Hybrid Spatial Structure in Ubiquitous Computing

A short paper for the Bigraphs Workshop, First Draft, June 2009

Steve Benford
Mixed Reality Laboratory
University of Nottingham
Nottingham
sdb@cs.nott.ac.uk

Introduction

Bigraphs have been proposed as a suitable approach to modelling ubiquitous computing systems, not least because their inherent support for both containment and linkage might naturally express the complex spatial structures that underlay such systems. In order to explore this potential, Milner has recently shown how Bigraphs can model an example ubiquitous computing system, a prototype educational game in which groups of children role-play being lions, hunting for prey on a virtual savannah that, through the use of GPS, appears to be overlaid on their school playing field [11]. This relationship between real and virtual space is central to Savannah and to several other recent systems, and Milner's paper raises the question of how this might best be modelled. This short paper aims to contribute some background to this debate by revisiting some previous work carried out at Nottingham on the structure of hybrid physical-digital spaces from the fields of ubiquitous computing, virtual reality and mixed reality (part of which involves exploring the relationships between these various fields). The aim of doing this is to introduce some general spatial concepts and structures that it might be useful to model using bigraphs. It addresses three themes in particular: approaches to modelling awareness and communication across space; how hybrid physical-virtual spaces are structured through overlay, containment and adjacency; and the need to model the presence and nature of 'seams' in the underlying ubiquitous computing infrastructure.

Modelling awareness and communication across space

An observational study of Savannah involved capturing video and audio of the children moving across the playing field and then comparing this with replays of system logs of their movements across the virtual savannah (from GPS) in an attempt to understand how they interacted with the system and explain some of the difficulties they experienced [5]. This study revealed several interesting issues that we will refer to throughout this paper, the first of which is the way in which the children would rapidly form and disband groups as they encountered prey in the virtual savannah. This typically involved one child finding some prey, stopping, and then shouting and gesturing to their colleagues so as to gather them into a circle so that they could all see the prey and coordinate their attack (the rules of Savannah required multiple 'lions' to attack the same prey within a short window of time). This raises the question of whether we need to model the nature of spoken, gestural and other forms of communication between participants who share a common physical and/or digital space. It might be tempting to assume that, for a spatially constrained environment such as a playing field, all such communication is globally available, i.e., that all participants can see and hear all others at all times, but this is probably not the case as visual communication requires line of sight and not all spoken or even shouted communication will be intelligible from one side of the field to the other. It is certainly not true where larger spaces are involved, such as in similar games that have been played across several blocks of a city (e.g., [3]).

Previous work form the field of collaborative virtual environments (online shared virtual worlds) has considered this problem as has developed techniques for managing awareness and communication across virtual space. While these techniques aim to model some aspects of spatial communication, they are also intended to be implemented in software so as to manage networked communication as participants move around within a virtual world. Thus, they are more concerned with providing expressive and flexible control over communications than they are with precisely modelling all aspects of human communication. One such model is the ‘spatial model of interaction’. The first basic version of this model considered communication between pairs of participants in a common virtual space [1]:

- A and B are inhabitants of a common virtual space
- They may communicate through multiple media (e.g., audio, text, 3D graphics).
- The communication that flows from B to A in a given medium, M, depends upon A’s awareness of B in M. The more aware A is of B the more information should flow from B to A.
- A’s awareness of B is a function of A’s focus on B and B’s nimbus on A. Put another way, A and B use A’s focus and B’s nimbus to mutually negotiate A’s awareness of B.
- Focus represents the ‘shape’ of A’s attention across the space. This may range from a simple discrete bounded space (B is either in A’s focus or not) or may be a more sophisticated spatial field.
- Nimbus represents the shape of B’s projection into the space, it is a mechanism that allows B to express where they do or do not want to project their information.
- By moving around the virtual space, A and B use their focus and nimbus to negotiate A’s awareness of B. For example, B may direct their nimbus at A in an attempt to interrupt them while A may direct their focus at B in order to better ‘overhear’ them. That this is a bi-partite negotiation is an essential feature of the model.
- The extent and shape of focus and nimbus also affect communication. A broad nimbus addresses many people while a narrow one may target a particular individual. Similarly, a broad focus attends to many while a narrow one is more selective.
- In turn, B’s awareness of A depends upon B’s focus on A and A’s focus on B. This need not be the same as A’s awareness of B, i.e.. awareness need not be symmetrical, modelling situations such as when I can see you but you cannot see me.
- Focus and nimbus and hence awareness are medium specific; they can have different shapes and extents in each medium of communication that is available. After all, I may be able to hear you when I cannot see you (if you are behind me) or see you when I cannot hear you (if you are a long way away from me). Figure 1 shows an example of focus and nimbus functions that were used to manage real-time video communication in MASSIVE, with awareness driving session management and also video quality of service (in this case frame rate).
- Focus, nimbus and awareness can be implemented in a system and used to control aspects of its operation such as session management, quality of service management, and security. An example is Greenhalgh’s implementation in the MASSIVE-1 collaborative virtual environment [X] where they were used to control graphics, textual and real-time audio communication.

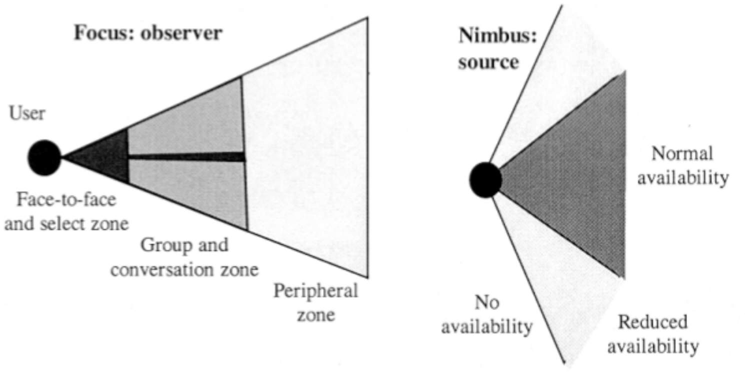

Figure 1: Example focus and nimbus functions from the MASSIVE system

In his PhD thesis [8], Greenhalgh extended the spatial model to account for the effects of independent (of the directly communicating parties) spatial structures and objects on awareness and communication, and also to support awareness among dynamic groupings and even crowds of participants (see also [2,7]). His key extension was the third-party object which was practically demonstrated in the MASSIVE-2 system:

- Third-party objects model the effects on mutual awareness of objects in a virtual space such as boundaries (e.g., walls, windows, one-may mirrors and others) and also communication tools (e.g., podia, microphones).
- A third party object can have two effects on awareness. It can directly adapt the awareness between two parties, through attenuation (e.g., a wall) or amplification (e.g., a podium). It can also aggregate information from multiple parties by mixing, summarising or selecting across multiple channels to create a new ‘group’ channel. These two effects can be combined, for example in a virtual room whose boundaries attenuate direct awareness between those who are inside and those who are outside, but that also replaces this with an abstract view of the activity that is taking place inside.
- These effects are triggered and controlled according to the awareness that the third party objects has of the communicating objects or vice versa.
- Third party-objects can themselves be mobile within the virtual space and may be dynamically created and destroyed. Greenhalgh demonstrated how the room example from the previous bullet could be extended into a mobile ‘crowd’ object that would be associated with a moving group of participants.

Although the initial impetus for the spatial model of interaction arose from the requirements of virtual Cartesian spaces such as those seen in the MASSIVE systems, the same concepts might be applied to other kinds of spaces including the everyday physical spaces of many ubiquitous computing systems, and potentially also to the spaces of bigraphs. Indeed, Rodden has previously presented an interpretation of the spatial model of interaction for graph- or network- spaces [14] rather than the continuous Cartesian spaces that were the basis of the initial implementations of the spatial model in the MASSIVE systems, an interpretation that would seem to be particularly relevant to biographical models of ubiquitous systems.

Hybrid spaces

One of the most distinctive features of Savannah is the manner in which it combines physical and virtual spaces, in this case appearing to overlay the latter on the former. There are a number of potential relationships between the physical and virtual that are worth considering here and that may ultimately require modelling using bigraphs.

The first possibility is to **overlay** the physical and virtual as in Savannah, so that movement through one is automatically related to movement through the other. The resulting experience may take several forms from a participant’s point of view as captured by Milgram and Kishino’s mixed reality continuum (Figure 2) [10]. At one end of this continuum lies everyday physical reality while at the other is an idealised form of virtual reality in which the participant is completely immersed in a computer generated virtual world. In between lie various hybrids: augmented reality, in which the participant primarily inhabits the physical world, but which appears to be overlaid with the virtual (Savannah lies here); and augmented virtuality, in which the participant inhabits an online virtual world, but which is made live with information from the physical world (e.g., with live weather, traffic flows or other data captured from sensors).

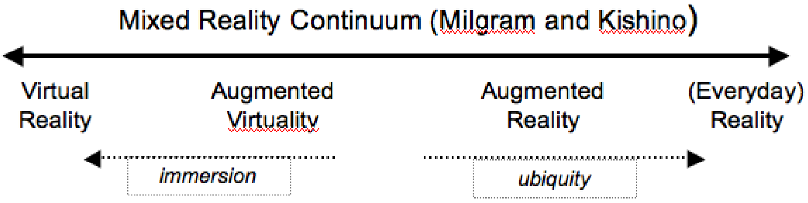

Figure 2: Milgram and Kishino’s Continuum

Some systems simultaneously occupy several points on this continuum. Can You See Me Now [3] was a game of chase in which online players in an online virtual world (a 3D graphical reconstruction of part of a city) were pursued by a second group of participants, the runners, who (equipped with handheld computers with GPS) had to run through the actual city streets in order to chase them. In this case, the online players inhabited an augmented virtuality, while the runners simultaneously inhabited and augmented reality. The same is true of Savannah, with the activity on the playing field being a form of augmented reality and the debriefing sessions back in the classroom in which the children replayed recordings of their movements and actions on a projected view of the virtual savannah being augmented virtuality, except here the same group of players moves between the different modes in sequence.

A second relationship between physical and virtual space is that of **adjacency**. Previous work on interfaces to virtual worlds has proposed the concept of a **mixed reality boundary** as a two way portal between the physical and virtual [4]. The basic idea is summarised in Figure 3. Participants in a physical space see a persistent projected view of a virtual space on the wall while participants in the virtual world see a persistent live video texture looking back into the physical world. The two are then aligned in such a way (and enhanced with an additional audio channel) to create the effect of a two way window so that each kind of participant can look from their space into the other. Multiple mixed reality boundaries can then be used to connect many physical and virtual spaces into a more complex hybrid structure, in much the same way that physical spaces are connected in a building. Research into mixed reality boundaries also considered a set of generic properties that such boundaries might possess including the ability to attenuate and amplify the information flowing across them, various degrees of symmetry and control by participants, and even the idea of traversability [9], in which participants experience the illusion of physically passing from one to the other.

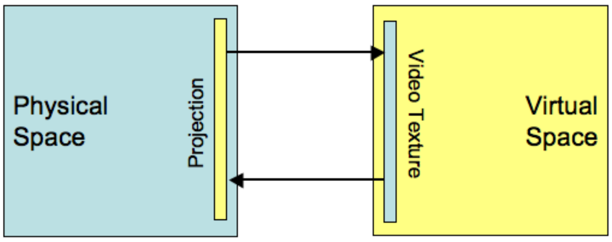

Figure 3: Schematic of a simple Mixed Reality Boundary

The third kind of relationship is that already explored in previous bigraph models, that of **containment**. Physical spaces can clearly contain other physical spaces and the same is true of virtual spaces. It is interesting to consider what it means for a virtual space to contain a physical space or vice versa.

In the general case, a ubiquitous computing system may involve many different physical and virtual spaces that are connected into a complex hybrid structure through relationships of containment, overlay and adjacency. We need to consider how each of the relationships can best be expressed using bigraphs.

Seams

A further issue raised by Milner when discussing Savannah is whether we should also model the properties of GPS, for example its variable accuracy and coverage as these appear to have had an impact on the players' experiences, contributing to the difficulty of successfully coordinating attacks on prey [5]. Other studies outside of Savannah have also described how the limited coverage and accuracy of positioning technologies and also the limited coverage of communication technologies had a profound impact on the end experience. For example, a study of the Can You See Me Now? game introduced previously revealed how the lack of coverage of both GPS and WiFi across a built up urban area had a major impact on the experience of online players and especially of the runners on the city streets (see Figure 4) [3].

These characteristics of the underlying ubiquitous infrastructure have come to be referred to as 'seams' in the infrastructure, and while they might ideally be invisible to users, they are often painfully apparent.

Seams are more prevalent than the occasional errors or bugs that plague many software systems; while bugs certainly have to be considered when designing reliable software, seams need to be taken on board as an ongoing fact of life – part of the fabric of ubiquitous computing. Recent work on authoring tools for location-based computing systems has proposed that the presence and nature of seams needs to be revealed to designers [13]. In effect, rather than being a two layer problem in which the virtual is overlaid on top of the physical, designers need to explicitly deal with the presence of a third ‘infrastructure’ layer (Figure 5). Mathew Chalmers and his team at Glasgow have gone a step further by creating a series of ‘seamful games’ in which seams become a resource to be used and discovered in a game, for example with players finding and then hiding in the ‘GPS shadows’ [6]. The main implication of such arguments for the formal modelling of ubiquitous computing systems is that seams – the operational characteristics of the infrastructure – will also often need to be explicitly modelled. The transparent or seamless operation of the technologies should probably not be taken for granted.

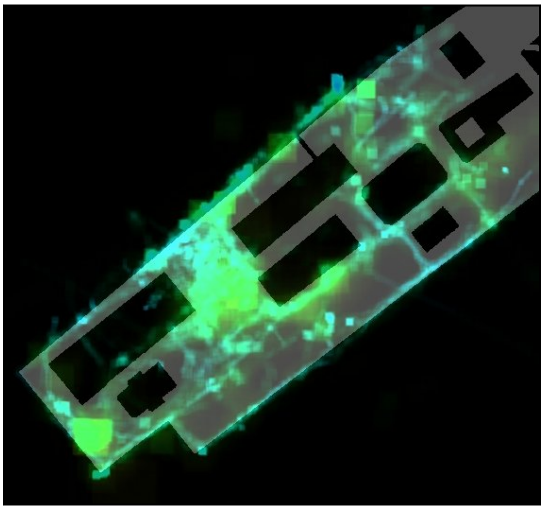

Figure 4: Visualisation of seams in Can You See Me Now? from [3] – an aerial image looking down onto a peninsular in Rotterdam overlaid with two hours worth of logged data; water and buildings shown in black; small blue squares show where high accuracy GPS was successfully transmitted back to the server; larger green squares for less accurate GPS; grey suggests no WiFi and/or no GPS coverage logged here throughout the two hours.

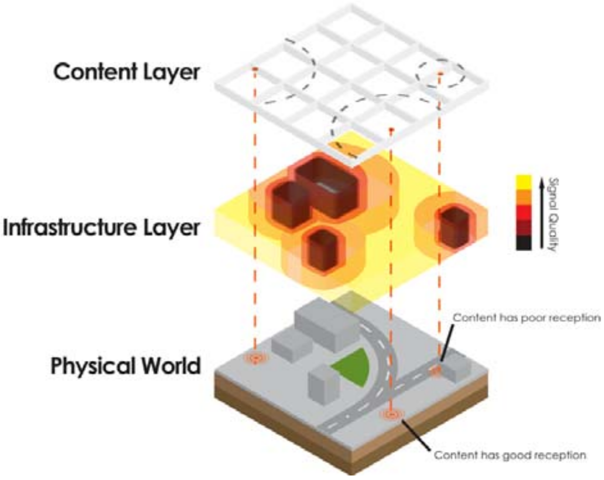

Figure 5: the spatial structure of a location-based application as a three layer structure from [13]

Conclusion

A review of previous work in virtual and mixed reality and ubiquitous computing has highlighted key aspects of physical and virtual space that may need to be modelled if we wish to create comprehensive formal descriptions of ubiquitous computing systems. These include:

- the spatial scoping and management of awareness and communication, as considered by the spatial model of interaction
- various possible relationships between physical and virtual spaces including overlaying, adjacency and containment
- the presence, nature and potential impact of seams in the underlying ubiquitous infrastructure

References

1. Benford, S. D. and Fahlén, L. E., "A Spatial Model of Interaction in Virtual Environments", *Proc. 1993 European Conference on Computer Supported Cooperative Work, ECSCW'93,* (edited by De Michelis, G., Simone, C and Schmidt, K.) 109-124, Milano, Italy, September 1993, Kluwer Academic Publishers.
2. Benford, S. D., Greenhalgh, C. M. and Lloyd, D, "Crowded Collaborative Virtual Environments", *Proc. CHI'97*, pp 59-66, Atlanta, Georgia, US, March 22-27, 1997, ACM Press.
3. Benford, S., Crabtree, A., Flintham, M., Drozd, A., Anastasi, R., Paxton, M., Nick Tandavanitj, N., Adams, M., Row-Farr, J., "Can You see me now?", *ACM Trans. Comput.-Hum. Interact*, 13(1), 100-133, 2006.
4. Benford, S., Greenhalgh, C., Reynard, G., Brown, C., Koleva, B., "Understanding and constructing shared spaces with mixed-reality boundaries", *Transactions on Computer-Human Interaction (TOCHI)*, Volume 5 Issue 3, September 1998, ACM.
5. Benford, S., Rowland, D., Flintham, M., Drozd, A., Hull, R., Reid, J., Morrison, J., Facer, K., *"Life on the edge: supporting collaboration in location-based experiences"*, Proceedings of the SIGCHI conference on Human factors in computing syste, CHI 2005, 721 - 730, Portland, Oregon, USA, ACM, 2005
6. Chalmers, M. and Galani, A. "Seamful Interweaving: Heterogeneity in the Theory and Design of Interactive Systems", *Proceedings of the 2004 ACM Symposium on Designing Interactive Systems*, August 2004, Cambridge, Massachusetts, ACM Press, 243-252.
7. Greenhalgh, C., and Benford, S., "Supporting Rich And Dynamic Communication In Large Scale Collaborative Virtual Environments", *Presence*, Vol. 8, No. 1, 14-35, MIT Press
8. Greenhalgh, C., *Large Scale Collaborative Virtual Environments*, BCS/CPHC Distinguished Dissertation Series, ISBN 1-85233-148-8, London: Springer-Verlag, 1999.
9. Koleva, B., Schnädelbach, H., Benford, S., Greenhalgh, C., "Traversable Interfaces Between Real and Virtual Worlds", *Proceedings of the SIGCHI conference on Human factors in computing systems, CHI 2000*, 233-240, The Hague, Netherlands, ACM, 2000
10. Milgram P., Kishino, F., 'A Taxonomy of Mixed Reality Visual Displays', IEICE Transactions on Information Systems, E77-D12, 449-455, 1994.
11. Milner, R., "Collaborative Behaviour in the Wild", February 2009
12. Milner, R., *The Space and Motion of Communicating Agents*, 1st Draft, December 2008
13. Oppermann, L., Broll, G., Capra, M., Benford, S., "Extending Authoring Tools for Location-Aware Applications with an Infrastructure Visualization Layer", Proceedings of Ubicomp 2006, San Jose, USA, September 2006, Springer Verlag
14. Rodden, T., "Populating the application: a model of awareness for cooperative applications", *Proceedings of the 1996 ACM conference on Computer Supported* Cooperative Work, CSCW 96, 87-96, Boston, Massachusetts, United States, ACM, 1996.