

The Properties of Mixed Reality Boundaries

Boriana Koleva, Steve Benford and Chris Greenhalgh
The University of Nottingham, Nottingham NG7 2RD, UK
{bnk,sdb,cmg}@cs.nott.ac.uk

Abstract: Mixed reality boundaries establish transparent windows between physical and virtual spaces. We introduce a set of properties that allow such boundaries to be configured to support different styles of co-operative activity. These properties are grouped into three categories: permeability (properties of visibility, audibility and solidity); situation (properties of location, alignment, mobility and segmentation); and dynamics (properties of lifetime and configurability). We discuss how each of these properties can be technically realised. We also introduce the meta-properties of symmetry and representation. We then describe and compare two contrasting demonstrations, a performance and an office-door, that rely on different property configurations.

Introduction - approaches to shared mixed reality

There has been a growing interest in techniques for combining real and virtual environments to create *mixed realities* – spatial environments where participants can interact with physical and digital information in an integrated way (Milgram & Kishino, 1994). Mixed realities may be shared, enabling people who are distributed across multiple physical and virtual spaces to communicate with one another. A variety of approaches to creating a shared mixed reality have been demonstrated, including augmented reality, augmented virtuality, tangible bits and mixed reality boundaries.

Augmented reality involves overlaying and registering digital information (e.g., text and graphics) onto a real world scene in such a way that the digital information appears to be attached to physical objects, even as they move about.

The physical scene might be the local environment, with the digital information being introduced via a see-through head-mounted display. Alternatively, it might be remote, being viewed on a video display that is then enhanced with digital information. Early examples of collaborative augmented reality include the *Shared Space* system (Billinghurst et al, 1996), in which users share virtual objects across a physical table top and *Studierstube* (Schmalstieg et al, 1996), in which virtual objects are also displayed in a physical space between multiple users. Both of these systems utilise see-through head-mounted displays. Systems based on video views of remote scenes are inherently sharable as the video display is usually located in a shared physical space.

In contrast, *augmented virtuality* (Milgram & Kishino, 1994) takes a virtual world as its starting point and then embeds representations of physical objects within this. These might take the form of textured video views, for example views of participants' faces on their avatars as in the *Freewalk* system (Nakanishi et al, 1996), or views of remote physical locations as in the 3-D media-space interface of (Reynard et al, 1998). Alternatively, telemetry data captured by remote physical sensors might be visualised using graphics, text and audio.

The approach of *tangible bits* (Ishii & Ullmer, 1997) involves the use of graspable physical objects called phicons to interact with digital information, for example moving physical models across a table top in order to access a digital map that is projected onto it. This may be coupled with the use of ambient display media such as sound, light and airflow to provide more peripheral awareness of background information, for example, showing the volume of network traffic as reflections of water ripples on the ceiling.

The approach of *mixed reality boundaries* involves joining together distinct virtual and physical spaces by creating a transparent boundary between them (Benford et al, 1996). With this approach, the spaces are not overlaid, but instead are distinct but adjacent. The occupants of the shared physical space can see into the next-door virtual space and can communicate with its occupants (e.g. avatars within a collaborative virtual environment). In turn, the occupants of the virtual space can see back into the physical space. A distinguishing feature of this approach is that it places equal weight on physical and virtual environments, considering how each can be accessed from the other. It also offers the potential to use multiple mixed reality boundaries to join together many physical and virtual spaces into a larger mixed reality environment in the same way that everyday boundaries such as doors, walls and windows are used to structure physical buildings.

Our paper is concerned with this last approach. Its departure point is the idea of a simple mixed reality boundary as described in (Benford et al, 1996). Figure 1 shows how such a boundary can be established. On the left of the figure is a physical environment into which are projected graphics and audio from the virtual environment on the right. In turn, a video camera and microphone capture video

and audio from the physical environment and this is transmitted back to the virtual environment over a computer network. The live video image is then displayed as a dynamic texture map within the virtual environment. The result is the creation of a transparent bi-directional window between the physical and virtual environments. This approach was demonstrated through an application called the Internet Foyer, in which a visualisation of an organisation's home pages on the World Wide Web, complete with representations of their visitors (a virtual foyer), was joined to its physical foyer using a mixed reality boundary.




Figure 1: a simple mixed reality boundary

This paper further develops this approach by identifying a set of properties that can be associated with mixed reality boundaries. These properties are intended to support the design of mixed reality boundaries for a broad range of potential collaborative applications. Applications as diverse as distributed meetings, performances, media-spaces, document editing and 3-D design will have varying requirements for managing awareness and privacy; for positioning a boundary and aligning it to different participants; and for scheduling its appearance. The set of boundary properties is also intended to provide an analytic framework for reasoning about how different boundary configurations (e.g., based on different combinations of projection and camera technologies) might afford different styles of co-operative activity.

The following section introduces our boundary properties and explores their technical realisation. We then present two contrasting demonstrations of how these properties can be configured to support different co-operative activities. The first is a distributed performance between a poet in a virtual world and an audience in a physical theatre. The second involves the use of a persistent boundary to allow remote visitors to "drop in" to an office over a network. We offer a property by property comparison of their design and conclude by considering how this work can draw on and contribute to related research areas such as tangible interfaces (Ishii & Ullmer, 1997) and co-operative buildings (Streitz, 1998).

Boundary Properties

We now introduce the fundamental properties of mixed reality boundaries, examining the utility of each and considering how each might be technically realised. Our choice of properties has been influenced by analogies with the real-world boundaries that partition physical space and also by our previous work on developing boundaries within virtual space (Benford et al, 1997a). Our proposed properties are grouped into three general categories: *permeability*, properties that define how information passes through a boundary; *situation*, the spatial properties of the boundary; and *dynamics*, the temporal properties of the boundary. We also introduce the meta-properties of *symmetry* and *representation* that apply to the other properties.

Permeability

Permeability describes how the boundary affects sensory information passing between the linked spaces. We break down permeability into visibility, audibility and solidity, based on the three primary types of information that can pass through the boundary. Our discussion focuses on vision, sound and touch because most current interfaces between the physical and the virtual are based on a combination of these. However, we note that smell and taste information might also be "transmitted" through mixed reality boundaries in the future.

Visibility – considers what visual information is permitted through the boundary and consists of two components: visual resolution and field of view. Visual resolution concerns the amount of visual information obtained through the boundary and is affected by factors such as the resolution of capture and display technologies and graphical level of detail. Field of view describes the volume of the connected space that is made visible through the boundary and is determined by factors such as the field of view/projection of (physical and virtual) cameras and projectors.

Audibility – considers what audio information is permitted through the boundary and is determined by factors such as the positioning and sensitivity of microphones as well as sampling rates.

Drawing on previous work on virtual boundaries and crowd representations (Benford et al, 1997a), we propose that visibility and audibility can be further described in terms of the combination of four effects:

- *attenuation* – for example, reducing video resolution or audio volume;
- *amplification* – for example, projecting audio in the manner of a public address system;
- *transformation* – for example, distorting audio and video to establish anonymity; and
- *aggregation* – summarising what lies beyond the boundary. For example, showing only the number of remote participants instead of each individual.

Solidity – refers to the ability to traverse the boundary. This includes metaphorically extending a limb through the boundary to manipulate or feel a remote object; pushing an object through the boundary so that it becomes available on the other side; and stepping through the boundary and assuming control of an avatar or physical proxy on the other side. Strictly speaking, this last case establishes a second mobile boundary between the spaces (see below) as the participant may not actually leave their local physical space behind. However, metaphorically, there is a sense of stepping through the boundary.

Traversal from the physical to the virtual can be realised using 3-D interaction devices and tracking technologies to manipulate virtual objects that appear on a projected display or to track local physical objects and update their virtual counterparts. Allowing the user to sense virtual objects is achieved through haptic devices such as those described by Fogg et al (1998) and Colwell et al (1998). Traversal from the virtual to the physical involves remote control of physical proxies such as mobile cameras and robots as in the *GestureCam* system (Kuzuoka et al, 1995). Digital documents can be pushed through the boundary by projecting them onto a desktop in the manner of the *Digital Desk* (Wellner, 1993) or by placing them directly on the boundary itself in the style of *Clearboard* (Ishii & Kobayishi, 1992).

The potential for combining different visibility and audibility effects with varying degrees of solidity allows the definition of a wide range of boundary types. These include analogies of familiar everyday physical boundaries such as windows, walls, curtains, fences, one-way mirrors and even lines on the ground, as well as new kinds of boundary that have no common physical counterpart. Furthermore, a systematic exploration of all possible combinations of visibility, audibility and solidity might identify the fundamental building materials that can be used to join together physical and virtual spaces.

Situation

Situation concerns the spatial relationships between the mixed reality boundary, the physical and virtual spaces that it connects and the participants and objects that these contain. This includes the location of the boundary, whether this location is fixed and whether the boundary is segmented. Between them, these properties determine the spatial understanding that the participants in one space have of the connected space.

Location – describes the placement of the boundary within the connected spaces. A vertical location involving projection onto a physical wall or screen or texturing onto a virtual wall or screen will establish the boundary as a window between the two spaces. Given a large enough display, the remote space might even be presented as a direct extension of the local space. A horizontal location involving projection onto a physical or virtual desk or board will establish the boundary as a shared drawing surface. The use of ambient display media as

proposed in (Ishii & Ullmer, 1997), could establish a more peripheral connection between the spaces.

Alignment – concerns the orientation of the boundary with respect to the different participants and objects and may establish different possibilities for turn taking and access to other participants. For example, the triangular alignment of figure 2 allows a physical performer to simultaneously address a physical and a virtual audience while allowing the two audiences to address one another.




Figure 2: a triangular boundary alignment

Mobility – describes whether the boundary assumes a static situation, thus offering connection between two fixed sections of the linked spaces, or whether the boundary can join different parts of the spaces over time. A mobile boundary is one that the participants can steer through the linked spaces or which follows a pre-programmed trajectory. In practical terms, mobility requires that the boundary components (e.g., physical or virtual cameras and microphones) can themselves be moved. There can be various restrictions on the movement of the boundary. For example, participants may only be able to rotate the boundary around a fixed point without translating its position.

Segmentation - a boundary can be segmented in terms of its properties and spatial location. The former refers to when a boundary is made up of one or more segments with distinct property sets. A spatially segmented boundary, on the other hand, links the two spaces through multiple non-adjacent segments (these can themselves be property segmented).

So far in our discussion of the situation properties we have assumed that the spatial co-ordinate system of one space is related to that of the other in a spatially consistent way. More specifically, that the connected spaces provide to some extent a unified frame of reference across which position, distance, orientation and perspective are consistent. However, there are two problems in establishing detailed spatial consistency. First, the use of a single fixed camera (physical or virtual) on any side will only provide an accurate perspective from one viewing position. A participant may move along the boundary or change their viewing angle, but will still retain the same view of the connected space. The use of

multiple cameras, mobile cameras and even stereo cameras may overcome this problem to some degree, but only for one participant at a time. Second, it may be necessary to locate audio information in a spatially consistent manner. This may require the use of multiple or mobile microphones and spatialised audio rendering.

The issue of spatial consistency is a very important one as it affects the level of detail to which participants on opposite sides can establish mutual orientation, a reciprocity of perspective and can use spatial language and gesture.

Combining these situation properties with the permeability properties described previously, defines the regions of each space that are public (i.e., accessible from the connected space) versus those that remain private to each space (i.e., are out of camera shot or microphone range).

Dynamics

Dynamics concern the temporal properties of the boundary, including its lifetime and its degree of configurability.

Lifetime – refers to when and for how long the boundary is in existence. Boundaries may be scheduled to appear at specific, even periodic, times to support the pre-planned nature of many activities (e.g., performances and meetings) or may be created on the fly. The potential duration of a boundary can be considered in the light of previous research into media spaces that distinguished different services ranging from persistent "office share" connections through to short-term "glance" facilities lasting for just a few seconds (Gaver et al, 1992).

Configurability – describes how dynamically the various boundary properties can be changed. Permeability might be adjusted in order to reflect changing privacy requirements. Configuring situation involves being able to move cameras and projectors and re-positioning furniture and other aspects of the connected spaces, for example in order to accommodate new participants. Finally, dynamic properties such as lifetime might also be directly configurable.

So far, we have established the fundamental boundary properties of permeability, situation and dynamics. We now discuss the two meta-properties of symmetry and representation that relate to each of these.

Symmetry

Symmetry refers to the extent that the properties of a mixed reality boundary are configured to be the same on both of its sides (i.e., from the physical to the virtual and vice versa). A degree of asymmetry may often be imposed as a result of the technologies used (e.g., where cameras and projectors differ in their field of view). In other cases, it may be desirable to deliberately create asymmetric boundaries in order to meet a specific communication need (e.g., using a one way

boundary to unobtrusively observe activity). Mixed reality boundaries may be asymmetric with respect to permeability, situation and dynamics.

Asymmetry introduces an additional dimension to the configurability of boundaries. We propose that participants should be able to configure their own side of the boundary and also set limits on the potential configuration of the other side as it affects them. For example, a participant may wish to set an upper limit on what the other side can see of them. To generalise, each control for configuring a boundary property might combine the ability to set the property in one direction and limit the property in the reverse direction.

Representation

Mixed reality boundaries are potentially complex technologies that may take many different forms. We argue that, in order to successfully use a boundary, participants will need to understand both the current and potential settings of its properties. In other words, the properties of mixed reality boundaries should be made visible (and possibly audible), be it explicitly through controls and labels, or implicitly through metaphor, interior design or architecture. Considering permeability, a boundary should indicate the current and potential settings for visibility, audibility and solidity at each side. Considering situation, the separation of public from private space should be clearly marked so that participants know how to position themselves in order to communicate or avoid those on the other side. For example, the view frustra of physical and virtual cameras could be made explicitly visible by marking them on the floor. Considering dynamics, participants might be notified in advance when a boundary is going to appear or disappear so as they may adjust their behaviour and/or appearance appropriately.

This concludes our introduction to the properties of mixed reality boundaries. The following section presents two examples of how these properties might be configured to meet different application requirements.

Demonstrations

We present two demonstrations of mixed reality boundaries that rely on different configurations of boundary properties to support different activities:

- a performance in which a performer on a virtual stage engages an audience in a physical theatre through a mixed reality boundary;
- an "office door" that establishes an open connection between a public virtual world and a private physical office. This necessitates the management of virtual visitors in relation to local physical activity, especially with regard to shifting privacy requirements.

For each demonstration we state its goals, describe its design and offer some initial reflections. We then compare the two in terms of the property configurations of their boundaries.

First Demonstration – a Mixed Reality Performance

Our performance demonstrator extends our previous experience of staging a poetry performance simultaneously in physical and virtual theatres as reported in (Benford et al, 1997b). This previous attempt involved poets performing in a conventional physical theatre and at the same time, appearing as avatars on a virtual stage in front of an on-line audience in a virtual environment. A view of the virtual environment was then projected as a back-drop to the physical stage. A key observation from this previous performance is that the event became fragmented into two parts – a conventional performance and a social-chat virtual environment. We have argued in (Benford et al, 1997b) that the nature of the projection of the virtual space into the theatre may have been a key factor in these problems. In particular:

- the projection created a one way boundary between the two spaces – the physical audience and performer could see their virtual counterparts, but not vice versa;
- for aesthetic reasons, the projection was rendered from a moving viewpoint. Consequently, there was no stable spatial relationship between the two spaces and it would be difficult for the participants in the physical theatre to establish any consistent reference or orientation to those in the virtual theatre.

The current demonstrator has therefore focused on the issue of whether effective social engagement can be established between real and virtual theatres. In this case, the performer (a poet) appeared on a virtual stage and attempted to engage the attention of an audience who were located in a remote physical theatre. The poet attempted to persuade the audience to join in the performance by answering questions, standing up and chanting as part of an improvised poem – essentially a test of whether they could exert sufficient social pressure on the audience. A key design goal was therefore that the boundary should be as invisible as possible, especially to the audience. Specific differences to the previous performance were that:

- the physical and virtual worlds were linked through a mixed reality boundary that allowed the audience to see and hear the virtual poet and vice versa;
- the boundary had a fixed spatial location with the intention that the virtual stage would appear to the audience as a conventional physical stage would;
- the poet was physically separated from the audience and was immersed in the CVE using a head-mounted display. As a result, the only option for communicating with the audience was via the mixed reality boundary.

Technical Realisation of the Performance

Figure 3 summarises the realisation of the performance.

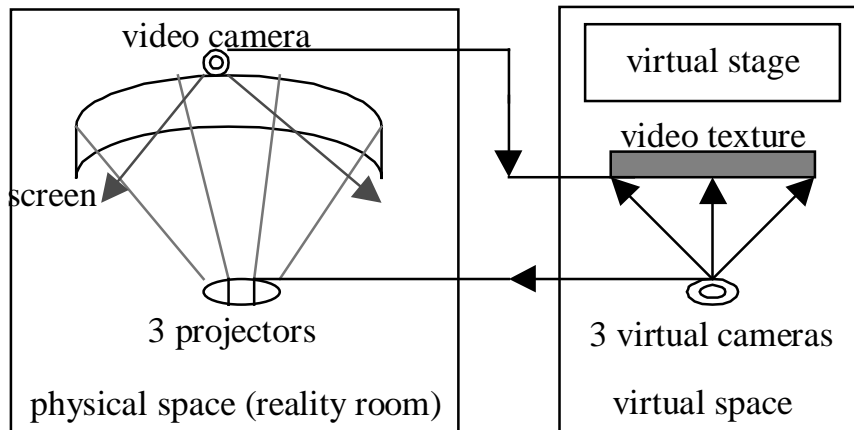


Figure 3: the design of the mixed reality performance

The physical audience were located in a Reality Room, a facility that allows the output of a high-performance computer (Silicon Graphics Infinite Reality Engine) to be projected onto a wide angle curved screen (using three separate projectors whose output is blended together). The poet was located in a separate physical room and used an Eyegen 3 head-mounted display to become immersed in the virtual and appear on the stage. Their view included a video texture looking back out into the audience space taken from a video camera mounted at the top and centre of the projection screen. Note that the video texture was transparent from the other side and did not interfere with the virtual cameras. A separate analogue circuit provided an audio link between the two spaces. By positioning a microphone in the audience space at the focal point of the curved screen we could easily pick up any noise made by the physical audience, a useful additional feature of placing a curved screen in a shared space. Finally, in order to introduce an extra element to the performance, an additional virtual actor was able to enter the virtual stage using a workstation that was located in the audience space.

Experience from the Performance

Our initial performance lasted for half an hour and involved one of the poets from the previous performance (Dave "Stickman" Higgins). The stage and poet's avatar were designed by the artist Derek Richards who also joined in the end of the performance as a supplementary actor. The performance began with the virtual poet entering from the wings and improvising a poem while the audience watched. After five minutes the poet directly addressed the audience for the first time, requesting that they stand up and asking them a series of questions. After picking on several other individuals in the audience, he then persuaded them all to rise and to clap and chant along with the poem. Figure 4 is taken from the

audience space and shows the poet avatar on the virtual stage addressing an audience member.




Figure 4: the poet interrogates an audience member

This initial test suggested to us that the poet was able to engage the audience to some degree. They did respond to his requests, although on some occasions (e.g., on first asking them all to stand up) this took several attempts, emphasising the importance of visibility through the boundary (it was clear that he could see when they had not responded to his request). The resolution of the textured video made it impossible for the poet to pick out details such as facial expression, but he was able to pick out gross physical features such as clothing and to spot large gestures. There were several problems with the performance, especially with the poet becoming disorientated. However, we argue that this simple test demonstrated a level of engagement between the real and virtual that was not achieved in our earlier performance.

In some ways, these observations reflect the current successful use of virtual actors to engage physical audiences through large screens at entertainment events and installations, see for example, (SimGraphics, 1999). However, in these events the human actor typically adopts an out of body view and sees the remote audience on a separate video monitor. Our experiment involved full-immersion with a textured video view of the remote audience being presented as a window in the virtual world.

Second Demonstration – the "Office Door"

In contrast to the previous demonstrator that focused on a specific event, the "office door" demonstrator explores how a mixed reality boundary might be configured to establish a persistent connection between a physical and virtual space. The aim of the demonstrator is to connect a private physical office to a public virtual corridor to enable remote visitors to drop by over a computer network. This has been inspired by previous work on media-spaces that introduced the ideas of "glancing" into remote offices and establishing long-term "office-share" relationships (Gaver et al, 1992). However, in this case, one of the connected spaces is a collaborative virtual environment. An important aspect of this demonstrator is that, in contrast to the performance, it raises the issues of dynamically managing and representing availability and privacy when using mixed reality boundaries.

Technical realisation of the office door

To create the office door boundary, we have projected a view of a virtual corridor onto the wall of a private office. At the same time, we have texture mapped the reverse video view of the office into the virtual corridor. By using a workstation on their desk, the occupant of the office can also step into the corridor, taking on the form of an avatar. Thus, the physical side of the boundary is not solid. Two potential views of the corridor are therefore available to the occupants, a permanent view looking out of their office (wall projection) and sometimes a mobile view from within the virtual corridor (using a workstation).

We have extended the basic mixed reality boundary design to offer varying degrees of visibility and audibility. On the virtual side, the volume of audio from the physical can be adjusted as can the resolution of video (down to no audibility or visibility). On the physical side, audio volume can also be adjusted as can the level of detail of the graphical view of the virtual corridor. For the latter, the current demonstrator supports four levels of detail: no visual information, an aggregate count of how many people are beyond the boundary, indication of the positions of these people (they are represented as simple blocks) and finally a full view of individual avatars. These levels of audibility and visibility are directly configurable at each side of the boundary through a series of interface controls that also indicate their current settings. Specifically:

- each side of the boundary provides a control for setting and indicating the visibility of the other side. A parallel control is provided for audibility.
- each side of the boundary provides a separate control for limiting the maximum visibility of this side to the other. Using this control, participants can set an upper limit on how visible they wish to be, including reducing their visibility to zero. A parallel control is provided for audio.

By using these controls, participants can dynamically negotiate degrees of privacy. It is important to note that levels of visibility and audibility need not be symmetric across the boundary.

Figure 5 offers a screenshot of these controls as seen from the virtual corridor. The two sliders at the side of the video texture set the desired audibility and visibility of the other side. The two sliders above the texture limit the ability of the other side to see and hear this side. We can see two avatars in the virtual corridor looking into the physical office and one person looking back at them. The part of the corridor that is visible from the physical office is shaded a different colour to the part that is not (although this is difficult to see in the greyscale image).




Figure 5: Controls for visibility and audibility

Experience of the Office Door

We installed the office door boundary in two offices within our laboratory, one that was shared by four researchers and a second that was the private office of our laboratory manager. The boundary was established as an open-connection between these spaces and the virtual corridor. It was also used during regular weekly laboratory meetings. The latter involved three or four members of the laboratory attending these meetings remotely and appearing in the virtual corridor, with the remaining participants (between five and eight) being physically present in the office. For the shared office, the boundary was placed diagonally across a

corner of the room so that the participants could arrange themselves in a circle when seated. For the private office, it was placed at the end of an existing meeting table.

Our initial reflections have raised a number of issues for further exploration. First, although virtual visitors appeared visually on a large projection and could be heard as soon as they spoke, the silent nature of movement in the virtual corridor was potentially disconcerting for those in the physical space. Sensing presence and movement near the boundary and indicating these through additional audio cues may prove useful, especially on the virtual side of the boundary. Second, several different avatars were tested for the virtual visitors. Those featuring a live video face seemed best suited to this particular application, perhaps because they offered a degree of symmetry in terms of providing a reverse video view back into their user's own physical space. In particular, they would make it possible for occupants of a physical office to tell when several remote users people were looking at them through a single shared avatar. Third, the circular arrangement of participants worked adequately for the laboratory meetings. Participants in the virtual space claimed that, at least some of the time, they could tell when they were being looked at and could identify individuals in the physical space (although, unlike the performance, the participants knew each other well). We suspect that swapping turns across the boundary as opposed to our usual progression around the circle may have promoted communication between the physical and virtual spaces.

Comparison and Discussion

We conclude our presentation of the two boundary demonstrations by directly comparing and discussing their properties. Our aim is to show how their different configurations reflect their intended uses. Table 1 summarises this comparison.

We begin with permeability. Visibility and audibility were fixed for the performance boundary with the aim of achieving the best possible visual and audio experience for both sides. In contrast, they were defined to be dynamically configurable for the office door boundary, allowing the participants to select a configuration best suited to their current requirements for awareness and privacy. In terms of solidity, both demonstrators allow one person to metaphorically step through the boundary from the physical to the virtual by controlling an avatar. In both cases, the boundaries were solid from the virtual to physical, due to a lack of locally available tele-presence technology (e.g., remote controlled robots).

There were also differences between the two boundaries in terms of their situation properties, especially alignment and location. With the performance, the aim was to link the whole physical room with all of its audience members to the virtual stage, making the two spaces appear to be direct extensions of each other. In other words, the boundary was intended to be as transparent as possible. This was facilitated by the use of a wide screen projection facility housed in a purpose

built room. Indeed, a key element of the performance was the moment of surprise when it was revealed that the performer could actually see the audience through the projection screen! In contrast, an important aspect of the office door alignment was the desire to deliberately keep one area of the office out of view of the boundary so as to retain a private area. The office door alignment was also limited by the physical shapes of the existing offices. This severely constrained the location of screens and projectors.

Property	Performance		Office door	
	Virtual	Physical	Virtual	Physical
Visual resolution	Video resolution 120 x 120 pixels	Projector resolution 3556 x 1024 pixels	Configurable video resolution from 120 x 120 to 0 x 0 pixels	Configurable graphical level of detail (4 levels)
Field of view	60°	175°	60°	65°
Audibility	Amplified		Variable – volume can be adjusted at both sides	
Solidity	Solid	one person can step through	Solid	one person can step through
Location	Vertical - establishes boundary as window	Vertical - establishes boundary as extension of space	Vertical - establishes boundary as window	
Alignment	Facing seats	Facing stage	Into part of office	Onto corridor
Mobility	Static		Static	
Segmentation	Property segmented		Not segmented	
Lifetime	Half hour		Persistent	
Configurability	None		Visibility, audibility	

Table 1: summary of boundary properties of the office door and the performance

The performance boundary raises an additional issue, that of property segmentation. As the table shows, there is a difference between the field of view offered by the projection (175 degrees) and by the video texture (60 degrees). In effect, this created a property segmented boundary consisting of a central segment which was roughly symmetrical with respect to visibility and audibility and two outer segments that were asymmetric (the audience could see the performer, but there was no reverse video texture by which the performer could see them). This segmentation was an accidental side-effect of the locally available camera and projection technologies. However, it could have had significant effects on social interaction, as it created areas of the boundary that were in effect one-way

mirrors. We anticipate that such situations might arise regularly where boundaries exploit existing local facilities. Furthermore, we propose that our boundary properties provide a useful analytic framework for predicting when such accidental effects are likely to occur.

A key difference between the two boundaries concerns their dynamics. In the case of the performance, the boundary was created at a specific time for a pre-planned event, it had a set duration and its properties remained static throughout its lifetime. The office door boundary, on the other hand, had an open-ended lifetime. The connection of a public space to a private space in which different activities could occur (including private consultations), necessitated the dynamically configurable nature of the boundary as was realised through the controls mentioned above.

Considering their symmetry, the performance boundary was intended to be symmetrical and any asymmetries were side effects of the technologies that were used. In contrast, the potential for asymmetry was deliberately designed into the office door. For example, the occupant of the office might have closed down their side of the boundary during a private meeting while retaining a view into the virtual corridor to see if any potential visitors were waiting.

Issues of representation of properties were particularly significant with the office door due to its configurability (i.e., it was necessary to convey the current and potential state of both sides of the boundary) and also due to the need to mark the distinction between public and private space. In contrast, the intended transparency of the performance boundary led to its properties not being directly represented.

Finally, considering spatial consistency, there was slight vertical misalignment at the virtual side of the boundary in both demonstrators. This is due to the video camera being located at the top of the projection screen for the performance and at the bottom for the office door as opposed to at its vertical midpoint.

Positioning the camera at the vertical midpoint behind the projection would have been ideal, but would have required a special screen that was transparent from the rear. In both cases, the correct alignment was achieved at the virtual side because the video texture was defined to be a one sided polygon through which the virtual camera could see.

In summary, we have presented two contrasting demonstrations of mixed reality boundaries that exploit different configurations of boundary properties in order to achieve their goals. The performance seeks to establish from scratch a boundary that is invisible to its users and that exists for a fixed period of time. In contrast, the office door requires that the boundary be integrated in an existing working environment and so must be visible, understandable and configurable. A further key point concerns the way in which purely technological factors (such as differences in the field of view of available cameras and projection equipment) can result in accidental asymmetries and segmentation of boundaries. We argue

that our notion of boundary properties provides a framework for understanding when these might occur and how they might affect participants.

This concludes our analysis of mixed reality boundary examples in terms of properties. In the next section we generalise the idea of boundaries and use it to briefly discuss other CSCW systems

Summary and Future Work

This paper has extended the basic idea of mixed reality boundaries by identifying the properties that such boundaries might have. These include properties affecting their *permeability* (visibility, audibility and solidity); properties affecting their spatial *situation* (location, alignment, mobility and segmentation); and properties affecting their temporal *dynamics* (lifetime and configurability). The central argument of this paper has been that through appropriate configuration and analysis of these properties, mixed reality boundaries might support a diverse range of applications. To support this argument, we have described and compared two contrasting demonstrators: the performance and the office-door. We conclude by raising issues for further work.

Implementing the full-range of boundary properties

Creating non-solid boundaries is currently difficult. Traversal from the virtual to the physical requires a greater integration with work on robotics and especially tele-operated physical proxies. In turn, traversal from the physical to the virtual might exploit non-solid projection surfaces that would require a participant to physically step through them in order to gain access to the virtual world beyond. In this way, a participant would leave behind their local physical environment as they stepped into the remote virtual environment. We are currently experimenting with this idea with the UK theatre company Blast Theory who have been experimenting with projecting images into a vertical curtain of water.

We believe that the technique of directly representing boundary properties, especially field of view, is an important one. Future work might consider how this could be achieved more subtly through furniture design, interior design and architecture. For example, the design of carpets, rugs and tables might encourage participants to align themselves appropriately to a boundary. We propose that similar techniques could be used by other applications that make use of cameras.

The idea of mobile mixed reality boundaries is currently unexplored. This raises issues of steering boundaries, integration with mobile cameras and the representation of mobile field of view and other situation properties.

Applying the boundary approach to other areas of CSCW

We propose that we can extend the approach of boundaries to cover not only CSCW systems that link the physical and the virtual, but also CSCW systems that connect two remote physical spaces or two synthetic environments. For example, videoconferencing, media-space and tele-presence applications can be seen as creating boundaries with particular property configurations between remote physical spaces. Could we use our framework of properties as a means of designing and analysing a wider range of CSCW applications? For example, previous work on using multiple video views in media-spaces identified problems in working with fragmented views of a remote scene (Gaver et al, 1993). Could this and other examples, be described in terms of our boundary properties of segmentation and location and the issue of achieving spatial consistency?

Adopting a still broader perspective, shared mixed reality is a topic of growing research interest, with some of the most radical recent developments being made in areas such as tangible interfaces and co-operative buildings. We propose that the approach of mixed reality boundaries can both draw on and contribute to this work. First, ubiquitous computing and tangible interfaces provide techniques for introducing digital information into local physical environments. These might be used to create more accessible mixed reality boundaries. Conversely, work on mixed reality boundaries suggests a greater consideration of how physical environments appear from the perspective of remote and digital environments (e.g., that of networked participants). Second, early work on co-operative buildings has explored how displays and computing facilities might be integrated with furniture, interior design and architecture. As noted above, mixed reality boundaries require the same consideration. Conversely, our approach of directly representing boundary properties within the connected spaces is relevant to the design of co-operative buildings.

In summary, the past few years have seen the development of new ideas for situating the digital in the physical. Our work on mixed reality boundaries is intended to balance this trend by also considering how the physical can be situated in the digital and by developing new techniques for joining the two domains of physical space and digital space on an equal footing.

Acknowledgements

This work has been carried out in the European Union (EU) funded *eRENA* project (ESPRIT IV I³ programme) and the EPSRC funded *Multimedia Networking for Inhabited Television* project (Multimedia Networked Applications programme). We gratefully acknowledge the EU and the EPSRC for their support.

References

- Benford, S. D., Greenhalgh, C. M. and Lloyd, D.(1997a): "Crowded Collaborative Environments", *Proc. CHI'97*, ACM Press, pp. 59-66
- Benford, S. D., Greenhalgh, C. M., Snowdon, D. N. and Bullock, A. N.(1997b): "Staging a Public Poetry Performance in a Collaborative Virtual Environment", *Proc. ECSCW'97*, Kluwer
- Benford, S., Brown, C., Reynard, G. and Greenhalgh, C. (1996): "Shared Spaces: Transportation, Artificiality and Spatiality", *Proc. CSCW'96*, ACM Press, pp. 77-85
- Billingham, M., Weghorst, S., Furness III, T. (1996): "Shared Space: An Augmented Reality Interface for Computer Supported Collaborative Work", *Proc. Workshop on Collaborative Virtual Environments (CVE'96)*, The University of Nottingham, Nottingham, UK.
- Colwell, C., Petrie, H., Kornbrot, D., Hardwick, A. and Furner, S. (1998): "Haptic Virtual Reality for Blind Computer Users", *Third Annual ACM Conference on Assistive Technologies (ASSETS 98)*, ACM Press, pp. 92-99
- Fogg, B., Cutler, L., Arnold, P. and Eisbach, C. (1998): "HandJive: Device for Interpersonal Haptic Entertainment", *Proc. CHI'98*, ACM Press, pp. 57-64
- Gaver, W., Moran, T., MacLean, A., Lovstrand, L., Dourish, P., Carter, K. and Buxton W. (1992): "Realising a Video Environment: EuroPARC's RAVE System", *Proc. CHI'92*, ACM Press, pp. 27-35
- Gaver, W., Sellen, A., Heath, C. and Luff, P. (1993): One is not enough: multiple views in a media space, *Proc. InterCHI'93*, Amsterdam, 1993, ACM Press.
- Ishii, H. and Kobayishi, M. (1992): "Integration of Inter-Personal Space and Shared Workspace: ClearBoard Design and Experiments", *Proc. CSCW'92*, ACM Press, pp. 33-42
- Ishii, H. and Ullmer, B. (1997): "Tangible bits: Towards Seamless Interfaces between People, Bits and Atoms", *Proc. CHI'97*, ACM Press, pp. 234-241
- Kuzuoka, H., Ishimoda, G. and Nishimura, T. (1995): "Can the GestureCam be a surrogate?", *Proc. ECSCW'95*, Kluwer, pp. 181-196
- Milgram, P. and Kishino, F. (1994): "A Taxonomy of Mixed Reality Visual Displays", *IEICE Trans. Information Systems*, Vol E77-D, No 12, December 1994.
- Nakanishi, H., Yoshida, C., Nishimura, T. and Ishida, T. (1996): "Freewalk: Supporting Causal Meetings in a Network", *Proc. ACM Conference on Computer Supported Cooperative Work (CSCW'96)*, Boston, pp. 308-314, ACM Press, Vol 16-20.
- Reynard, G., Benford, S., and Greenhalgh, C. (1998): "Awareness Driven Video Quality of Service in Collaborative Virtual Environments", *Proc CHI'98*, ACM Press, pp. 464-471
- Schmalstieg, D., Fuhrmann, A., Szalavari, Z. and Gervautz, M. (1996): "Studierstube - An Environment for Collaboration in Augmented Reality", *Proc. Workshop on Collaborative Virtual Environments (CVE'96)*, The University of Nottingham, Nottingham, UK.
- SimGraphics Vectors (1999): <http://www.simg.com/vactorlbe.html> (web page as of January 1999)
- Streitz, N. (Ed) (1998): *Proc. International Workshop on Cooperative Buildings (CoBuild'98)*, Springer, Germany, February 1998.
- Wellner, P., (1993): "Interacting with Paper on the DigitalDesk", *Communications of the ACM*, Vol. 36, No 7, July 1993, pp. 87-96