Where On-Line Meets On-The-Streets: Experiences With Mobile Mixed Reality Games

Martin Flintham, Rob Anastasi, Steve Benford, Terry Hemmings, Andy Crabtree, Chris Greenhalgh, Tom Rodden
The Mixed Reality Laboratory,
The University of Nottingham
Nottingham, NG8 1BB, UK
{mdf, rma, sdb, tah, axc, cmg, tar}@cs.nott.ac.uk

Nick Tandavanitj, Matt Adams, Ju Row-Farr
Blast Theory
Unit 43a Regent Studios
8 Andrews Road
London, E8 4QN
{nick, matt, ju}@blasttheory.co.uk

ABSTRACT
We describe two games in which online participants collaborated with mobile participants on the city streets. In the first, the players were online and professional performers were on the streets. The second reversed this relationship. Analysis of these experiences yields new insights into the nature of context. We show how context is more socially than technically constructed. We show how players exploited (and resolved conflicts between) multiple indications of context including GPS, GPS error, audio talk, ambient audio, timing, local knowledge and trust. We recommend not overly relying on GPS, extensively using audio, and extending interfaces to represent GPS error.

Categories and Subject Descriptors: H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems - Artificial, augmented, and virtual realities; H.5.2 [Information Interfaces and Presentation]: User interfaces; K.8 [Personal Computing]: Games.
General terms: Design, Human Factors
Keywords: Mobile & wireless games, context, GPS, audio

MOBILE MIXED-REALITY GAMES
We explore the design of experiences in which mobile participants collaborate with those who are online. Potential examples include remote maintenance, command and control, museum visiting and games. The key to such applications is establishing a relationship between humans who are operating across radically different contexts. The role of the system is to provide them with sufficient information, presented in an appropriate way, for them to be able to understand the context of the other. This focus on human-to-human understanding and use of context sets such applications apart from many current mobile applications such as tourist guides in which the system 'decides' to trigger information for users (see [1]).

CHI 2003, April 5–10, 2003, Ft. Lauderdale, Florida, USA.

Our particular focus is on 'mobile mixed reality games' in which online players play with or against those on the streets. Elements of such games can be seen in the growth of online games, wireless games and early demonstrations of augmented reality games. The latter include the use of handheld computers and augmented reality displays to play games while moving through buildings and the city streets [5,10,11], and also augmenting traditional physical games with new digital content [8].

In this paper, we compare two early experiences of mobile mixed reality games. In the first, members of the public who were online played against professional performers who used mobile technologies on the streets. The second reversed this relationship, putting the mobile technologies in the hands of the public.

These games have allowed us to explore the different forms of collaboration that can take place between those on the city streets and those online. They have also provided us with a way of investigating the deeper question of how humans use and interpret different contextual cues when collaborating in such a situation. Specifically, we explore:

- the impact of GPS inaccuracy on the use of *location* as key component of content and context information;
- the significance of other sources of information, especially real-time *audio*, in supporting and sometimes contradicting location information.
- the strategies employed by participants to use these sources of context to achieve successful collaboration;
- ways in which interface and game designers can respond to these issues and support these strategies.

CAN YOU SEE ME NOW?
Our first experience is a chase game called Can You See Me Now? (CYSMN) that was staged as a public event in Sheffield in the UK over a weekend in late 2001. CYSMN was designed to be a fast-paced game in which up to twenty online *players* (members of the public using the Internet) were chased across a map of the city, by three *runners* (professional performers) who were moving

through the actual city streets. The main objective of CYSMN was to engage and excite the online players by giving them a sense of the runners' experience of the city, and of how their online actions could affect events on the streets. Runners and players shared an online map. Runners had a global view showing all avatars, and players a local view showing avatars in their local vicinity. The runners' positions were determined using GPS. The players' positions were controlled through their online interfaces. Runners and players avatars were marked out through the use of different colors. The system determined that a runner had caught a player if their avatar moved to within five virtual meters of the player's avatar on this map.

In order to make the players more aware of the runners' experience of the city streets, we provided a real-time audio channel from the runners to the players. The runners talked to one another over a shared radio channel using walkie-talkies. This talk was then digitally encoded and streamed to the players over the Internet. In return, the players talked both to one another and to the runners using a text channel. We choose to use text because the current Internet cannot reliably support the transmission of many audio streams among many parties, especially over dial-up connections. The game took place over a central area of the city that was roughly half a mile square and that consisted of a mixture of open spaces and narrow streets lined with tall buildings.

The player interface

A player's experience began at the CYSMN webpage where they entered a name for themselves and then joined a game queue (we restricted access to a maximum of twenty players at a time). When it was their turn to play, they were dropped onto a 2D schematic map of the city. They could use the arrow keys on their keyboard to move across this map, but were unable to enter solid buildings. Figure 1 shows an example of the player interface. A player was represented as a pair of icons on the map. A white icon showed their current position according to their local client, providing immediate feedback as to their movement whenever they pressed a key. A blue icon showed their position according to the game server. This would trail behind the white icon with a lag of about one second due to the communication delay between client and server and the time taken to process players' movements at the server. Other players were represented as blue icons and runners as orange icons. The players continued to move and text until a runner got sufficiently close to them that they were caught. At this point they were removed from the game and offered a chance to re-enter the queue.

The runner interface

The runners' interface was delivered to them on a Compaq iPAQ handheld computer from a server in a nearby building over a 802.11b wireless local area network. A GPS receiver plugged into the serial port of the iPAQ registered the runner's position as they moved through the streets and this was sent back to the server over the wireless network. The iPAQ and GPS receiver combination was attached to a weatherproof board to improve ruggedness, and ease of carrying. Given the small screen size of the iPAQ, the runners' map allowed them to zoom between a global view and a close-up local view centered on their current position. The runners used walkie-talkies with earpieces and a head-mounted microphone. Finally, they carried digital cameras so that they could take a picture of the physical location where each player was caught. These pictures appeared on an archive web site after the event [2]. Figure 2 shows one of the runners equipped and ready to go. Figure 3 shows the equipment that they carried. Figure 4 shows an example of the interface from the iPAQ. A summary video of CYSMN is available from www.equator.ac.uk.

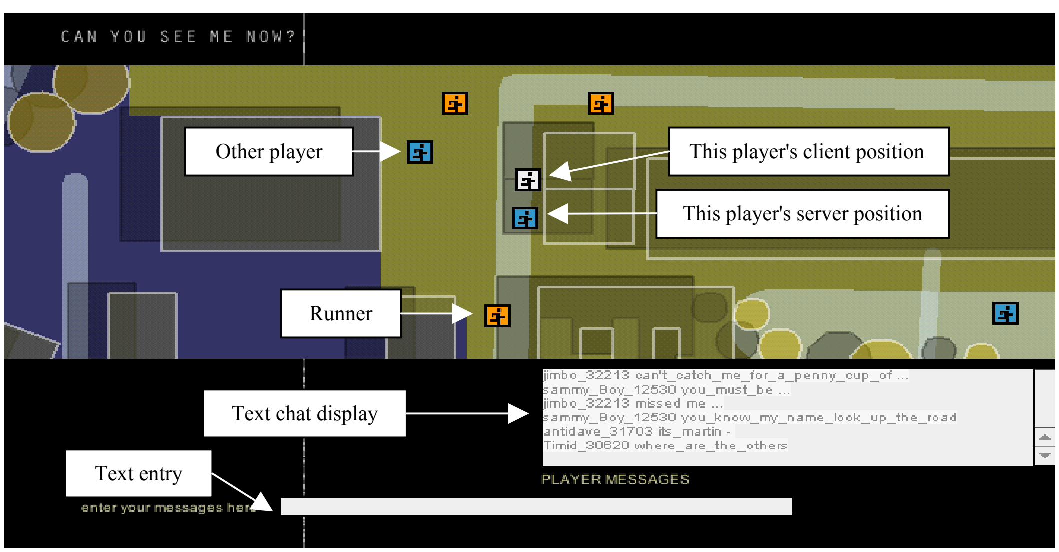

Figure 1: The online player interface from Can You See Me Now

EVALUATING CAN YOU SEE ME NOW?

Can You See Me Now? was live for six and a half hours during the weekend of Friday 30th November and Saturday 1st December 2001. 214 players took part over the Internet. Of these, 135 of these were caught by the runners, 76 logged off and 3 were never caught. The best 'score' (time without being caught) was 50 minutes. The worst was 13 seconds. The game was staged alongside a new media festival that was taking place in the same area of the city, and several of its delegates logged on from public terminals. Others came in over the wider Internet, including several over intercontinental connections.

Our analysis of CYSMN draws on three sources. First was offline feedback from players via email and face-to-face conversations (for those known to us) and also debriefing meetings with the project team. Second was ethnographic observation (utilising video and field notes) of the activities of the runners, players and behind-the-scenes production crew. Third was statistical and manual analysis of system logs of GPS positions, GPS errors and text messages so as to reveal broad patterns of activity.

Our overall impression was that CYSMN was broadly successful. Player feedback was largely positive and there were some moments where we seemed to have created genuine tension and excitement for players, offering glimpses of the potential of this game format. As one player put it in a subsequent email:

> "I only managed to get on to the map once for about 15 minutes. I can't remember the name I used, but it was pretty un-nerving first hearing my name said"

The following analysis focuses on how players and runners were able to successfully collaborate with one another to play the game, and in particular, how they exploited multiple channels of information.

Figure 2: a runner ready to go

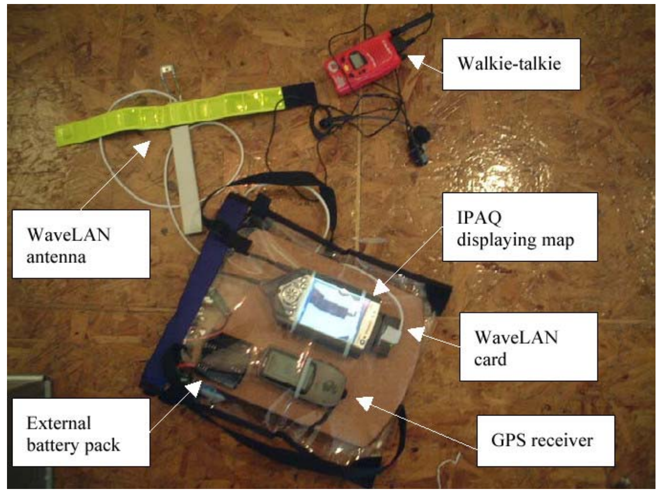

Figure 3: a runner's equipment

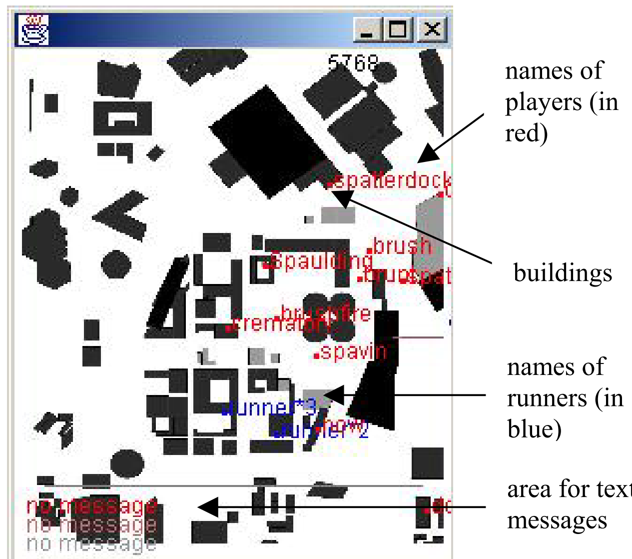

Figure 4: the runner's interface on the iPAQ

Exploiting GPS inaccuracy

Location, as sensed by GPS, was part of the core content of CYSMN. Location was used by the system to determine when key events had occurred (capturing players). It was also displayed directly to all participants on the shared map.

Previous attempts to take augmented reality outdoors have highlighted GPS inaccuracy as key issue [3]. While GPS is a versatile positioning technology for outdoors applications, it can also be a problematic one, particularly with regard to inaccuracy that can vary with location on the Earth's surface, time of day, proximity to buildings, and weather. Accuracy ranges from a few centimeters up to tens of meters, and is often worse in urban environments and when using budget GPS receivers (both factors in CYSMN). We anticipated that GPS would be a significant factor.

Analysis of system logs shows that, from a technical perspective, GPS was indeed quite inaccurate. Estimated errors ranged from 4m to 106m with a mean of 12.4m and a standard deviation of 5.8m. Error varied according to location in the game area, with some of the more open spaces exhibiting typically only a few meters error while the more narrow built up streets suffered considerably more. The more extreme errors were most likely due to multi-path reflections or temporary losses of satellite visibility.

These errors would have had a variety of effects. They would have shown runners being in different locations on the map compared to their actual physical locations on the streets. They would also have resulted in the runners' avatars making sudden unfeasible jumps across the map. Indeed, when we subsequently plotted the runners' positions as reported by GPS (i.e., from the system's point of view rather than their actual physical locations) alongside the locations where players were captured, we could see several examples of captures that seemed to have occurred as a result of sudden and unlikely jumps in GPS positions. Such captures could be considered to be 'unfair'.

We now consider how these inaccuracies were actually experienced by the players. It was interesting to note that the runners on the street and online players experienced GPS inaccuracy differently. The runners actually exploited a growing knowledge of GPS inaccuracy in relation to the geography of the game as part of their gameplay. Without this knowledge, the online players seemed to be much less aware of its presence and possible effects.

GPS inaccuracy experienced in the street

Over the course of the two days play the runners became increasingly aware of the effects of GPS inaccuracy and also where on the city streets it was most likely to be experienced. By the second day's play, they had begun to exploit this knowledge as part of their tactics, as shown by the following conversation between a runner and one of the development team:

Crew: So your tactics: slow down, reel them in, and get them?
Runner: If they're in a place that I know it's really hard to catch them, I walk around a little bit and wait till they're heading somewhere where I can catch them.
Crew: Ambush!
Runner: Yeah, ambush.
Crew: What defines a good place to catch them?
Runner: A big open space, with good GPS coverage, where you can get quick update because then every move you make is updated when you're heading towards them; because one of the problems is if you're running towards them and you're in a place where it slowly updates, you jump past them, and that's really frustrating. So you've got to worry about the GPS as much as catching them.

The development team had felt after the first day that the players were too quick for the runners and had debated changing the game server to slow them down (although this turned out to be impossible to do overnight). However, this new tactic of exploiting GPS inaccuracy, combined with a more collaborative approach by the runners, significantly changed the balance of game, so that by the second day the problem had disappeared.

GPS inaccuracy experienced by online players

It is less clear what effect GPS inaccuracy had on the players. Although runners did make occasional references to GPS over the audio channel (e.g. "I'm waiting for satellites"), it was not perceived to have been a major factor influencing the players. They didn't mention sudden jumps, unfair captures, or indeed GPS in their text messages or in feedback after the event. We suggest several possible reasons for this:

- Players could see the reported positions of the runners on the on-line map, but could not compare these with observations of where they were in the physical world. It was therefore difficult for them to spot positioning errors. Runners on the other hand, would be able to directly compare reported and actual positions.
- Players may have been unfamiliar with GPS and the nature and causes of its inaccuracies.
- Even if familiar with GPS, they would probably have been unfamiliar with the local terrain, especially the locations of tall buildings and narrow streets, where GPS might be expected to be perform badly.
- They weren't exposed to the game for as long as the runners and so didn't have time to build up a working knowledge of GPS problems.
- It would have been hard to distinguish jumps caused by GPS error from those caused by variable delay over Internet connections (which was probably a more familiar problem for the players).

As we have seen, this ended up placing online players at a disadvantage as runners in the street exploited this difference. This raises the question as to whether the technology should have made them more aware of the characteristics of GPS? However, the shared map was only one source of useful information. Another was audio.

Revealing context through talk

The audio stream from the runners to the players proved to be highly significant. First, as has been reported from previous studies of online experiences [4], the audio channel formed a core part of the content. In particular, audio was capable of heightening the tension for the online players (see the earlier quote for an example).

In addition, audio was a major source of context information. The runners frequently and deliberately revealed information about the local environment, the status of the technology, and their own physical status through their talk. Phrases such as "I'm waiting for a Green Man" (meaning I'm waiting at a pedestrian crossing) revealed the presence of traffic on the roads at key locations. Mentions of hills and other features provided knowledge of local topography. There were also references to batteries being low and other technical problems. The background audio provided further contextual cues such as the heavy breathing of the runners when running up hills and the presence or absence of traffic noise. Feedback from players suggests that some were tuned in to this information and were able to exploit it tactically. As one player put it:

> "I figured out pretty quickly what was uphill and downhill. I also figured out which was the main road to cross"

More generally, this use of audio illustrated a number of interesting features of context.

The highly localized nature of contextual information

Audio references to particular places highlighted the very situated nature of the interactions in the street. The runners' talk often contained references that were based upon their local knowledge of the area and that the players could not be expected to know. For example, requests to meet other runners at the 'Glowing Mushrooms', a colloquial name for a group of local buildings, would have made little sense to the players as this label did not appear on the shared map (and probably not on any map of the area).

Location context therefore was determined by the knowledge of the runners on the street built up over the two days of the event. This provides our first example of the potential problems involved in resolving different channels of information, in this case the map and the audio – a theme that we return to later on.

The richness of talk in developing context

In spite of such difficulties, audio was clearly a rich source of context, especially when compared to the rather minimal schematic map with its (often inaccurate) representation of locations. This, of course, is not at all surprising. After all, talk furnishes the practical context and content of action in a wide variety of circumstances and settings [7].

It is interesting, however, to speculate on the potential of different media to establish a shared context between on-line players and those on the street. Perhaps we could have used real time video from the players to paint an even richer picture of events on the streets. However, this might also have also revealed the inaccuracies of the GPS much more clearly (further heightening the conflict with the map?). Perhaps audio is appropriate here because it is rich but also positionally imprecise, somewhat vague and open to interpretation, and because participants can choose how much information to reveal?

Given these initial observations on the nature and use of location and audio information in supporting collaboration between online players and mobile performers, we now turn to a second example, a game called Bystander, that reverses the arrangement, putting the public players on the street and the performers on line.

BYSTANDER

In Bystander, a *local player* takes a journey through the city on the trail of a mysterious person whose name and picture they have only briefly been shown. A partner, an *online performer*, collaborates with them and guides them in the search. Between them the two participants simultaneously travel through the city streets and across an online map in search of the mysterious target person.

The role of the online performer is to steer the local player through a series of key physical locations. At each, they manually trigger the release of a clue that hints at where they should go next. Each clue takes the form of a short video clip that suggests a route or destination, and that is viewed on the iPAQ. Eventually, the local player encounters a ringing payphone; they answer it, to hear a message from the person that they have been following. They can then leave their own final message.

As a follow on to CYSMN, Bystander is interesting for several reasons:

- The mobile technology is now in the hands of a public player who is out on the city streets. They are now at the 'sharp end' of the technology and have access to the broader setting of the city while the performer driving the experience does so from a distant and relatively isolated on-line interface.
- The key action takes place in the physical world. Success depends upon finding and viewing video clips at actual physical locations. Furthermore, the clips are designed to make sense when viewed from precise physical locations (i.e., within just a couple of meters of a specific vantage point).
- The game is much less transient than CYSMN. It engages the local player and the online performer in a longer-term relationship (roughly forty minutes) that passes through several critical moments.

The technologies for Bystander are based on those from CYSMN. The local player uses a similar mobile interface and both players see a shared online map where the positions of avatars are automatically generated from sensor data (GPS). Audio is streamed from the local player to the online performer (this time from the iPAQ's onboard microphone) and text messages pass back the other way. The online performer now has a zoomable 3D view of the map. Within this, their viewpoint is tethered to the position of the local player (as reported by GPS), so that they can only see the area directly around their reported position. The online performer is therefore not omniscient and they have to collaborate with the local player to carry out the game. Figure 5 (right) shows the online player's interface, including the (restricted) view of the 3D city model, buttons for zooming this in and out, and an area for composing and viewing text messages. Figure 5 (left) gives an overview of the 3D city model that they explore.

EVALUATING BYSTANDER

At the time of writing, Bystander has still to be presented to the general public. So far, we have conducted a series of initial tests consisting of six trial runs. Our local players have been drawn from the development team and associates, but have included inexperienced players with little knowledge of the game or technologies. All, however, had considerable local knowledge that is important to the provision of contextual information. Even with these relatively knowledgeable players, testing reveals some interesting lessons for context. The following analysis revisits our discussions of location and audio.

GPS inaccuracy revisited

In CYSMN we saw the way in which GPS errors were largely invisible to the online players and even exploited by the runners. In contrast, GPS errors in Bystander become much more noticeable and problematic for both online and mobile participants.

Local players repeatedly commented on the inaccuracy of the map, for example, reporting that it had placed them on the wrong street. In contrast to CYSMN players errors became much more obvious because they were now directly aware of their location in the physical world and could directly contrast this with that shown by the map. And yet, these errors were not obviously exploitable. Rather they were taken as evidence of the system's failure.

In early tests, the online performers appeared to largely ignore the issue of GPS inaccuracy (as had the players in CYSMN), even though they were familiar with the issues in principle. However, this also turned out to be much more problematic. Performers appeared to place too much trust in the map, responding to the position and orientation of the local player's avatar in great detail as if it were accurate. Because performers were seeking to cause effects in the real world the resulting misunderstandings also became much more problematic for the on-line participant. For example, instructions from the online performers frequently included detailed spatial references such as "go west" or "turn left and go up the road". GPS errors made such instructions practically meaningless (particularly when combined with network delays that meant that the local player could easily have moved by the time the received the instruction).

Because they were no longer on the streets and so able to directly see the consequence of GPS errors, *performers were now denied the information about inaccuracy of the GPS system that they had so successfully used as a resource within CYSMN*. This version of the design hid one of the main resources used to build the context of the game by the performers.

Presenting inaccuracy as a resource

In response to this observation, we designed new representations of the local player in the online performer's interface to explicitly convey the presence, amount and history of GPS error.

Figure 6 shows two early designs. The three images towards the left show (at different levels of detail) a design in which the location of the online player is represented by three components. First is a cross that marks the point of the most recent GPS update. Second is the walking avatar figure that is initially placed at the first GPS update and then continually walks towards the most recent update. Third is a series of octagons that show the positions of the last few GPS updates and whose size represents the GPS error and that fade over time. This representation seemed to greatly help the online performers in later tests (see below).

However, we felt that a more direct representation would be required for future online players. The right hand image of figure 6 (over) shows a simpler alternative. Here the local player is represented as a dark, semi-transparent blob, the size of which increases with GPS error and whose outline continually shifts. The aim is to convey the message: "the system thinks that the online player might be somewhere around here". It also hides the area of the map under the representation, making it difficult for the viewer to give precise directions that refer to the area of uncertainty.

Audio and context revisited

Given the difficulties associated with location information, audio took on an even more significant role. In later tests, the online performers moved away from giving precise spatial instructions and instead exploited the audio channel to gain the contextual information needed to make the game work. Their strategies included encouraging local players to describe what they could see and asking the local players to slow down or even stop so that the online performer could get a fix on their location. Although this mitigated problems with GPS error, the local players voiced some frustration at being made to wait by their online partner.

These two tactics are interesting in that the on-line player is exploiting the audio communication to address problems with the location information. Moreover, by slowing them down, they are also manipulating the local player's interaction with the underlying GPS technology.

Acting on contextual information.

Another difference between CYSMN and Bystander is that the performer rather than the system reacts to the contextual information by triggering events (downloading a new video clip to the local player). In a sense, the performer was playing the role often played by the system (e.g., in context-aware mobile tourist guides).

It is useful to see the strategies that the performers adopted. A key strategy involved comparing the audio information with the location information and in particular, hoping for a stable moment where they seemed to agree with each other. The most experienced online performer described his strategy for triggering a video clips as follows:

> " …try to get [the player] right on top of them. Wait for a small GPS error and audio confirmation."

However, such moments did not always occur and so the performers would also have to resolve conflicts between

Figure 5: the 3D city model (left) and online player's interface (right) for Bystander

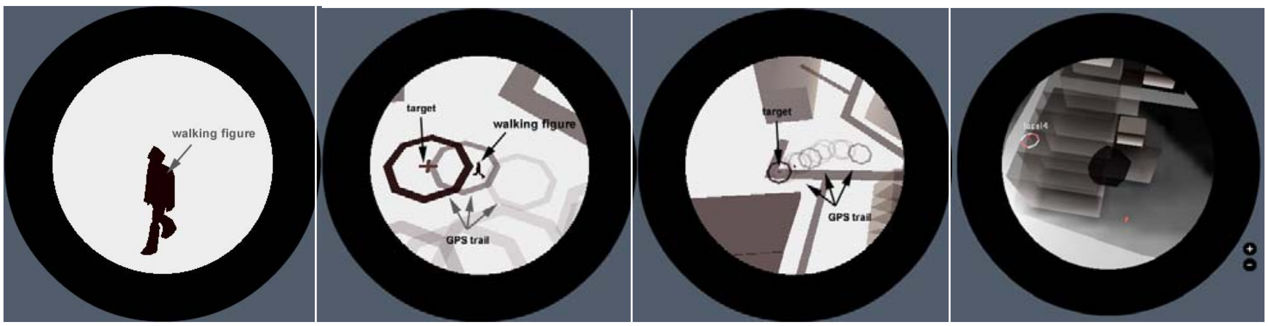

Figure 6: Interface designs for representing GPS uncertainty in Bystander

these different sources. Another described why she had chosen to trigger a particular video clip as:

> "He said he was at the bus stop [the target location] even though the map didn't show that. I trusted him"

Furthermore, these were not the only contextual factors to be taken into account and resolved. The pressure of time was also a factor. The same performer accounted for triggering a different clip as follows:

> "I felt he was wandering around. He was going to be there for a long time. I was being a bit sloppy with the rules."

In short, understanding the context of the local player sufficiently to decide to trigger a video clip involved the online player in comparing and resolving multiple contextual cues including GPS, GPS error, talk over audio and time.. Triggering a clip at the right moment is one thing, but receiving it is another. The responses of the local players over the audio channel (confirmed in subsequent discussions) suggested that that they had often moved on to a new location before they actually looked at a clip. So after all that effort, there was still no guarantee that a clip would be seen in context.

DISCUSSION

Both CYSMN and Bystander tell us something about the nature of context as it applies to mobile wireless systems, especially those in which online participants collaborate with those on the streets. We begin with three broad observations concerning the general nature of context.

Context is socially constructed.

First, in these experiences, context is more socially than technically constructed. The essence of both of them is whether the online player can construct an appropriate understanding of the mobile player's experience of the city streets. It was difficult to see any boundary between the sensing of context and the activities of those in involved. In fact as Dourish highlights [6]:

> "context and content (or activity) cannot be separated. Context cannot be a stable, external description of the setting in which activity arises. Instead it arises from and is sustained by the activity itself. "

This indeed is the case with our experiences. The system provides channels for context information between the participants, but makes little use of it directly. What is important however, is providing the correct channels.

Context judgments draw upon many resources in tandem

Context-aware judgments are complex, with human participants relying on a combination of location, errors in location, talk, ambient audio, timing, local knowledge and even trust in making their decisions. In CYSMN and Bystander, this information comes through two key channels: sensed location information shown on the online map, and a real-time audio stream.

Context judgments may involve conflicting information

The different sources of information often conflicted. A major task for the participants was to resolve these conflicts. Indeed, this is potentially the key challenge and therefore source of pleasure and engagement for both experiences. In some cases, resolving these conflicts involved manipulating participants' interactions with the underlying sensing technology.

Beyond these general observations on context, our experiences raise some more specific points in relation to the recommended use of location and audio information.

GPS and Location

GPS can be problematic in urban environments (due to limited visibility of the sky and the likelihood of multi-path errors). This was an issue for both CYSMN and Bystander, although with different consequences. It is worth stressing that the resolution of GPS is typically much higher than its precision. However, as we saw in Bystander, there is a danger that users will relate to it as a precise position (resolution) rather than as a likely region. In addition, estimates of GPS error are an extremely important piece of contextual information. The runners in CYSMN were able to exploit knowledge of GPS error that had been built up over a lengthy time. Deliberately representing GPS error in the interface was a useful strategy in Bystander.

It is also worth stressing that the significance of GPS errors can vary remarkably with task and geographical context. Some video clips in Bystander were only interpretable within a specific area of a couple of meters. On the other hand, in CYSMN a runner pursuing a player across an open space was effectively only using relative positional changes (moving the virtual cursor) – the lack of physical features and constraints permitted the decoupling of absolute physical and virtual position.

Audio and Talk

The mobile user's talk gave many clues to their context, especially in relation to named streets, shops, and other buildings of a describable appearance. Dialog was used extensively to elicit explicit contextual information and to test hypotheses about the mobile user's position, orientation, understanding and status. The audio channel was also a rich provider of less direct contextual information, e.g. whether the mobile user was out of breath, their apparent emotional state, and background environmental noise such as traffic. The interpretation and construction of speech and text messages drew heavily on the local geographical knowledge of the online user. Readily nameable elements of the physical environment were not captured by the system or (generally) represented in the interface, however they were extensively used to establish absolute positional information and directions.

Design challenges

We finish with key challenges and recommendations for the developers of future mobile mixed reality games.

- Don't be over-reliant on GPS, especially in urban environments. Don't mistake resolution for precision.
- Design interfaces that explicitly communicate the presence and nature of GPS error (see figure 6) and that encourage participants to see location as just another (possibly unreliable) source of information.
- Seriously consider using real-time audio as a rich but not overly precise source of context.
- Consider ways to exploit richer forms of contextual information, especially temporal characteristics of the experience and participants' local knowledge.

As a final note, in spite of the many difficulties that we have encountered, both CYSMN and Bystander have offered glimpses of how mixing online with on-the-streets can create compelling experiences. Our own plans involve carrying these issues forward in the development of a full public experience for 2003.

ACKNOWLEDGEMENTS

This work has been supported by the EPSRC through the Equator IRC (grant GR/N15986/01, www.equator.ac.uk). Additional support for Can You See me Now? has been provided by the Arts Council of England, BBC Online and b.tv. Additional support for Bystander has been provided by the Arts and Humanities Research Board (AHRB).

REFERENCES

1. Akoi, P. & Woodruff, A., Improving Electronic Guidebook Interfaces Using a Task-Oriented Design Approach, *Proc. DIS 00*, 319-325, Brooklyn, NY, ACM
2. www.canyouseemenow.co.uk - verified 1st Jan 2003
3. Azuma, R., The Challenge of Making Augmented Reality Work Outdoors, In *Mixed Reality: Merging Real and VirtualWorlds*, 1999, Springer-Varlag.
4. Benford, S., Greenhalgh, C., *et al*., Inhabited Television: broadcasting interaction from within collaborative virtual environments, *ACM ToCHI*, Dec 2000, ACM
5. Bjork, S., Falk, J., Hansson, R., Ljungstrand, P. (2001). "Pirates! Using the Physical World as a Game Board". *Interact 2001.*
6. Dourish, P., What is we talk about when we talk about context, Available from http://www.ics.uci.edu/~jpd/
7. Garfinkel, H., Studies in Ethnomethodology, Englewood Cliffs, Prentice-Hall.
8. Ishii, H., Wisneski, C., Orbanes, J., Chun, B., & Paradiso, J. (1999). PingPongPlus: Design of an Athletic-Tangible Interface for Computer-Supported Cooperative Play. *CHI '99*, 394-401, Pittsburgh.
9. Salber, D., Dey, A. & Abowd, G., The Context Toolkit: Aiding the Development of Context-Enabled Applications, *CHI'99*, 434-441, Pittsburgh.
10. Starner, T., Leibe, B., Singletary, B., & Pair, J (2000b), "MIND-WARPING: Towards Creating a Compelling Collaborative Augmented Reality Game", *Intelligent User Interfaces (IUI) 2000*, pp. 256-259.
11. Thomas, B.H., Close, B., Donoghue, J., Squires, J., De Bondi, P., and Piekarski. W., ARQuake: A First Person Indoor/Outdoor Augmented Reality Application, Journal of *Personal and Ubiquitous Computing*