MUD Slinging: Virtual Orchestration of Physical Interactions

Mark K. Thompson, Mark J. Weal, Danius T. Michaelides,
Don G. Cruickshank, David C. De Roure

Intelligence, Agents, Multimedia,
University of Southampton, UK
{mkt,mjw,dtm,dgc,dder}@ecs.soton.ac.uk

Abstract

In this paper we look at two scenarios of physical-digital interaction. The first is an Ambient Wood, forming a sub project of the Equator IRC researching playful learning through novel interaction. The second, a meeting room, constructed for the FEEL project in which the management of intrusive notifications is desired. We look at how these scenarios can be modelled using a Multi-User Dungeon (MUD), providing orchestration tools for the interactions taking place, with consideration to other mechanisms available.

1 Introduction

In various projects within our research group, we are discovering a need to provide infrastructure support for the co-ordination, arbitration and orchestration of resources within an environment. For example, the automatic configuration of the audio and light levels in a meeting room for a particular group's use, or the selective and timely notification of lab members by various appropriate means of events such as scheduled meetings, important emails or opportunistic coffee-breaks.

The nature of our multi-disciplinary and multi-project research lab means that a wide spectrum of interfaces exist to the various resources that we might wish to control, as development of the various components has been driven by different research goals. 'Virtual' representations of physical resources exist in numerous different forms, and an on-going research activity within the lab is exploring distributed, event-based infrastructures that enable the orchestration of these resources.

In this paper, we examine two different scenarios in which people interact with physical artifacts and virtual information. One is from the EPSRC Equator sub-project Ambient Wood[1], a mixed reality space constructed in a Sussex woodland. The other from the IST FEEL project[2], investigating the management of intrusive services in meeting room environments.

[1] http://www.cogs.susx.ac.uk/projects/equator/equator.htm#Ambient

[2] http://www.feelproject.org/

Figure 1: Viewing images in the wood.

Using these two example scenarios, we illustrate our adopted approach for orchestrating the interactions between participants using MUD software, more commonly used for gaming or as a teaching abstraction [5].

1.1 AMBIENT WOOD

In the Ambient Wood, a number of activity spaces have been defined, each offering its own aims with a focus on the different kinds of technologies and activities. The experiments within the wood are targeted at the Keystage 2 Science curriculum, focusing on hypothesis testing. The infrastructure developed is designed to be flexible in that it can be re-used for other parts of the curriculum.

Pairs of school children are instructed in the use of various devices and tasked to collaboratively discover information about plants and animals living in the different habitats in the wood, some of which have been 'invisibly' augmented with technology. Their experiences are later reflected upon in a 'den' area where the children share their findings with each other. They then construct hypotheses about what would happen to the wood under various conditions over a period of time, for example lack of water or the introduction of a new species.

As the title suggests, a key requirement of the Ambient Wood scenario is for the technology to be as invisible as possible. Whilst, inevitably, the devices with which the children collect data about their surroundings form part of their attention focus, the other technological devices are camouflaged where possible.

A simplified usage scenario illustrating the types of physical-digital interaction that we require to orchestrate could be:

> Team Squirrels walks into the small clearing and uses a probe device to take a reading of the level of moisture in the ground. Their hand-held device shows that the ground is very dry. As they walk toward the large tree, they hear the sound of a badger in the undergrowth. An image of badger tracks appears on their hand-held device (See Figure 1) and a voice-over narrates a short description of the typical habitat requirements for that species.

A more detailed description of the Ambient Wood project and an exploration of the issues motivating the research can be found in [12].

Figure 2: The augmented Meeting Room.

1.2 FEEL

The primary objective of the FEEL project is to explore issues surrounding the intrusiveness of today's mobile technology, and how work in local environments can be enhanced by introducing the idea of non-intrusive services.

Disruptions to collaborative tasks, for example meetings, often occur due to external stimuli such as telephones ringing, e-mail arriving or people physically interrupting a conversation. Within the FEEL project, an intrusion is defined as *"an occurrence of a process or task that is not intimately related to the current task of a group and that interferes with the realisation of that task"*.

The role of technology within FEEL is to manage the way in that interruptions intrude on activities. This includes the development of different mechanisms for capturing the information being conveyed by an intrusion, assessing its importance and then selecting an appropriate means of delivering it. The project exercises the notion that different participants may each assert a different level of acceptable intrusion.

An example scenario would be:

> A group of researchers sit in a meeting (see Figure 2) to report on their project progress for the month. One of the participants is expecting an important telephone call relevant to the meeting, but the group otherwise wish to remain undisturbed by people not in the room and so set their 'intrusive knobs' to reflect their desire. All other telephone calls are then automatically deflected to the Group secretary.
>
> Having scheduled the meeting in advance using the room calendar system, the audio-visual equipment is already configured for the meeting, with the agenda and only relevant progress reports available on the projector for presentation and annotation.

This scenario extends beyond FEEL and into other research projects within our group, and as a result there is a diverse set of technologies deployed in our experimental meeting room.

2 Interactions

Whilst the above examples are by no means exhaustive scenarios within the two projects, they serve to illustrate some of the interactions that we wish to orchestrate. The following sections pull apart the scenarios to examine the specific nature of the interactions taking place at both device and information levels.

2.1 Ambient Wood

In the realisation of this scenario, children are walking round with hand-held devices, with 802.11b networking and a second small hand-held unit that serves as a 418MHz radio receiver.

Placed within the wood are small radio transmitters that periodically broadcast a unique identifier, a 'Pinger'. These provide physical contextual anchors for the hand-held's position, and therefore the position of the children, within a region of the wood.

There are other 'probe' devices that provide contextual information over the short-range radio link, including small moisture and light probes and an enclosed GPS position pinger, carried in backpacks worn by the children.

Also hidden within the environment are wireless speakers providing an interface through which audio cues and event sonifications can be made to occur as ambient sound.

In the Ambient Wood scenario outlined above, there are various physical interactions going on that require some level of orchestration. The devices in the scenario can be divided into four categories, Sensors, Tangibles, Interfaces, and Infrastructure.

- Sensors. *Pingers and probes, communicating to hand-held devices by radio, capturing physical context cues such as location, light-level, etc.*
- Tangibles. *Physical objects with RFID tags that the children can place at various points around the wood on disguised tag readers*
- Interfaces. *Jornada hand-held devices serving as a direct audio and visual interface to the children, but also acting as a proxy between the sensors and other infrastructure components*
- Infrastructure. *Speakers for sonification, and laptops on which the software infrastructure is executed*

Networking in the Wood is achieved through the deployment of Linux-based laptops acting as 802.11b base stations in Host AP mode, providing one flat IP-based network with coverage of the entire Wood area.

In addition to the interactions between physical devices, the children also interact with information in the form of triggered events. These can be modelled as virtual objects which the children interact with causing the information to be transfered to the appropriate device (hand-held, speaker etc.)

Figure 3: The Intrusive Knob.

2.2 FEEL

The FEEL scenario contains less physical context detections and 'invisible' interactions than the Ambient Wood scenario. It is more concerned with the automation of local meeting room artifacts and the controlled interactions between them. It also features close integration with other third-party information services, including calendaring software, document management databases and the Web.

The meeting room supports a variety of public and private visual displays and audio spaces; cameras and microphones for participatory input; and automated lighting and sound synthesis for notification. The purpose of the room is to experiment with the management of intrusiveness, initially focusing on audio and visual display device intrusion. The audio facilities are production quality and entirely under software control, coupled into a software framework. The video facilities are typically digital videoconference grade, software controlled and networked using PC hardware.

The 'intrusive knob', its current form shown in Figure 3, is a tangible interface through which users can express a desired level of acceptable intrusion, on a scale of *"I don't care whether anything interrupts us"* through to *"Under no circumstances would any intrusion be acceptable"*.

Controlling entry into the room can be achieved in a number of ways. For example, with an electric lock or electronic signage that indicates whether a particular person, whose identity is detected can enter or not.

The types of 'devices' present in the meeting room are:

- Sensors. *RFID tag readers, card-swipe readers and other sensors that detect context of the meeting, e.g. what papers are on the table, who is present in the room*

- Intrusion channels. *Things that enable intrusions to occur and therefore require controlling, e.g. mobile phones, door to the room, email and instant-messaging clients*

- Displays. *Things that intrusions might be displayed or conveyed through, e.g. projector, ambient light, personal audio speakers*

- Information systems. *Sources of information that assist in the assessment of whether an intrusion should be permitted or not, for example web-based calendars and community of practise databases*

- Infrastructure. *Audio/video equipment, including various types of microphone and speakers, enable us to define both public and private audio spaces, e.g. for video conferencing and notification delivery*

3 Mechanisms for Orchestration

Given the two scenarios, there are a number of candidate approaches that would be applicable to orchestrate the interaction between the physical devices and information structures. A representative selection of these is presented below.

3.1 Manual orchestration

The simplest form of automated orchestration in these scenarios is to not provide any. That is, manual manipulation of physical resources without any systems-level awareness of the state of the world. In both scenarios, this would imply the involvement of a human facilitator who observes the actions of the participants in the scenario, and manually updates device states, for example causing a sound to be played, or a telephone call to be deflected.

The benefits of this approach is its degree of flexibility. However, this flexibility is at a high cost in that there is no automation, requiring explicit human intervention.

3.2 Direct connection

One approach to scenario orchestration would be for each participant device to maintain its own mapping of events and interactions. This would have the benefit that the co-ordination takes place at the point of interaction. In a scenario where network connectivity is sporadic, devices would be able to continue to operate in an autonomous fashion when disconnected.

However, the maintenance and complexity of mappings and interactions could be difficult to manage. For each change in scenario, there is potential for every device to require re-coding. Each individual device would need to be manually configured and named, and if that name changes every other device that interacts with it would need to be modified individually, too. Also, the orchestration logic moves to the devices themselves, increasing the complexity of the software they will need to run.

Tracking what is happening in such a distributed setup is also problematic, with each device having to maintain its own event logs which would then need to be assembled in a coherent fashion after the fact before any global analysis could take place.

3.3 Purely event-based

One mechanism for decoupling the direct links between interactive participants in the scenarios would be to generate notifications (events) for interactions between devices or services. Processes register interest in events.

For example, a new participant enters the meeting room and their presence is detected by an RFID tag reader. A process monitoring the status of the tag reader generates a notification. A process maintaining a list of attendees reacts to the notification by updating its list.

In this approach, no global world state is maintained in any one place. Components subscribe to events about which they are interested and ignore everything else. It requires a well defined schema for notifications, enabling interoperability, extensibility and adaptability.

As with the directly connected approach above, there is a lack of global state and the problems that accompany that.

3.4 Dataspace co-ordination

A different approach is for each participant device in the scenario to maintain a virtual model of its own state within a shared dataspace, for example a tuple-space system such as TSpaces[16], or a spatially-aware shared dataspace system with an event model as described above, such as EQUIP[7].

Orchestration processes could then be written that subscribe to dataspace events, such as indication of new state or a new event, to observe the modelled state of the world. Processes on, or working on behalf of, the devices in the scenario would watch for changes to the model that would be of interest to them so as to physically realise the changes, for example play a sound or deflect a telephone call.

This is a very flexible arrangement, and is one that has also been adopted for another implementation of the FEEL scenario, and also for the EQUATOR sub-project City[9].

The virtual model of the physical world, maintained in the data space, can be analysed and, if needed, manually manipulated should devices fail or the scenario require adaptation on-the-fly. However, it does delegate the responsibility of data update and interaction orchestration out to the devices. In the Ambient Wood scenario, for example, individual processes would need to be written for each participant device, which, as in the case of the location pingers and probes, would need to be executed on the hand-held devices as proxy processes. As the scenario evolves, the codebase on the mobile devices would need modifying.

3.5 Hypermedia Linkbases

A similar approach to dataspace co-ordination, where, rather than employing software processes to parse and adapt the tuple-based model of the world, hypermedia linkbases are used that containing structured transitions (links) that define the different world states available from a particular context.

Orchestration could be carried out by identifying links between states and events. For example in the Ambient Wood scenario, a particular pinger being received by the children could be taken as a source anchor for a hypermedia query. Links could then be created from this source to different possible outcomes, for example a sound being played on a speaker, an image being displayed on the handheld device. The client application would query the linkbase and receive a number of links back. It would then choose one or more of these and carry out the required action, for example if the destination of the link is an image, the child's handheld would be told to display the requisite image.

The links would be manually authored at scenario-definition time and it would be responsibility of the various client applications to invoke the actions required by link traversal (e.g. invoking a sound on a remote device). The benefits of this approach are comparable to the above use of tuple-watching processes, with the added advantage that of additional structured layer above the tuple space using well understood techniques borrowed from the Hypermedia systems domain.

This is a simple form of orchestration that separates the logic into a linkbase, using an open hypermedia model. Simple interactions can be easily created, however more complicated interactions would require increasingly sophisticated clients to enact links traversal. Also, there is no state kept in the system so events that evolve over time would be harder to construct.

An alternative hypermedia approach to orchestration would be that of sculptural hypertext [3, 15]. Here, each link has a source that is a context item as opposed to an anchor on a node. The link service is queried with a context and any links with a matching source context are returned.

The act of following a link causes events to occur and also modification to the user's context, using the concept of link 'behaviours'. Sculptural hypertext allows a more state-based approach to be taken. The query asks 'What links are available as the system stands currently?'

3.6 Agent processes

For the purposes of this paper, we consider agents as software processes that model or 'take ownership' of particular aspects of the scenario exclusively, interacting with other agents as appropriate. Agents would dynamically discover each other and therefore be exposed to different facets of the scenario at run-time and interact as appropriate. This would require the definition of the discrete models encapsulated into agent processes and the definition of the interactions permitted on the model, incorporating a performative vocabulary and a means to express the semantics of that vocabulary to other agent processes.

Orchestration would be the composition of agent processes to manage interactions between them, and therefore the model and the scenario. In a trivial case, this would be the restriction of what performatives are permitted to be invoked between different types of agents.

For example, within the FEEL scenario an agent might offer a software interface to control and observe a MIDI slider that is associated to particular audio channel in the meeting room's mixer. The agent's local model maps the slider to the audio level of a remote participant in the video conference, modelled in the conference's agent process. The MIDI slider agent exports an interface through which other agents can request the status of the audio level, or perhaps to request that it is changed. An orchestration task performed would be to restrict which other processes can request audio level changes, for example, by a process of negotiation dependent on the state of other models in the scenario [10].

3.7 Virtual Worlds

The hypermedia approach models the different states of the system as nodes, with the links as events triggering changes in state. The Virtual Worlds approach takes this one stage further, with the model being a virtual re-creation of the physical environment. Interactions between devices in the physical environment have corresponding interactions *within* the model. At any given moment, the virtual world is an 'attempted replication' of the state of the real world. We use the phrase 'attempted replication' to emphasise the fact that the virtual world is dependent on sensed information regarding the physical world. For example, it can only model the real world participants' location based on the last time it received information about them.

This approach is a centralised one in that the virtual world has to be aware of and able to model all interactions; it is an all-encompassing model.

Virtual Worlds range from interactive text-based adventure games through to fully 3D virtual reality environments such as MASSIVE3 used to produce inhabited TV events[8].

4 MUD Orchestration

Of the various techniques illustrated above, we chose to adopt a variant of the Virtual World approach. In the next sections, we describe our chosen tool, the MUD, and how we have used this technology to realise the Ambient Wood and FEEL scenarios.

4.1 Introduction

The first MUD was created by Richard Bartle and Roy Trubshaw at the University of Essex in 1979–80 [2]. Since then different kinds of MUDs have evolved, with various implementations. The MUD Frequently Asked Questions document states the following[14]:

> A MUD (Multiple User Dimension, Multiple User Dungeon, or Multiple User Dialogue) is a computer program which users can log into and explore. Each user takes control of a computerised persona (or avatar, incarnation, character). You can walk around, chat with other characters, explore dangerous monster-infested areas, solve puzzles, and even create your very own rooms, descriptions and items.

Not all MUDs involve monster-infested arenas, nor the 'combat features' that characterise some. In MUD terminology, the style of MUD environment discussed here is a derivative of LPMud, named after the original author, Lars Pensjö. The MUD mechanics are programmed in an object-oriented C dialect called LPC.

Our use of a MUD is a different approach to other research-oriented uses, where the focus is typically on providing direct support for collaborative work, for example task co-ordination [6, 4]. Rather, we envisage the MUD as an invisible, pervasive orchestrator of the physical world, reacting to, and generating events as appropriate.

The key factors for our interest in deploying a MUD architecture for world modelling and orchestration were as follows.

- Metaphor of real life. *In a MUD, people and things exist in a place, and people interact with their environment as they would in real life*
- Programming model. *The object-oriented, event-based programming model enables us to attach behaviours to virtual representations of real-world objects, and model the state changes on a per-object basis*

- The facility to intervene. *The virtual model can be changed explicitly from within to affect events in the physical world*
- Extensible in real-time. *Not only the ability to create new locales and objects as the usage scenario requires, but also to adapt the functionality (the 'game logic') of existing interactions whilst the MUD is running, enabling dynamic and efficient modelling*
- (Centralised) Orchestration. *For the initial phase of experimentation at least, maintaining a central model of all of the state of the various sensors, tangibles and real-world objects facilitates scenario engineering*
- Decomposition after the fact. *Through maintaining extensive event logs, different parts of the scenario can be analysed, or even replayed, virtually*

4.2 Realising the Ambient Wood scenario

The Ambient Wood MUD is structured such that every interaction of interest that occurs in the physical world is modelled in the virtual world. The model includes all of the regions of interest within the wood, the participants and all of the physical devices, such as probes and speakers.

4.2.1 Interfacing between the real and virtual

To enable the event-driven nature of the scenario whilst minimising the processing load placed on the mobile devices, a lightweight subscription-based notification service is used, Elvin [1, 13]. By focusing on the delivery of messages based on content, as opposed to explicit addressing, the complexities of naming and locating different devices is avoided.

Central to the approach adopted within Ambient Wood is a process called MEAP (the MUD-Elvin Application Proxy). Elvin notifications form the primary interface between the real and the virtual from the MUD's perspective in that notifications are generated as event triggers based on physical interactions. These events are interpreted by and enacted by a Wizard character[3] within the MUD, updating the model and generating new events as a result. Think Laurence Olivier in 'Clash of the Titans'

Figure 4 shows some of the objects that were modelled in the MUD for the user trials at the Ambient Wood in September 2002.

4.2.2 Modelling the wood

The physical locations within the wood have to be modelled virtually in the MUD, becoming rooms. Location presence is determined by cues from sensor devices, such as the pingers placed at various points in the wood, and also from GPS data being collected. The real wood is shown in Figure 5 with the locations of the physical pinger devices marked.

Figure 6 shows the rooms of the MUD and the connections between them. Although arranged to roughly mimic the real wood in the diagram, the topology is the only thing of interest to the MUD. If the pingers were moved in the actual wood the topology of the MUD would remain intact, which is one of the benefits of using a MUD.

[3] Game characters able to modify the internal model of the MUD

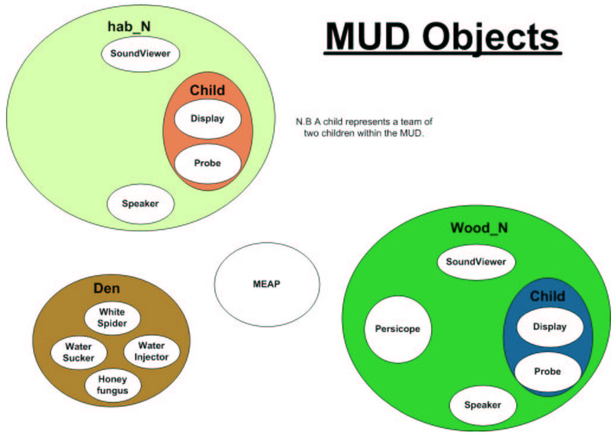

Figure 4: The Objects in the MUD.

When modelling the team's positions within the MUD, we included a fuzziness threshold so that, should a positional cue not be received within a certain time, the team's virtual representation would be teleported to a special 'limbo' room, indicating the lack of confidence in placing the team within the physical space. In practise, this facility was not used due to the excellent network coverage within the wood, and the frequency of the location cues' detection.

Also related to position is the mechanism by which actual locations of teams were determined. It was possible, in fact more often the case than not, that the MEAP proxy would receive many location cues from different sources (e.g. GPS pinger and fixed location, labelled pinger). In some cases, the mapping between sensed position according to the different cues and the real physical position would be in conflict.

For the September Ambient Wood trials, the GPS-based data was ignored as far as the modelling of team position was concerned within the MUD, due in part to inaccuracies arising due to drift and the lack of a differential-GPS receiver on-site. In some physical locations, where radio propagation and labelled pinger placement permitted, radio pings would either overlap, resulting in the loss of the message, or interleave with the receiver reporting two distinct locations for the same team. Without a degree of smoothing, the MUD model of the teams' positions would be erratic. A degree of sensor fusion, for example by assigning levels of confidence or accuracy weightings to the different sensing technologies, is required for future experiments with this model.

4.2.3 Participants

The Ambient Wood scenario contains a number of children that actively interact with their real environment. They do this by walking around the wood, probing the ground and taking light readings and placing physical artifacts in areas of the wood. To model these interactions the children have to be modelled in the MUD. MUDs allow two types of characters, Player Characters (PCs) and Non-Player Characters (NPCs).

The children are modelled as NPCs as they will not be interacting directly with the MUD, instead the MEAP Wizard will move the NPC representations

Figure 5: The real wood and device locations.

around the MUD based on location detection events. A child in the MUD is representative of a team of children in the wood (two). This simplification arises from the practicality of there being only one pinger receiver per team, so the MUD must assume that the two children are at the same location. When MEAP receives a message to say a particular pinger has been received, the NPC representing the team of children is moved to the appropriate room in the MUD.

Because MUDs perceive little difference between PCs and NPCS, we can test the scenario by entering the MUD as a PC, allowing us to move around, use the probes and experience the scenario first hand without the need to set up the physical devices in the actual wood. This provides a useful debugging tool allowing for thorough laboratory testing, minimising the setup time in the actual wood.

4.2.4 Virtual interactions

When the children change location in the MUD by coming into range of pingers in the real world, events can be triggered. These are modelled within the room in the MUD and used to allow the children to discover invisible artifacts. These artifacts can manifest themselves in a number of ways, perhaps a sound from a hidden speaker, or an image appearing on the display they are carrying.

The virtual object in the MUD, when discovered, can be observed in the real world. To achieve this, the location can trigger a sound-viewer object to play a specific sound through a speaker or to to display an image or sound on the teams' hand-held. A number of alternative ways of modelling this in the MUD could be used. We chose to model the sounds directly in the locations. An alternative which might provide greater flexibility would be to model the

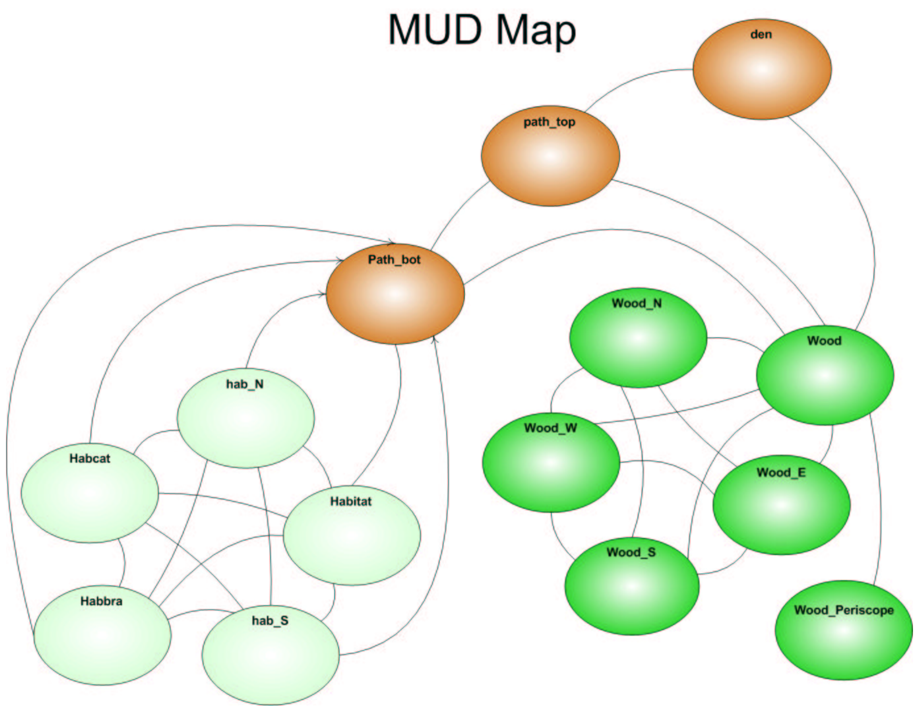

Figure 6: The room map for the Ambient Wood MUD.

events as objects that are placed at locations and triggered by the arrival in the location of an NPC.

4.2.5 Physical interactions and Tangibles

As part of the Ambient Wood experiment, a periscope device was constructed that allowed the children to view movies illustrating processes that were not possible to view directly. For example, how the wood might look in other seasons, what the wood might look like if different species were introduced etc. To simulate the introduction of new species in the wood, tangibles were used. In this case, tagged physical objects that the children could place on a disguised RFID tag reader. The placement of the objects on the reader triggered events that caused the representational object in the MUD to be moved to the periscope's modelled location.

The presence or absence of objects in the periscope room in the MUD caused different sounds to be played in the wood after the movies had been viewed. For practical reasons, the tag reader was not explicitly modelled in the MUD, rather MEAP responded to movie events from the periscope.

4.2.6 Record and replay

All the activity in the MUD was recorded in log files, as were all of the Elvin notifications. Further research aims to develop 'journals' that allow children to replay their activities in the wood in the form of an interactive multimedia presentation that will identify key events from the log, replay the sounds and movies that the children experienced, allowing them to relive their field trip or recount their activities to classmates.

4.3 Realising the FEEL scenario

The FEEL MUD takes a slightly different approach to that developed for the Ambient Wood scenario, but still includes representations for the participants, their devices and the other infrastructure in the meeting room.

4.3.1 Interfacing between physical and virtual

The people in the meeting room are modelled in a similar manner to the children in the Ambient Wood MUD, as NPCs, as their interactions with other artifacts within the MUD model's interest is indirectly observed.

Presence detection of participants is achieved through the use of RFID tags placed on coffee cups, log books and laptops, with the tag reader unobtrusively mounted on the underside of the meeting table. There is no problem therefore with more than one representation for a participant being detected (e.g. their personal coffee mug and their log-book both on the table), however, problems arise when people lend their mug or their laptop to someone else to take to a meeting.

RFID tags are also used in this scenario to tag paper reports and documentation tabled at the meeting, providing additional contextual information as to the current activity enabling automatic adaptation of intrusion levels. For example, the presence of the participants in the room but with no documents on the table might indicate that the meeting has yet to begin or is about to end. This could permit a different level of intrusion compared to when papers and logbooks are out on the table.

Environment controls, such as the volume levels of the room's audio sources and their routing to different speakers, are controlled in software through the adoption of MIDI[4] as a control interface.

The control of various visual outputs, for example interactive signage and data projectors, is achieved through various means including remote desktop control (VNC)[11] and purpose-built control software.

In the example FEEL scenario there is no notion of attention focus insofar as a requirement to model which participants are interacting with each other, unlike with other CVE systems. However, given an appropriate sensor technology, the model can be trivially extended to incorporate this data. This extension would be useful, for example, to make assessments about the intrusiveness of a pending notification dependent not only on collaborative task (as with the current scenario where the entire room is considered as engaged in one task), but also for users to say that *"Its okay to intrude whilst X is talking"*.

4.3.2 Interfacing between virtual and virtual

The FEEL scenario is much more dependent on external information sources than Ambient Wood. There are two axes of virtual 'device' that are participants in this scenario.

[4] *Musical Instrument Digital Interface* - see `http://www.midi.org/`

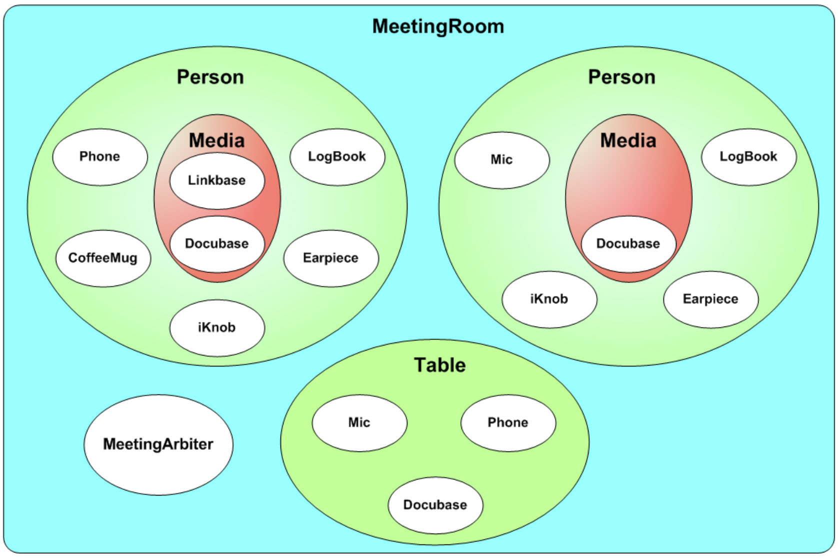

Figure 7: An example MeetingRoom container in the MUD.

Firstly, information systems such as the calendar system and the research group's project and publication databases are virtual participants that have a role to play in the orchestration of notifications and other interactions within the room.

In many respects, the modelling of such information systems within the MUD as 'information objects' is similar to the use of a software agent approach as discussed above. However, the nature of the MUD model means that pieces of information can be modelled as belonging to a particular participant using the familiar metaphor of containment: Dave's documents are information objects inside the NPC container that represents Dave's presence in the room, such as shown in Figure 7 which illustrates some of the MUD objects and their containment relationship within one meeting room.

This naturally extends into the modelling of people and information that are presently outside of the room, but that may still have interactions that we wish to model, c.f. the 'limbo' room in Ambient Wood. This extension comprises the second axis of virtual device, in particular, the consideration of multiple meeting rooms with distributed participants and the modelling of interactions between them.

Currently, there is little to distinguish a remote participant in a meeting from a local one, they are both NPCs and both can contain or interact with resources in both physical worlds, abstracting away from the underlying implementation, e.g. the videoconferencing tools.

5 Discussion

The following section discusses some of the more interesting issues and unforeseen aspects regarding the deployment of a MUD for orchestration in the two scenarios.

Decoupling and spoofing

Physical world events in the Ambient Wood scenario were realised as multicast notification messages, that were routed to interested processes depending on the content of the notification. For example, the interaction between labelled, positional pinger and Team-Badgers' hand-held generated a location notification that would be delivered to any process registered as interested (subscribed) to team location updates. The adoption of the lightweight and flexible Elvin notification infrastructure meant that it was possible to efficiently decouple physical interactions from the processes that observe or orchestrate them and thus reduce the processing requirements of the various devices in the wood.

An anticipated and much welcomed benefit of the decoupling of physical interaction from virtual model was the ability to spoof, or 'Wizard of Oz', interactions in software in the lab to carefully test and analyse all modelled interactions without being dependent on the physical hardware. Physical interactions could be mimicked in the MUD both by generating the notification by hand, or by direct manipulation of the MUD's model. This was especially important given the relatively short time from experiment specification through to field trials in the Wood. Furthermore, during the experiments, it was possible to compensate for any physical device failure by directly manipulating the MUD.

As regards integration of Elvin with the MUD, the use of the MEAP proxy process was sufficient to explore the issues around orchestration of the model. However, a more thorough implementation would be to integrate the notification service into the MUD's internal event model. This would be achieved by extending the functionality of the driver code such that MUD objects would register handler functions that the MUD driver would call upon receipt of appropriate (matched) notifications. This integration activity is an on-going effort.

Scenario definition

The nature of the MUD data model means that the virtual representation of the scenarios closely resemble the physical environment, simplifying the task of conceptualising scenarios for experimentation. The MUD also offers the facility to abstract the physical world sensing and actuating devices such that the scenario coder can rely on events such as *"object X is here"* rather than low level event data such as *"ping received from $LWood4;"*.

In both of the scenarios, the relationship between the physical world, its 'sensors' and 'actuators' have been manually codified in the MUD's logic. If only for simplicity of use, an abstraction for describing the physical scenario to configure the virtual representation and define the set of interactions and transitions available is required. An example of how this might be achieved would be to build an editing and visualisation application that prescribes configuration parameters to objects that are to be co-ordinated between the physical and virtual worlds, resulting in an XML model description document that is rendered into MUD game code, rather than requiring scenario editors to be proficient in programming languages such as LPC.

In the case of the Ambient Wood scenario, this would enable teachers and classroom assistants to more readily visualise the environment in which the experiments would take place both before and after the actual visit to the wood. The text-based simplicity of MUDs, combined with such a tool, would enable rapid prototyping and experimentation by educationalists as opposed to computer scientists.

Rich information

Having demonstrated that the MUD approach to orchestration can model the desired interactions between physical and virtual environments, we are keen to extend the domain of information for which the MUD has responsibility.

As well as modelling contextual information, such as participant or object location, the MUD could model 'information', applying Adaptive Hypermedia techniques to control the presentation and delivery of that information based on the current activity of participants within the scenario. In the Ambient Wood experiment, the MUD triggered different pre-rendered voice-overs to be played on the team's hand-held based on repeated visits to a particular location in the wood. Modelling the information in the MUD would enable adaptive delivery based on where the team had already been, or what they have previously interacted with. This would enable more complicated scenarios to be defined.

Distributed Orchestration

In both scenarios, the MUD forms the only 'centralised' infrastructure component; there is one MUD orchestrating the scenario in both environments. There are scenarios in which it would be beneficial for the orchestration facilities to be realised across a distributed environment. For example, two sets of school children experimenting in two different woods, sharing readings and observations, where the actions of one school party might affect the readings and observations of the other. Whilst this could be modelled and orchestrated with one centralised MUD, the delegation of ownership to localised MUDs might reveal interesting properties.

Likewise with FEEL's Meeting Room scenario, the distributed orchestration of multiple meeting rooms introduces new modelling and interaction challenges for a MUD-based infrastructure. It would require co-ordinating both local and remote resources in the interactive meetings, requiring more complex models of intrusiveness reflecting the different requirements.

6 Summary

Our initial experiences with the use of MUD software to orchestrate the physical and virtual interactions in two different information-rich environments has been both successful and rewarding with regards several key observations:

- The MUD model is a simplified representation of the objects, locations and interactions that were are interested in, as opposed to an exact replication of the real world. This allows us to focus on particular things of interest.
- Further work needs to be done on abstracting the scenario logic from the MUD. This will allow simple scenario construction by those not wishing to understand the MUD programming language.

- The abstraction provided by the MUD helped separate the construction of the devices from the development of the scenarios, i.e. scenarios detailed 'this event' happening when the participant was doing 'this task' without requiring knowledge of how that task information was acquired.
- During occasional device failure, the MUD allowed the scenario to continue with a virtual participant modifying the MUD manually to simulate events in the real world. This is only possible with a decoupled approach.

We are currently working on extending the model across distributed MUDs to connect multiple meeting rooms with different requirements. We are also working on developing scenario definition tools and more fully integrating information systems into the environment.

Acknowledgements

This Technical Report is a work-in-progress.

This research is funded in part by EPSRC Interdisciplinary Research Collaboration "EQUATOR" GR/N15986/01, and by project IST-2000-26135: Non-Intrusive Services to Support Focused, Efficient and Enjoyable Local Activities, "FEEL".

The authors wish to acknowledge their partners on the Ambient Wood and FEEL projects. No badgers were harmed during the trialling of the orchestration tools.

References

[1] ARNOLD, D., SEGALL, B., BOOT, J., BOND, A., LLOYD, M., AND KAPLAN, S. Discourse with disposable computers: How and why you will talk to your tomatoes. In *Usenix Workshop on Embedded Systems (ES99)* (1999). Cambridge, MA.

[2] BARTLE, R. Interactive multi-user computer games. Tech. rep., BT Martlesham Research Laboratories, Dec. 1990. Annotated version available from `http://www.mud.co.uk/richard/imucg.htm`.

[3] BERNSTEIN, M. Card shark and thespis: exotic tools for hypertext narrative. In *Proceedings of the twelfth ACM conference on Hypertext and Hypermedia* (2001), ACM Press, pp. 41–50.

[4] CHURCHILL, E. F., AND BLY, S. Virtual environments at work: ongoing use of MUDs in the workplace. In *Proceedings of the International Joint Conference on Work Activities Coordination and Collaboration (WACC'99), San Francisco, CA* (1999), D. Georgakopoulos, W. Prinz, and A. Wolf, Eds., pp. 99–108.

[5] DE ROURE, D. C. Introducing the declarative dungeon. In *Proceedings of Programming Languages: Implementations, Logics, and Programs. 9th International Symposium, PLILP'97. Including a Special Track on Declarative Programming Languages in Education* (1997), H. Glaser, P. Hartel, and

H. Kuchen, Eds., no. 1292 in Lecture Notes in Computer Science, Springer Verlag, pp. 407–17.

[6] Evard, R. Collaborative networked communication: MUDs as Systemtools. In *Proceedings of the Seventh Systems Administration Conference (LISA VII), Monterey, CA* (Nov. 1993), pp. 1–8.

[7] Greenhalgh, C. Equip: a software platform for distributed interactive systems. Tech. Rep. Equator-02-002, Nottingham, April 2002. Available as http://www.equator.ac.uk/papers/Ps/2002-greenhalgh.pdf.

[8] Greenhalgh, C., Bowers, J., Walker, G., Wyver, J., Benford, S., and Taylor, I. Creating a live broadcast from a virtual environment. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques* (1999), ACM Press/Addison-Wesley Publishing Co., pp. 375–384.

[9] MacColl, I., Millard, D., Randell, C., and Steed, A. Shared visiting in EQUATOR City. In *Proceedings of Collaborative Virtual Environments (CVE2002)* (Sept. 2002), Sheridan.

[10] Ramchurn, S. D., and Jennings, N. R. Multi-agent architecture for non-intrusive mechanisms. Tech. Rep. Deliverables 3.1 and 3.2 of the IST FEEL Project, Department of Electronics and Computer Science, University of Southampton, June 2002.

[11] Richardson, T., Stafford-Fraser, Q., Wood, K. R., and Hopper, A. Virtual network computing. *IEEE Internet Computing 2*, 1 (Jan. 1998), 33–8.

[12] Rogers, Y., Price, S., Harris, E., Phelps, T., Underwood, M., Wilde, D., Smith, H., Muller, H., Randall, C., Stanton, D., O'Neale, H., Weal, M., Thompson, M., Michaelides, D., and Cruickshank, D. Out of the classroom and into Equator: Reflections on the Ambient Wood project. Tech. Rep. to appear, EQUATOR Technical Report, October 2002.

[13] Segall, B., and Arnold, D. Elvin has left the building: A publish/subscribe notification service with quenching. In *Proceedings of Queensland AUUG Summer Technical Conference* (1997). Brisbane, Australia. Available as http://www.dstc.edu.au/Elvin/doc/papers/auug97/AUUG97.html.

[14] Smith, J. *The MUD FAQ*. The MUD Connector. Continually revised, available at http://www.mudconnect.com/mudfaq/.

[15] Weal, M. J., Millard, D. E., Michaelides, D. T., and Roure, D. C. D. Building narrative structures using context based linking. In *Proceedings of the Twelfth ACM Conference on Hypertext and Hypermedia* (2001), pp. 37–38.

[16] Wyckoff, P., McLaughry, S., Lehman, T., and Ford, D. TSpaces. *IBM Systems Journal 37*, 3 (Aug. 1998), 454–74.