have changed over their development lifecycles), and did not allow users to utilise a variety of sensors or to control the distances over which messaging could be conducted. Similar to these projects, Skyline (Multaharju *et al.*, 2004), Egraffiti (Burrell and Gay, 2002) and PlaceMail (Ludford *et al.*, 2007) are public messaging systems that allow users to view and post messages at their current location. These systems are both based on wireless LAN (WLAN) technology, and so again only support indoor messaging at a single proximity range on a single sensor type. The Urban Tapestries project (Lane, 2003) also allows users to create and access public annotations using GPS location-enabled networked handheld computers to build ‘urban histories’ of places.

Messaging systems need not be wholly electronically mediated through sensor discovery, although this can increase the cognitive and physical effort associated with marking and receiving messages. For example, the Yellow Arrow project (Yellowarrow.net) allows mobile telephone users to generate publicly accessible content that is linked to spatially located and physically embodied messaging objects. Participants stick uniquely identifiable arrows at different locations and objects in the world, and messages are associated with the arrows by sending a text message (SMS) from their mobile telephone to a specific telephone number beginning with the arrow’s unique code, along with a message in the text that becomes linked to the location where the sticker is placed. When another person encounters the arrow, he or she can send the unique code on this arrow to the visible number and will immediately receive the associated message on their mobile phone as a text reply. An

associated website extends this location-based exchange by allowing participants to annotate their arrows with photographs and maps in an online gallery of tagged yellow arrows.

Geotagging allows geographic metadata to be added to electronic information such as photographs (e.g. flickr.com) or websites (e.g. geocaching.com). This is necessarily specific to a fixed location, rather than a proximal relationship, although an element of communication may be involved, either physically in the location itself (as a physical message or graffiti), or in messages associated with the geotagged location that can be accessed over a web connection. A more interactive experience of geotagging can be seen in Hewlett-Packard's mscape system (mscapers.com) presents a 'mediascape' that attempts to deliver context-aware services, such as media elements that are associated with tagged GPS locations, and which can be accessed (or more typically, media are automatically played) as users move around those spaces on their mobile devices. Typically, elements within these mediascapes are not intended to be generated in real-time, and so they offer a weak interactive element necessary for communication, but this platform (although tied as it is to GPS location data) does offer potential for rich media messaging, albeit one that is currently unrealised.

One other project is worthy of mention: the SLAM (Scalable Location-Aware Monitoring) system at MIT, featuring the Cricket indoor location system, which is intended to support wide range of tracking and controlling applications (Priyantha *et al.*, 2000). Although this is not tailored for messaging services, the system uses a similar technique to the one we have proposed to resolve the restricted coverage limitations in location monitoring systems. It achieves this using multi-sensor discovery from different technologies such as ultrasonic, GPS and RFID sensors to provide coverage across various environments and distances. The clear distinction between SLAM and what we discuss here is that SLAM uses purposely designed ultrasonic beacons to offer high accuracy in indoor environments whereas we have based our design on wireless sensors that are already a part of our existing electronic environmental and personal infrastructures. Thus the similarities between SLAM and our own system are more to do with the nature of sensor discovery than the purposes to which the systems are put. In order to explain the function and architecture of our own system, we now turn to describe this in detail.

The 'Marked Space' System

The proximity-based messaging system that we have developed is known as the 'Marked Space' system. It is intended to allow users to generate messages in a variety of media (text, images and audio) 'in situ' on their mobile devices and to associate them with electronically identifiable entities or places. These messages can then be accessed asynchronously by their recipients when triggered by their coming into proximity with the entities or places associated with the message. Messages can be associated with three levels of proximity; these levels of proximity can be generally interpreted as the zones of the body, the social and the locality (note that sensors may not be available for all three levels of proximity at any one time due to their physical availability or presence). Of course, messages may never be accessed by their recipient – unlike an SMS/text message that is likely to be delivered irrespective of its recipient's location, a Marked Space messaging system user with messages awaiting them may not come into range of their message triggers. This is intentional, and should not be seen as a disadvantage of the system – messages are intended to be connected to people, places and things, and if these are not physically encountered, then their content is unlikely to be highly relevant.

In its design, our system is very different to the other forms of messaging system described: we have

deliberately focused on user interaction with the interaction being drive by proximal relationships (and not by location). Within this notion of proximity-driven interaction, we are interested in scalar issues of proximity, this being a relational association that is naturally derived from spatial distance. Having three levels of distance will allow us to explore both the technical issues related to the implementation and how convergence operates in practice, as well as the social and interactional repercussions of different levels of proximal distance as users struggle to make sense of how others will interpret their messages and how they can generate their messages to be used in the way that they are intended, and how messages will be interpreted and made use of by their recipients. Nevertheless, this is an ongoing research project, and the current research objectives are to develop the architecture of the system and explore its potential for sense-making and its application in interaction design. At the core of this is the development and exploration of proximity-based interaction on mobile devices, with a secondary element looking at how this can be supported over different levels of proximity that are enabled through a multisensor architecture. Utilising a multisensor architecture supports a degree of robustness, so that messaging is possible in a variety of environments where the availability of sensor infrastructures is unpredictable.

In order to understand the operation of the Marked Space system, and how this will lead to implications for its use and interaction design, we need to examine its underlying design. In this respect, the Marked Space system is composed of three notable functions with respect to its proximity-sensitivity: device discovery, marking, and proximity matching. We discuss each of these in turn below.

Device discovery

In order to offer proximity-sensitive messaging among mobile device users, it is necessary to discover these devices in the environment. Mobile devices are already capable of using integrated sensors (e.g. Bluetooth) to proactively find themselves in relation to other sensors that are in close enough proximity (see Treu and Küpper, 2005, for a discussion on the distinction between proactive and reactive proximity detection). We have therefore explored the possibility of using the proximity between devices to piggyback proximity-sensitive messaging services to mobile device users without building a new infrastructure. To provide such messaging services across all environments it was necessary to use sensors that are capable of offering coverage in various environments. No current single sensing technology is capable of providing coverage across indoor and outdoor environments and across a different proximity ranges. We have exploited current sensing technologies and their distinctive coverage capabilities to resolve this problem. In the proposed system, we bring together different types of sensors that work in and across different environments into a unified system that provides pervasive coverage in both outdoor and indoor environments.

Many sensors and sensor-enabled devices offer the opportunity for users to understand more about what the sensors are connected to. For example, Bluetooth sensors (and to a lesser extent WLAN-enabled/802.11 devices and access points) allows users to discover more details about their context than acting as a simple identifier: as well as providing a unique MAC address, they can be queried to provide additional details that allow us to determine (or at least make a reasonable guess about) their location, whether they are fixed or mobile, their forms of use or the identity of their user/s. It is worth emphasising that this information is already available to mobile device users to support such messaging, and does not require any information to be added to the sensors to enable this. This information allows senders and recipients to make reasonable judgements about what or who

is being tagged and allows senders to form an opinion on how their tagged messages might be interpreted by their recipients. Location data from GPS co-ordinates are less specific as to their contextual relationships, but due to the physical size of the area that can be marked (this is not limited to co-present sensor range), this is less of a problem, as long as this wider messaging range is clear to message senders and recipients. Similarly, RFID and other very short-range and manually-selected sensor types (including printed barcodes) require users to be physically very close to tagged items, and so the relationships between the object and the sensor and the are likely to be clearly evident from this. Nevertheless, these sensors typically offer up little recoverable contextual information about their relationships (for example, RFID tags are not normally labelled with information about who they were placed by, or what specifically about the object that they are placed on they relate to). Senders may therefore need to manage their message content to allow their recipients to make appropriate connections between the message and the entities or places that the message relates to.

Marking

Marking involves making electronic annotations using a variety of media on people, places and objects. This marking facility gives the opportunity to mobile device users to choose a person, place or an object and to leave a message attached to this for other users of the system to retrieve. Whilst several systems like ComMotion (Schmandt *et al.*, 2000) have addressed aspects of this marking in the past, we have added more flexibility into our system by allowing users to choose any person, place or an object by simply discovering them in the environment through the use of integrated sensors.

In the Marked Space system, once the sender's device discovers other sensors and devices, it identifies them and users can attach messages to them. Messages are not physically stored on the marked device or sensor, but are tagged against a unique identifier (such as a MAC address), these details are stored on a central messaging server's database. When the message recipient's device discovers a marked sensor, and they are identified as its recipient, message delivery is automatically triggered. For reasons of optimising connection to the remote messaging server, messages are downloaded and stored locally on recipients' devices until their devices discover a marked sensor. At this point message delivery is triggered to the mobile device, and the message becomes accessible in the user interface. Notably, this message caching is purely intended as a practical solution to the problems of network latency and connection costs; given a faster network connection and more economical data connectivity charges, we would have used a real-time, proximity triggered message delivery system, as this would have a number of advantages, including problems arising as the number of potentially large size multimedia messages scaled up with their attendant memory demands on the mobile device.

If the sender's device fails to discover any devices to mark people, place or an object to attach messages, a user can place a small, non-complex and cheap sensor, such as passive Radio Frequency Identity (RFID) tags, on the object to be marked. The increasing prevalence of already existing RFID tags within the commercial and consumer environments means that it is increasingly likely that such sensors are likely to be increasingly useful and available for supporting proximity-sensitive services. Similarly, if the environment or nature of the message are not well supported by local device or sensor co-discovery, we have used the Global Positioning System (GPS) to provide us with spatial co-ordinates that we can associate with messages, allowing users to send and receive messages when they are able to pick up a GPS signal.

Proximity matching

Proximity matching involves delimiting and determining the proximity range of a marked physical area through the use of sensor discovery. There are many instances where people want to deliver messages to other mobile device users within a particular proximity range: by limiting message delivery to a particular range users can post messages where they are expected to be most relevant. For the purpose of our design, we have grouped proximity range into short range (a few centimetres) utilising RFID technology, medium range (up to 10 metres) utilising Bluetooth signals, and long range (beyond 10 metres) utilising the GPS system. No single type of transponder or sensor has the potential to accurately cover all three ranges and we therefore took advantage of sensors with different proximity ranges and integrated them into a single system to work in a complementary manner. Notably, this choice of sensors to associate messages with also allows users to make use of the particular properties of the various sensors beyond that of simple sensor range.

System design

In order to explain how the prototype Marked Space system has been developed to support proximity-sensitive messaging, we need to explore the basic architecture of the system. This section explains the generic system architecture and its functionality without going into the details of its implementation. We have used an ‘open’ architecture designed to adapt to and accommodate the different kinds of sensing technologies available for mobile devices. Our rationale for keeping the architecture open is so that future developers are free to choose sensor technologies and to integrate them into the system to meet their different messaging requirements and to keep pace with technology developments and sensor availability. The current user interface to the system is presented below, showing both the send and receive functionality (fig. 1). This interface is being used for testing purposes, and does not include all of the functionality and features that the architecture of the system has the potential to support: the reason for its inclusion in the paper at this stage is to visually demonstrate the functionality of the system relating to the proximity and message-related information that is available to users.



fig. 1. Prototype interface to the Marked Space system.

As can be seen, figure 1 shows three radio buttons at the top of the interface that become activated according to the services available: thus if a GPS signal is available 'Long' will be enabled, if another Bluetooth device is discovered, 'Medium' will be enabled, and if an RFID tag is detected, 'Short' will be enabled. The user can choose to 'attach' a message an appropriate environmental entity to using these radio buttons. If a user wishes to send a message, they will need select an enabled radio button, then touch the 'To:' menu to choose a contact name to send the message to. This will list potential recipients' phone numbers for the user to select, allowing an SMS message to be sent (in the prototype they may also select an email address). Please note that in the system's current manifestation, SMS has been used as the primary medium through which messages are sent, although email is used for testing purposes. We are open minded as to how messages can be best propagated to recipients, and this use of the short messaging service is a pragmatic solution to our system testing needs. The 'Tag To:' menu is a dynamic menu that provides a list of discovered sensors or the current GPS location (this is dependent on the choice of range/sensor type that users have selected from the radio buttons. Thus, if the user has chosen to send a message to a medium range sensor, it will display all the Bluetooth sensors available, and will include both its MAC address and 'friendly' name). The text box below this ('Write your message here') allows user to enter the text of their message. The Marked Space architecture allows other forms of media to be sent other than text, but for development and testing purposes, this has been limited to text annotations. Below this, the interface elements relate to received messages.

The first box in the received messages part of the interface lists all received messages that relate to the current environment. This box displays the sender's address (either a phone number or email address), alongside details of the sensors that have been detected in the environment with their identifying details. Users can highlight a message from this list and see further details of this in the next two boxes: the first box displays sender details,

while the second box shows the content of the message. It needs to be recognised that this interface design is not intended as a final product, and has not undergone user testing and does not include the full range of functionality that we have planned, and although this not present in the interface, it is supported by the software and system design and can be relatively easily added at a later date. These include additional features such as a list of sensors that have been previously encountered, with their associated messages. These will be developed as the project progresses. The information available at the interface is drawn from the underlying messaging architecture, and this is described in the following section.

Messaging Architecture

Four functional components are responsible for the operation of our proximity-sensitive messaging system:

1. The Discovering and Marking Component (DMC) is responsible for finding the signals from proximate environmental sensors and for associating messages with the unique identifying characteristics of these sensors.
2. The User Interface (UI) provides a front end to the application, allowing senders to identify proximity range that messages are associated with and for recipients to make sense of what received messages are attached to, as well as authoring, filtering, editing and personalising messages. This can be seen in fig. 1.
3. The Database Engine (DE) provides data management necessary for storing and maintaining messages and their associated information on a remote central messaging server. It also manages message routing between sender and recipient mobile devices.
4. The Link Processor (LP) subsystem provides an internal interface between various sensor and other components of the system within the mobile messaging device. It takes its input from the UI and passes it to the DMC to initiate device discovery then gathers information about any discoveries, feeding this back to the UI. The LP also takes care of communication between the UI and DE, connecting wirelessly with the external messaging server to deliver and receive messages.

The interrelationships between these four components are shown in Figure 2, which gives an overview of the generic system architecture. The arrows in the diagram indicate the direction of information flow between components. The DMC, UI and LP components are embedded into the mobile device while the DE is hosted remotely on a server, and can be accessed by multiple mobile devices.




Figure 2: Messaging Architecture

For the purposes of this paper, we are interested in exploring the characteristics of the three sensors to describe their role in proximity-sensitive messaging, and this can be best examined through looking more closely at the DMC component. Figure 3 shows a more detailed technical architecture for the DMC design. The DMC controls three different types of sensors on its user's mobile device to discover and mark at different levels of proximity. These sensors offer short, medium and long range coverage and are referenced in the diagram as *integrated* sensors. Sensors attached to people, places and objects in the environment are distinguished from the Marked Space device in the diagram as *environmental* sensors, or in the case of GPS, as an environmental signal. The DMC application controls short (RFID), medium (Bluetooth) and long range (GPS) integrated sensors and ensures that they discover and communicate with any short, medium and long range environmental sensors (respectively) that they encounter.




Figure 3: Technical architecture of the Discovering and Marking Component (DMC)

The DMC performs a location update at regular intervals and keeps records of all the sensors that are in close proximity, i.e. within the coverage range of the mobile device's various sensors. If the device is in the known presence of a sensor (e.g. a visible RFID tag), a user may initiate a sensor update to identify it before sending or receiving a message.

Different Communicative Affordances

Whilst they offer similar potential for marking, the three sensors that the system utilises are very different in terms of their operation. Two of these sensors (RFID and Bluetooth) are 'environmental', in that they require two sensors to be in physical proximity to one another, whilst the other (GPS) is carried by the user and does not require a 'carrier' device to be in proximal range. It is therefore location-based, though proximally triggered. We discuss the sensors that we have used in our prototype (and the communicative affordances for each of them) in turn below:

The short-range sensor selected for use (passive RFID) has a coverage of only a few centimetres. Each RFID tag has a unique identifying code. In terms of communication, the advantages of RFID lie in its precise sensitivity,

and as it is an un-powered sensor, it allows relatively cheap ‘fit-and-forget’ placement. In order to make contact with a passive RFID tag, users need to bring their devices into very close range with the tag, most likely in a deliberate act of swiping over it (although chance proximal connections are possible). This has an interesting side-effect in that the requirement for such deliberateness of interaction has a presumption of interest for the recipient: they are not likely to enable such a proximal interaction without their having an interest in finding out whether a message has been left for them. Thus, the automaticity of message delivery is itself mediated through a deliberate act of ‘finding out’ if a message has been left. Senders of messages are therefore likely to (at least if they understand how the system needs to be used) develop their messages in a form that fits with this act of deliberate interaction. They are also able to ensure (to a certain level) that their recipient is clearly able to relate the received message to the entity to which that message is attached, and that it will not simply ‘appear’ on the recipient’s device without an indication of *what* it is attached to.

The medium range sensor selected for use (Bluetooth) has an intermediate coverage of several metres, although they require a power source. Bluetooth signal transmission may be weakened by metal structures (which may include some walls and other building structures), an apparent disadvantage that has a positive side effect of allowing room-level location sensitivity in some instances. Many devices in the environment are already Bluetooth enabled (e.g. cellphones, wireless headsets, vehicle navigation systems, media players, notebook and desktop computers, printers, mice, keyboard, medical devices, gaming devices, camcorders, projectors and digital cameras). As of November 2006, there is an estimated installed base of over 1 billion Bluetooth products in some form or another (Bluetooth SIG, 2006) giving a huge existing infrastructure for message placement: this is a huge advantage over other, less numerous wireless forms of sensor. Moreover, it is possible to distinguish a good deal of information from the Bluetooth identifying signal beyond its unique identifier (its MAC address), including the type of device (‘device class’ e.g. mobile telephone, PDA, headphones, etc.) and a user generated ‘device name’ for that device (e.g. “Mark’s Palmtop”). Such information allows senders and recipients of messages to make better-informed judgements of when and under what circumstances their messages will be received and how to interpret their content. One of the particular features of Bluetooth sensors of relevance to our work is that these are commonly used in mobile telephones and for practical reasons they are likely to be ‘discoverable’, thus enabling person (or more accurately, intimate personal device) discovery and marking.

The long range sensor selected (GPS) can locate a signal at varying ranges of precision. It is unlike the other two sensors which need to detect other devices within signalling range. While the location information provided by RFID and Bluetooth is *relative* (because they are only able to determine their location in relation to the signal recipient), the GPS sensor positioning is *absolute*, in that the location of the sensor (unlike the other two forms above) is mapped within geographic space: relationships between the device and its associated proximal message identifier are therefore inferred, rather than directly sensed. This has an interesting side-effect that is again different to the other two forms of sensor: for *relative* sensors, the sensor and the marked spatial location need to be within range to activate—this is not the case for an *absolute* positional sensor, which is able to specify its current proximity to a tagged place or object (if it is immobile), independent of any environmental markers. One potential disadvantage of (unassisted) GPS is that it is sensitive to signal loss within structures, although as we later discuss, this may have a communicative relevance to users. To allow a broad, yet meaningful scale of proximity, we have bounded the GPS range to within 100 metres by using matched x and y co-ordinates ± 50

metres. This range of course has scope to be adapted (possibly across multiple or even user-defined ranges) following user evaluation in the field, although this has yet to be completed.

The differences between these sensor variations means that applications built on top of them can be put to different uses – the reasons why a user would tag a space to a range of a few centimetres are likely to be different to why they would do so where the range and accuracy of the tag is in the order of a radius of 5, or 100 metres. Whilst each of these sensor technologies may be used independently, they can also be used in conjunction with one another, due to the different capabilities and constraints that they carry.

There are other types of sensor that could be added to the system with different constraints to those that we have discussed, some of which may carry additional constraints and offer new opportunities for developing proximity-sensitive services. Indeed, the wireless environment has an incredibly rich existing infrastructure of uniquely identifiable resources, and provides a ready-made platform for developing proximity-sensitive applications. Examples of this within an indoor environment include potentially integrating infra-red and ultrasound beacons (e.g. Randell *et al.* 2002) and another related technology, audio networking (Madhavapeddy *et al.*, 2003), in which spaces and things can be easily and cheaply demarcated with inaudible audio signals running off commonly available sound hardware, and which can be picked up with microphone on mobile devices. These technologies offer room-level precision, as infra-red and audio signals are bounded by walls, and this may be a useful property for controlling the range of the broadcast signal. As these are not already commonly available in the environment, this infrastructure would need to be created; however, this is not to say that such devices would not necessarily become more mainstream at some later date. Another technology that could be integrated is Wireless LAN, as has been seen in a number of existing location-based projects as noted earlier: WLANs have a greater range than Bluetooth, although they are likely to be less frequently encountered, are less likely to be mobile (and therefore act as a personal signifier), and like ultrasound beacons and audio networking, typically have less metadata with which to make interpretations about the context of the message (or proximity-related service) discovered. Even the humble barcode offers opportunities for marking objects, and has the advantage of being extremely cheap to place. Indeed, like the barcode, any uniquely identifiable and easily captured information media can be used for marking, ranging from written and manually inputted telephone numbers to visual tags (e.g. Madhavapeddy *et al.*, 2004), icons or pictorial images that can be automatically recognised by the camera in a mobile telephone. Where electrical power is available, it may be possible to physically mark environmental sensors to either support precise sensor discovery (e.g. determining what messages relate to), or simply to determine that a message has been associated with a sensor, and that it should be investigated further (e.g. by scanning a passive RFID tag). If such sensors were networked, this interactivity could be made specific to individual users, as sensors recognised that they had an association with a message recipient. This set of suggestions is not to say that we would recommend such developments, rather they serve to illustrate the different properties of messaging techniques and sensors, and to illustrate the resultant implications for messaging for use that they might engender: most have different communicative affordances, offer different opportunities for use and would subsequently give rise to different opportunities for framing the communication. Clearly, the particular constraints and properties of sensors used is likely to have a major impact on the ways that they are incorporated in use practices, and how we might need to consider adapting interaction designs to accommodate these different proximity-supporting media.

Discussion

The provision of access to information *in place* and *relating to that place* is poorly supported in existing mobile and ubiquitous computing technologies. Current computer technologies tend to serve as central repositories for electronic data and services, housing our email, instant messaging and web pages, while mobile telephones and handheld devices attempt to provide mobile support that is also centralised on a device. These existing electronic hubs provide us with access to all of our electronic information—whether it is relevant to our current context or not. Proximity-sensitive systems provide a solution to this disconnect. Of course, developing proximity-sensitive systems is not a simple matter for all manner of technical (e.g. battery limitations, device and infrastructural limitations, vested commercial interests) and interactional reasons (e.g. how to represent sensors to senders and recipients), some of which we have encountered during the development of our own prototype, others of which are likely to fall out of detailed user studies or through the commercialisation of such proximity-sensitive systems. Below, we draw some lessons in developing proximity-sensitive systems drawn from our experience so far of developing the marked space system:

- 1) We have attempted to make use of existing environmental sensors for two reasons, and would recommend this to other researchers and developers. The first of these is cost: the expense of developing a rich and pervasive sensor infrastructure is far more expensive than simply using an existing set of sensor resources. The second reason for this choice is that there are an increasing numbers of sensors that are already available on mobile devices that can detect signals that are already pervasive within the environment, and this means that people do not need to add to their existing complement of devices that they carry around with them: proximity-sensitive applications can be developed relatively easily on top of the currently and commercially available devices without costly and complex customisations, and the reliability problems that these are likely to introduce.
- 2) There is value to proximity-sensitive systems in supporting different levels of proximity. The notion of proximity is a relative one – and depending on the form of the marked entity, the notion of ‘nearby’ is only meaningful at different level of scale. Furthermore, utilising a single level of proximity is communicatively limiting for reasons of access (it provides fewer opportunities for discovering a proximal relationship), spatial association (only a single relational distances between sensors could be associated with a service) and its communicative affordances (a single level of proximity affords less variations on content interpretation). However, these different levels of proximity should be clearly made known to users, so that they are able to make appropriate associations as to their use or meaning.
- 3) The user interface to proximity sensitive systems should allow users to make better informed judgements about the constraints of the various sensors (e.g. robustness and range of signal reception) that they are utilising. This is not something that is naturally visible to users, particularly as these signals are wirelessly transmitted and often invisible, and this needs to be made apparent in the interface of the mobile device in some way so that it is possible to determine these constraints and so use and interpret the system appropriately. Although this could be done explicitly, through user manuals, or on-screen instructions, this could also be achieved by allowing the system’s users to easily visualise the constraints of the sensors, for example, showing them when GPS signals are lost when inside buildings.

- 4) Placing and locating messages is not a simple matter: senders may not fully realise or understand where or on what they are placing a message; similarly, recipients may be awaiting a message and not know precisely where to receive it, or they may receive a message, but not know where or what it is attached to. We can build on top of existing metadata, such as Bluetooth's device class and name, which allows users some insight into the message's relationship with the physical world actually is, but much of this must be inferred. This is likely to be easier with GPS signals and RFID sensors, although, for example, users need to know that their GPS location has been updated and is current, and given that sometimes slow signals or where the GPS signal has been blocked by physical structures, this is not always the case; similarly, it is possible that in some cases, such as multiple RFID sensors stuck in the pages of a book, several sensors may be triggered simultaneously, and it will be hard to determine precise message-sensor relationships. There is a clear design relevant point to make about making as much information available to users as to the nature of the signals: making visible to the user the type of sensors that can be connected to messages or that messages are attached to, their signal range, how recently the sensor was detected, and making relevant metadata relating to the message sensor available are clearly valuable in supporting users in making meaningful interpretations about the message. This point can be extended beyond proximity-sensitive messaging to other proximity-sensitive services. Some of these design considerations are clearly dependent on whether these are multi-sensor systems, but making the nature of the service-sensor relationship is important in determining precisely what, where and how this relationship is embodied.
- 5) Just as we can associate messages with people, places and things, so can we use this information for tracking whereabouts, introducing potential problems in intrusions into personal privacy and the ability to track where people have been and what they have come into contact with. This is not an issue wholly unique to this technology, and has also been considered in other location-based technology developments such as Place Lab (Hong *et al.* 2003), and Reno (Consolvo *et al.*, 2005). However, unlike location-based systems, the notion of location within Marked Space is inferred from proximity (other than by its GPS component), which adds a degree of fuzziness to a user's whereabouts. Whilst proximity relationships may be equally interesting and useful to nefarious users of the system, this means that the precision as to where users of the system have physically been is less easy to determine. Of course, may be possible to tell what entities that users have come into contact with from accessing the Database Engine, although in practice, as the system is currently designed, metadata about environmental proximity sensors is managed at the level of the local device, and the only information sent to the remote server (and used in the Database Engine) are the abstract, unique identifier details, such as RFID codes and Bluetooth MAC addresses that are much harder to identify with a particular person, place or thing. It is important to note that we are not making any more information about people or devices available than is already present: we are simply utilising the information that users are already making available through, for example, their use of a Bluetooth enabled cellphone. Users are already publicly broadcasting this information, and we are simply making use of this as an enabler for determining proximity. Further to this, in the system that we have developed, there is explicitly no feedback to message senders that their messages have been received (as can be requested in SMS/text messaging for example) so that user activities cannot be traced. GPS data may transmitted back to the remote server (although not in our current design which relies on caching), but again, this is not stored as would be required for a tracking application, rather, incoming GPS location information would merely be

matched to the database to see if a messaging match had been made and then discarded. There is no need to store information on a user's location for the purposes of this application, so that privacy concerns should be allayed; of course, whether users will be aware of this model of information use by the system is yet to be determined.

One point of caution must be made that relates broadly to privacy: by allowing proximity-based tagging, we make it possible to attach messages to people in a way that they might be uncomfortable, to give a puerile example, it may be analogous to pinning a 'kick me' message onto someone's back. The person to whom the message is attached is likely to be completely unaware that they have been tagged or labelled, and may have no mechanism to remove such messages (other than to turn off or put down their devices) even were they to realise that such a message had been left. This is something of a hard problem to resolve. Solutions may be to allow people to examine and edit messages that are associated with sensors that they are responsible for, although this is itself open to abuse. It is likely that were the system to become widespread that we would see an etiquette developing around such proximity-sensitive messaging and other services that would provide a set of acceptable forms of behaviour.

- 6) The idea that sensors could be 'improved' so that signals were more accurate or could be received in a wider range of environments may not be as useful as it initially appears. The very fact that constraints are imposed on system use as it currently operates may allow users to make better informed judgements on how their message might be interpreted or about what the sender had intended: that a message could not be received and is understood not to be receivable in certain contexts (or vice versa) is important in enriching the meaningfulness of communication, and ensuring uniformly pervasive reception across sensor types may have detrimental communicative value. The utility of this seamfulness (see Chalmers, 2004) of sensor reception may therefore actually carry value to users, although this would need to be examined in field evaluations. This final point on seamfulness in proximity-sensitive systems is particularly important in that it offers the potential to both form and aid users' understanding about the nature of the proximal connection with the world, and allows its users (as well as designers the of future services on such systems) to be creative with the ways that personal communications and other services are provided by making use of these seams in connectivity to support interpretation. This is stark contrast to systems like Place Lab, in which all of the sensors in the system are fused into a single notion of location and this does not allow users to make any particular interpretations of *how* those places might have been understood or selected by the sender, to convey their message.

Returning to the broader theme of the paper, we have developed three key concepts relating to the use of proximity on developing context-sensitive services and messaging and these can be seen in the features developed as a part of the Marked Space system described. The first of these concepts lies in highlighting the importance and unique nature of proximity as a relational property in mediating the relationship between initial communicative intention and information retrieval. It follows that proximity-sensitive technologies are a distinct and nuanced form of locating, and offer different communicative affordances to those presented by location-based technologies, with different design considerations that need to be accounted for. Second, by giving users a differentiated view of proximity offers those users more meaningful communicative functionality that operates at a more human scale of distance: 'nearby' has a very different notion of scale depending on users' physical contexts (in the countryside versus in the office), their activities (walking around versus reading a book) and

their intentions (warning someone about something dangerous versus informing someone about something that might be interesting). To set this at a single arbitrary distance, as has been done in almost all of the existing location and proximity-based systems developed so far, is to artificially and arbitrarily cut across these levels of scale, and makes the system useable only under a particular set of conditions, thereby reducing the set of circumstances under which the system will be useful. There is a secondary effect to this following a reduction in the possible circumstances of system use in its potentially falling below the critical mass of use necessary to maintain interest in the system, and its falling into disuse. The third key feature is that different mechanisms of detecting proximity offer different resources and constraints in their operation and thus offer different communicative affordances to users. The particular communicative affordances of the sensors offer system users a variety of mechanisms to complement the mediated proximal relationship between the generator of information and its recipient. The use of different sensors also allows us to utilise the existing infrastructures in the environment in a way that is more richly pervasive than relying on a single form of sensory input. Together these three features of the Marked Space system described offer what are relatively simple techniques in terms of the technology being utilised, but which offer a unique and powerful mechanism for communication and the delivery of services to users.

We envisage proximity-sensitive systems and services to have application across areas as diverse as *in situ* access to product information, targeted mobile advertising, mobile guides, context-sensitive instructions and user manuals of various kinds, support for product inspections and servicing, mobile educational tools utilising situated and exploratory learning, mobile games and interactive entertainment, and awareness-based collaborative systems amongst others. The three concepts identified above apply are likely to apply to all of these areas to a greater or lesser extent depending on the form of the application. Once proximity-sensitive software is present in mobile devices, many of these applications can be relatively easily developed on top of these with relatively little effort beyond the provision of novel content and the services themselves, allowing developers to concentrate on their own offerings and without having to roll out and introduce the technology from scratch. At an interactional level, these techniques also offer techniques for user interface design and content provision based on physical proximity. By utilising the natural zooming effect of users approaching and moving away from sensor-enabled interactive devices, information can be presented in different forms or at different levels of detail, for example, situated displays (cf. O'Hara *et al.*, 2003) that change their font size or display interactive features as users approach them. This paper presents a first attempt to explore these issues systematically (within the context of asynchronous proximity-sensitive messaging) and, although this area is clearly still in its infancy, there seems great potential to employing the approach across a broad range of application areas that have so far received little attention.

In conclusion, we have discussed some of the implications for proximity-sensitive messaging systems and services and the notion of differentiated proximity sensitivity, but this field is ripe for further investigation and offers enormous opportunities for novel forms of data access and communication services.

References

1. Kindberg, T., Barton, J., Morgan, J., Becker, G., Caswell, D., Debaty, P., Gopal, G., Frid, M., Krishnan, V., Morris, H., Schettino, J., & Serra, B. People Places, Things: Web Presence for the real world. Tech Report HPL-2000-16, Hewlett-Packard Laboratories, 2000.

2. Eagle, N. and Alex Pentland, A. (2004) Social Serendipity: Proximity Sensing and Cueing. MIT Media Laboratory Technical Note 580, May 2004.
 3. Schmandt, C., Marmasse, N., Marti, S., Sawhney, N. and Wheeler, S. (2000) Everywhere messaging. IBM Systems Journal, 39, (3-4), 660-677.
 4. Espinoza, F., Persson, P., Sandin, A., Nyström, H., Cacciatore, E. and Bylund, M. (2001) GeoNotes: Social and navigational aspects of location-based information systems. In Proc. Ubicomp 2001, Eds. Abowd, et al., Atlanta, Georgia. Berlin: Springer, 2-17.
 5. Multaharju, M., Koskinen, K., Spacil, R., Ikonen, J. and Porras, J. (2004) Floating note - a location based messaging application. 2nd Workshop on Applications of Wireless Communications, p. 9-16, 5th August 2004, Lappeenranta, Finland.
 6. Burrell, J. & Gay, G. (2002) Egraffitti: evaluating real-world use of a context aware system. Interacting with Computers, 14, 301-312.
 7. Lane, G. (2003) Urban Tapestries: Wireless networking, public authoring and social knowledge. Personal and Ubiquitous Computing, 7, (3-4), 169-175.
 8. Yellowarrow.net (2006) Yellow Arrow. Last viewed on 9th Jan 2007. Available at <<http://yellowarrow.net/index2.php>>
 9. Priyantha, N.B., Chakraborty, A., Balakrishnan, H. (2000) The Cricket Location-Support system. In Proceedings of 6th ACM MOBICOM, Boston, MA, August 2000.
 10. Treu, G. and Küpper, A. (2005) Efficient Proximity Detection for Location-based Services. In Proc. 2nd Workshop on Positioning, Navigation and Communication (WPNC'05), Shaker Publishing, Hannover, Germany, March 2005.
 11. Chalmers, M. and Galani, A. (2004) Seamful Interweaving: Heterogeneity in the Theory and Design of Interactive Systems. In Proc. ACM DIS 2004, Cambridge, MA, USA, August 1-4, 243-252.
 12. Perry, M., O'Hara, K., Sellen, A. Harper, R. and Brown, B.A.T. (2001) Dealing with mobility: understanding access anytime, anywhere. ACM Transactions on Computer-Human Interaction, 8 (4), 323-347.
 13. Nokia (2007) Nokia Sensor. Available at: <<http://www.nokia.com/sensor>>. Last viewed on 18/07/2007.
 14. Bluetooth SIG (2006) Bluetooth® wireless technology surpasses one billion devices. Available at: <http://www.bluetooth.com/Bluetooth/Press/SIG/bluetooth_wireless_technology_surpasses_one_billion_devices.htm>. Last viewed on 19/07/2007.
 15. Ludford, P.J., Frankowski, D., Reily, K., Wilms, K. and Terveen, L. (2007) Because I carry my cell phone anyway: functional location-based reminder applications.
 16. Randell, C., MacColl, I., Muller, H. and Rogers, Y. (2002) Exploring the Potential of Ultrasonic Position Measurement as a Research Tool. In *IEE Communication Networks and Services Professional Network, editors, First European Workshop on Location Based Services*, pages 171-4, Savoy Place, London WC2R 0BL, September 2002. Institution of Electrical Engineers.
- Madhavapeddy, A., Scott, D., Sharp, R. and Upton, E. (2004) Using Camera-Phones to Enhance Human-Computer Interaction. *Adjunct Proceedings of the Sixth International Conference on Ubiquitous Computing*

(UBICOMP 2004).

Madhavapeddy, A., Scott, D. and Sharp, R. (2003) Context-Aware Computing with Sound. *Proceedings of the Fifth International Conference on Ubiquitous Computing (UBICOMP 2003)*.

Björk, S., Falk, J., Hansson, R. and Ljungstrand, P. (2001) Pirates! - Using the Physical World as a Game Board. Proceedings of the 8th IFIP TC13 International Conference on Human-Computer Interaction (INTERACT). Tokyo, Japan, July 9-13.

Falk, J., Ljungstrand, P., Björk, S. and Hansson, R. (2001) Pirates: Proximity-Triggered Interaction in a Multi-Player Game. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2001). 119-120.

Redström, J., Dahlberg, P., Ljungstrand, P. and Holmquist, L. E. (1999). Designing for Local Interaction. In: Nixon, P., Lacey, G. & Dobson, S. (Eds.) *Managing Interactions in Smart Environments: 1st International Workshop on Managing Interactions in Smart Environments (MANSE '99)*, Dublin, December 1999, pp. 227-238. Springer-Verlag.

Bell, M., Chalmers, M., Barkhuus, L., Hall, M., Sherwood, S., Tennent, P., Brown, B., Rowland, D. and Benford, S. (2006) Interweaving Mobile Games with Everyday Life. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2006), pp. 417-426, ACM Press.

Holmquist, L.E., Falk J. and Wigström, J. (1999) Supporting group collaboration with inter-personal awareness devices. *Journal of Personal Technologies*, 3(1-2), 13-21.

Iwatani, Y. (1998) Love: Japanese Style. In Wired News, 11 June 1998. Available at <http://www.wired.com/culture/lifestyle/news/1998/06/12899>. Last viewed on: 25/07/2007.

Consolvo, S. Smith, I.E., Matthews, T. LaMarca, A. Tabert, J. and Powledge, P. (2005) Location disclosure to social relations: why, when, & what people want to share. In Proceedings of ACM CHI 2005, April 2-7, 2005, Portland, Oregon, USA, pp. 81-90.

Hong, J.I., Boriello, G., Landay, J.A., McDonald, D.W., Schilit, B.N., Tygar, J. D. (2003) Privacy and security in the location-enhanced world wide web. In Proceedings of Fifth International Conference on Ubiquitous Computing: UbiComp 2003 (Workshop on UbiComp Communities: Privacy as Boundary Negotiation), Seattle, WA.

O'Hara, K., Perry, M., Churchill, E. and Russell, D. (Eds.) (2003) Public and situated displays: social and interactional aspects of shared display technologies (edited collection), Dordrecht, the Netherlands: Kluwer International CSCW series.