A Portal for Interacting with Context-aware Ubiquitous Systems

Don Cruickshank
School of Electronics and Computer Science
University of Southampton, Southampton UK
+44 23 8059 4490
dgc@ecs.soton.ac.uk

David De Roure
School of Electronics and Computer Science
University of Southampton, Southampton UK
+44 23 8059 2418
dder@ecs.soton.ac.uk

ABSTRACT
In a project to collect medical data from devices in wearables, the contextual information is essential for use by medics in interpreting the data. We have developed a portal which provides interactive access to the state and context of the devices and to the information collected. The portal is based on a MUD, which provides an appropriate spatial metaphor to reflect the context in the physical world.

Keywords
Guides, instructions, author's kit, conference publications

INTRODUCTION

In this paper we describe the portal we have developed for a project involving the collection of medical data from devices carried by patients after leaving hospital. The portal gives access to 'live' information about the state and context of the devices as well as access to the information being collected. The portal provides an interactive "digital world" interface to the physical world containing the devices.

For practical purposes we have developed the prototype portal using a variant on a 'multi-user dungeon' (MUD) [1]. The reason for the comfortable fit with a MUD relates to the natural mapping of context from the physical world into that particular digital world. The implementation of this mapping is the focus of the main part of this paper.

The next section describes the project. We then go on to discuss the requirements of the portal and its realisation in a MUD. This is followed by a discussion about context mapping and representation.

MEDICAL DEVICES

Hospital patients in intensive care are heavily monitored. Their context is very well known – it is their hospital bed, and they are probably lying in it. If monitoring can be maintained when they become more mobile then it may be better for the patient's recovery and, for the hospital, it might not be necessary to occupy hospital beds for so long at high cost.

Portable devices for monitoring physiological signals are available. However, these signals are not meaningful unless their context is known. For example, the medic needs to know if the patient is in bed, sitting, walking, running, at home, at work or travelling. This contextual information can be provided by additional devices, such as accelerometers and location tracking.

In the *Grid Based Medical Devices for Everyday Health*, the physiological signals and context information are relayed back to a grid-based system which performs the appropriate signal processing and makes the data available to medics, identifying episodes which need priority attention. The information is fed back to medics via portable devices.

The devices and sensors that we are dealing with typically have limited computational power and storage, they may be turned off to conserve battery power, and they only have intermittent network connectivity.

One of the objectives of the project is to address an infrastructure research question, i.e. to what extent can the Grid services paradigm adopted in the 'back end' be applied in the direction of the devices. The project has been conducted using Globus Toolkit 3, which provides a service-oriented model. Although in some cases the devices may be capable of hosting Grid Services, generally they interface via a Grid service proxy. The information – data and context – is then available for processing and storage.

Portal requirements

Our focus here is the means by which medics access the information. This is achieved through a portal, which is a point of interactive access to the information from and about the deployed devices. In turn, access to the portal may itself be from devices carried by the medics. Although the initial focus was on information access by the medics, it is apparent that the portal also has an important role in providing information to patients and also to those responsible for configuring and maintaining the devices (the 'pervasive support desk').

As well as access to the medical data, the portal has other ambitions within the pervasive computing research agenda. The first is that it should be easy to design and build the 'experience' for the users, and this should be achievable without needing a team of highly qualified computer scientists. Secondly the portal should be applicable to a broad range of projects. Finally, it should be a vehicle for addressing research issues in "record and reuse", i.e. it should facilitate experiments in replaying the captured data in novel ways. This stems from a need to consider both anticipated and unanticipated reuse of the data.

Our portal design has been driven by two scenarios. The first is natural language interaction, such as a clinician using a phone:

> ***"look at Tom"***
> *"Tom's patient ID is 12345. He is wearing a jacket. There is archived data."*
> ***"Look at jacket on Tom"***
> *"His blood pressure and heart rate are monitored"*
> ***"Look inside archive"***

The second involves interaction using instant messaging from a mobile device such as a cellphone or PDA, supporting a similar style (but perhaps more terse) interaction. This has been our test scenario for the prototype system.

One can envisage other interfaces, such as Web access, an RSS feed or even a Blog. In general there needs to be support for integration with other collaborative tools through use of open standards. There also needs to be access for non-human agents.

Implicit in these scenarios is the requirement for the portal to be both remotely accessible and supporting multiple users. Although we are not focusing on it here, the authentication and authorisation issues are also paramount, especially given the nature of the data.

The portal also needs to accommodate a varying number and diversity of devices. For research purposes it was desirable to be able to simulate additional devices, since deployments are inevitably limited in early trials.

Choice of portal technology

Given these requirements we were faced with a number of options. Clearly we could build a system from scratch. We sketched such a design, considering both a tuplespace-approach and an approach based on an RDF triplestore. We concluded that this was a viable approach. We also considered some of the portal technologies that are in use in grid projects and learning management systems, and the new JSR168 portlet protocol – the standards being adopted in these areas could provide interoperability with the portal.

We had a positive experience with a MUD in a previous project [3] [4] and knew that this already met some of the conceptual requirements including the room metaphor, a mechanism for representing environment, a notion of roles and multiple kinds of users. In support of research we wished to conduct, a MUD provides a basis for aspects of record and reuse, and a mechanism for proximity-based eventing. A MUD also brings a number of very practical benefits:

- Open source
- Scriptable (in the LPC language, a small object oriented C-like language)
- Rich support of network protocols (ftp, http, pop3, …)
- Standard clients available
- Support a variety of intermud protocols

We are also familiar with MUD deployment through the Access Grid, a room-based videoconferencing facility in which multimedia streams are routed to a video wall and speakers in the meeting room space; the operators of access grid 'nodes' communicate through a MUD, and the underlying methods of addressing the multimedia streams are hidden. The operators meet in agreed locations within the Virtual Venues, and the Access Grid software automatically locates the correct multimedia streams depending on their location – the metaphor is that the streams belong to the locations within the 'Virtual Venues'.

However, we were concerned that MUDs traditionally have not used a precise spatial representation of all objects against a coordinate system, that they provide no special security features and no grid integration.

We decided to adopt a MUD in our prototype for investigative purposes, with a view to using this experience to inform later decisions to build from scratch or to enhance a MUD.

The MUD

Figure 1 shows a map of the physical space which we represented in the Portal.

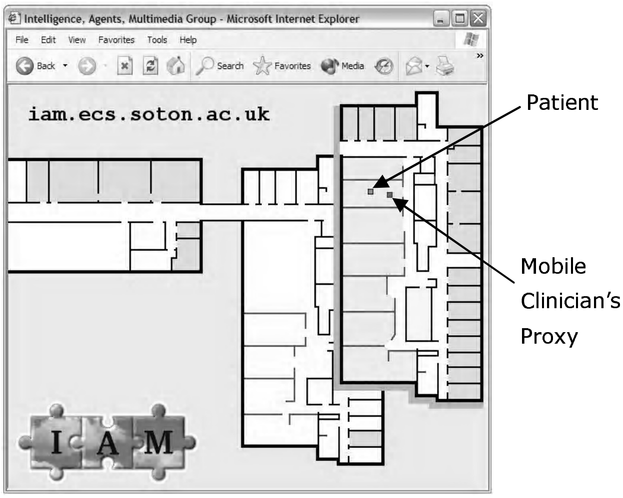

Figure 1: floorplan

Location

Objects within the MUD are positioned according to the stream of location information being received. MUDs are typically made with locations that describe the area they represent. Modern MUD technology (such as MudOS, see http://www.mudos.org/) is capable of accurately representing real space, where each location can have custom dimensions, and the relative 3D displacement between locations can be specified on each possible transition. It is trivial for MUDs to build their own internal map structure dynamically, and thus provide their own translation from an absolute coordinate space to the location system.

In the Portal, each location is able to define its own extent in the environment by use of a containment function. We use the support functions from the 'Discworld' MUD for axis aligned boxes to describe area, and these areas are now interpreted as axis aligned bounding boxes. A new function is used to test particular points for inclusion in a particular location.

In describing the physical space for the trial, we chose a polygon based solution to describe the areas used for locations. The polygons are z aligned, and a simple point-in-polygon function is used to test if a point is inside or outside the location. In conjunction with the bounding boxes (to implement floor and ceiling), this provided a mechanism to describe areas of office space. The two buildings described are positioned at an angle with each other, and neither building is aligned with North, thus relying on axis aligned boxes will not suffice.

As well as position, orientation is very important contextual information. MUDs use compass directions for location exits and the most common MUD libraries have support for locations of variable size (e.g. 10m x 8m) and exits have a default size depending on their type, and they are usually configurable on a per location basis. Orientation is also present in most libraries, and it is typically derived from where the user last came from. One common use for this information is for players who enter a location to get an idea about the direction where other users came from. Some MUDs also have the ability to allow users to "position" themselves within a room, e.g. "Fred is sitting on the sofa", or "Bill is standing next to the pillar".

An alternative method to location within MUDs is to dispense with strict containment rules, and allow a more freeform method of object visibility and manipulability. Every time a person or object is given a new coordinate, then the room for that object is generated dynamically containing the description of what they can see from that point in space. Other people and objects close by have a representation in the room that can be manipulated similarly to MUD environments with strict containment rules.

The different forms of physical location technology have consequences on the accuracy of object locations. GPS for example produces relatively low resolution information compared to ultrasound pingers. It is often not feasible to deploy ultrasound pingers on a large scale such as an entire city, however inside a building GPS will not work at all. Our system therefore has to take into account location data from a number of different information sources. Not all location technologies provide coordinate information. Inside a building, we may not have any radio or ultrasound detection and rely on RFID or smartcards. If the most recent information is that a person has walked through a doorway, we know that the person is in a particular room. Any previous information is useful as a trail, but does not assist us in locating the current position of the person.

In the case of a person walking through a doorway, we might not be able to deduce the precise coordinates of the person, but in the least we know which room they are in. For algorithms that work solely on coordinates, we might choose to use the room origin as the coordinate for that person. If the rule simply relates to membership of the room, then coordinate approximation is neither made nor used.

Many of the notable characteristics of an object in the MUD are dependent on the type and purpose of the object's location. In the Ambient Wood, the trail of interactions of each object is recorded, and a journal [5] in the form of a hypertext is produced. Location within the journal is depicted not by an absolute form such as a GPS coordinate, but in a symbolic form, e.g. "The large clearing." MUDs can also use relative positions for objects. It is often not necessary to locate an object with another in absolute coordinate space, but to denote the relationship using the prepositions such as "on", "in", "at" or "beside". For example, a person can sit *in* an armchair, or sit *at* a table.

Mud descriptions

In the MUD, objects are typically described in natural language. It is possible to access varying levels of descriptions depending on how focused the observer is on that particular object. For example an object may be briefly described as "a medical jacket". A user might decide to look at the jacket in further detail and see a longer description of the jacket, and that is has sensors in it. The user can then choose to look at the sensors in the jacket, etc.

The brief descriptions are often compounded into sentences such as "Two medical jackets and a medicine dispenser". To assist the narrative flow of events that occur to MUD objects, the MUD programmer has the choice to choose whether objects are described by "a/an", "the" or "one of the" when referring to objects. The last form conveys extra information about the number of objects that match the description. If the MUD tells us that "Don configures one of the medical jackets", then we can infer that there must be at least two jackets in our location.

A further step in increasing the quality of narrative flow is to compound similar statements from the MUD. The common purpose for this is to convey how many objects are being manipulated and how many others there are in the same location. If Don configures a second wearable shortly after configuring the first one, the MUD might notice that similar messages were the result of subsequent verbs and compound the messages to produce "Don configures two wearables". If there are exactly two wearables in the location, this will be evident from the message "Don configures both of the wearables". This is an example of how contextual information can be provided succinctly within a crude natural language interface.

In addition to offering multiple levels of descriptions and for those to be compounded automatically, each MUD object can offer a tailored representation that is dependent on the observer. The tailored representation includes the descriptions of the object and the nouns and adjectives that can be used to refer to the object. For example, a medic in the room will be able to recognise medical symptoms and communicate with appropriate terminology, while the patient receives a patient-friendly version.

Figure 2 illustrates an interaction using instant messaging on a phone.

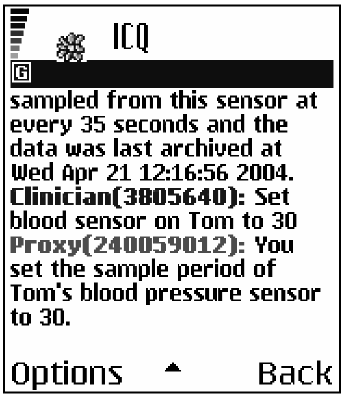

Figure 2: interface on phone

Persistent state

Objects have state within the MUD. Objects that can be carried by a user (i.e. in the user's "inventory") have a serialised form. Due to the large number of objects that a MUD can handle, the serialised form might not contain all of the variables that the live objects contain. For example, the maximum number of devices that a wearable can contain is not present in the serialised form because they are constant for all objects of that type. Player objects in the Discworld MUD have approximately 250 separate variables, and some of those are complex datatypes, e.g. arrays and hashmaps.

MUDs deal with serialised objects all the time. Users that log out and log back in will usually have their inventory recreated. Also objects that offer permanent storage will typically use the serialised form to keep the objects in a save file.

Context

Context is represented at various levels within the MUD. Location is a strong contextual concept. Objects can exist within the inventory of other objects, e.g. a torch inside a backpack, inside a player's inventory that is inside a location.

The parser maintains the context of the objects that you have recently referred to. For example, words such as "him", "her", "it" and "them" will refer to the last object or objects that were referred to that match those words. These are kept separately and "it" and "them" are used depending on how many objects are being referred to.

Another concept is the idea of an object "trail". A history of locations is kept within each object, including the timestamp and the cause of each change of location. This information is sometimes kept within gaming MUDs for administrative purposes, but it can be useful for more generic applications – e.g. to retain an audit trail. The length and content of the trail can be dependent on how long ago the events occurred and also the relative importance of the individual trail data items.

Manipulating real world artefacts

In addition to manipulating entirely virtual entities, MUDs can also be used to locate real world artefacts. The strong link with virtual location and navigation can be used to interact with objects in the real world. In the medical devices scenario, we use the MUD to configure sensors on the medical jacket. Again, we chose to not address the individual devices and sensors using Grid handles, but using expressive locators like "blood sensor on Tom". The context of the clinician's phone is used to identify which Tom we are referring to. In the case where there are several people that share the name "Tom", we can configure the MUD to force the user to be more specific if that case should arise.

Event processing

The event stream from an object is broadcast to all objects in the same environment, and also the environment object itself. Event data can range from simple notifications to communication streams. Communication streams are not limited to text. Audio and video streams are also possible.

In the Portal, we aim to not only perform session control of live information streams (*cf* multimedia streams), but also to carry the multimedia packets and filter them using the event channels within the MUD.

An existing mechanism of the Discworld MUD, known as the Broadcaster, can be used to broadcast events more widely than a single environment, but within a certain radius of a location within the MUD.

DISCUSSION

In [2] six forms of contextual knowledge needed to provide personalized services are given:

1. The role of the person in the meeting
2. The type of services that the person has access to
3. The type of the devices that the person carries
4. The type of non-computing objects the person's vicinity (e.g. the type of clothes the person wears and the type of objects that the person holds)
5. The time at which the person enters the room or joins the meeting
6. The identity of people to whom the person is talking

We observe that all of these are represented explicitly in the MUD.

However, the power of the MUD-based portal is not simply that it captures context-sensitive information but rather that the *stateful interaction* with the portal matches the stateful interaction of users with corresponding devices in the physical world. The portal is a natural link between the physical and the digital world.

This study has also illustrated the power of the natural language interface in conveying contextual information, some of it implicitly. This is effective even with the very terse interface that we have tried.

We are currently populating a location ontology to describe the space referred to in this paper, as a step towards the application of knowledge technologies to the portal. By modelling the physical properties of an environment using the ontology and loading this model into a MUD, we will have a rigorously described model with a location based event and behaviour system. With a range of sensory inputs, the MUD can position people and devices within the environment, and behaviours can be applied depending on the context to each object.

We are in the process of applying the portal to a sensor network project. We are also interested in collaborative aspects, whereby a community of users interact via the portal, which were not paramount with the medical devices. The application as part of our "record and reuse" research is also a subject of active discussion.

CONCLUSION

We have presented a system which enables a user to interact with virtual devices which correspond directly to physical devices in the real world. The relevant context is preserved as completely as possible, and we have discussed how we have achieved this. Significantly, not only the contextual information but also the stateful interactions are faithful between the digital 'model' and the physical reality.

The prototype has been based successfully on a MUD and we are currently looking at the incorporation of a location ontology within this. The portal is being explored across various projects.

ACKNOWLEDGEMENTS

This work was supported in part by the UK Engineering and Physical Sciences Research Council under grants 'Grid based Medical Devices for Everyday Health' (GR/R85877/01) and the Equator IRC (GR/N15986/01).

REFERENCES

1. Bartle, R. Interactive Multi-User Computer Games, BT Martlesham Research Laboratories, 1990.
2. Chen, H. and Finin, T. An Ontology for Context Aware Pervasive Computing Environments.
3. Rogers, Y., Price, S., Fitzpatrick, G., Fleck, R., Harris, E., Smith, H., Randell, C., Muller, H., O'Malley, C., Stanton, D., Thompson, M. and Weal, M.J., Ambient Wood: Designing New Forms of Digital Augmentation for Learning Outdoors. in *Third International Conference for Interaction Design and Children (IDC 2004)*, (Maryland, USA, 2004), 1-9.
4. Thompson, M.K., Weal, M.J., Michaelides, D.T., Cruickshank, D.G. and De Roure, D.C. MUD Slinging: Virtual Orchestration of Physical Interactions, ECS, University of Southampton, 2003.
5. Weal, M.J., Michaelides, D.T., Thompson, M.K. and De Roure, D.C., The Ambient Wood Journals - Replaying the Experience. in The fourteenth conference on Hypertext and Hypermedia (HT03), (Nottingham, UK, 2003).