Life on the Edge: Supporting Collaboration in Location-Based Experiences

Steve Benford, Duncan Rowland, Martin Flintham, Adam Drozd
Mixed Reality Laboratory
University of Nottingham,
Nottingham, NG8 1BB, UK.
{sdb, dar, mdf, asd}@cs.nott.ac.uk

Richard Hull, Josephine Reid
Hewlett-Packard Laboratories
Bristol, BS34 8QZ
{richard.hull,
josephine.reid}@ hp.com

Jo Morrison, Keri Facer
NESTA Futurelab
1 Canons Road
Bristol, BS1 5UH
{jo.morrison, keri.facer}
@nestafuturelab.org

ABSTRACT
We study a collaborative location-based game in which groups of 'lions' hunt together on a virtual savannah that is overlaid on an open playing field. The game implements a straight-forward approach to location-based triggering in which players must be in the same spatial locale in order to share information and act together. Comparison of video recordings of physical play with system recordings of game events reveals subtle and complex interactions between highly dynamic player behavior and the underlying technology. While players exhibit a fluid approach to group formation, the system embodies a more rigid view, leading to difficulties with sharing context and coordinating actions, most notably when groups of players span virtual locale boundaries or initiate actions while on the move. We propose techniques for extending locales to support more flexible grouping and also discuss the broader implications of our findings for location-based applications in general.

ACM Classification: H1.2 User/Machine Systems; H5 Information Interfaces and Presentation; C2.4 Distributed Systems;

Keywords: Location-based, Shared-context, Mobile, Ubiquitous, Collaborative, Education, Games, Study.

INTRODUCTION
Location-based applications exploit positional information to support mobile interactivity in domains as diverse as tourism [1,8], information retrieval [7], resource discovery [10], workplace awareness [13, 14], performances [5,9] and games [6,11]. An important focus for HCI research in this area has been the user's experience of positioning technologies, especially the impact of their inherent uncertainties. Positioning technologies can suffer from various limitations, the most significant of which concern availability, resolution, latency and accuracy. As an example, GPS, one of the most widely used outdoor positioning technologies, is notorious for black-spots where a receiver cannot see sufficient satellites to obtain a reading (e.g., in built-up urban areas), offers a typical resolution of a meter or less, a latency of a second or so, and a typical inaccuracy of several meters or sometimes many more.

Recent studies have revealed how these uncertainties can affect a user's experience. Users may become aware that reported positions do not correspond to actual positions, may observe apparently quirky behaviors such as when representations of others jump about, may be unaware of whether they are successfully being tracked or may have to build up a detailed working knowledge of the technology over many hours in order to successfully interact [9]. More generally, this relates to an ongoing concern in HCI as to how users can meaningfully interact with invisible sensing systems where they may be unaware of how they are being tracked, may be unsure how to address the system, and may not know how to avoid and rectify mistakes [3].

This paper presents a further study of how users experience location-based technologies, this time focusing on collaborative experiences in which several co-located participants have to coordinate their movements in a fine grained way. While the now familiar limitations of positional technologies do feature to some extent, we shall see that a focus on co-located collaboration brings new issues to the fore. In particular, complex interactions between player behavior and the underlying technologies inhibit players from establishing a shared context, even though they are physically co-located, leading to difficulties in coordinating their actions.

AN OVERVIEW OF SAVANNAH
We have developed and studied a collaborative location-based educational game called Savannah in which children learn about the ecology of the African savannah, especially about lion behavior. Groups of six children at a time role play being lions by exploring a virtual savannah that appears to be overlaid on an empty school playing field, an open grassy area of roughly ninety by sixty meters. Equipped with iPAQ PDAs with WiFi and GPS capabilities the children move around the playing field, exploring the varied terrain of the

CHI 2005, April 2–7, 2005, Portland, Oregon, USA.

savannah and discovering the resources that lions require to survive.

Our primary reasons for designing a game to be played on a featureless open space, as opposed to one where digital content is attached to the features of a specific environment, were: portability, the desire for the game to be easily deployable on any playing field with minimal effort; and configurability, the ability to radically change the locations of virtual content throughout the experience. We also anticipated that both GPS and WiFi would perform relatively well in such an environment.

The rules of the game

The game is divided into three levels, each of which last for approximately eight minutes. In the first level, the pride has to explore and scent mark its territory. In the second and third levels the pride hunts for food, first in the bountiful wet season and then in the harsher dry season. Each individual needs to maximize their health, with health points being gained for eating and drinking and lost for being starving, dehydrated or attacked by other animals. It is also possible to die, in which case the dead lion has to return to the physical starting point and suffer several minutes time penalty before being resurrected and sent back into the savannah. Points are awarded according to each lion's health at the end of each level.

The game requires collaboration; just as a pride of lions has to work together to survive, so players have to work together to win the game by deciding which animals to attack and how many lions must attack together in order to succeed. For example, a lone lion can bring down a stray wildebeest calf, whereas several lions are needed to bring down a mature buffalo, and even all six lions will not bring down a full-grown elephant. Unsuccessful attacks drain health, so that careful decisions are required as to the costs and benefits of a given number of lions attacking.

Inspired by previous location-based learning projects [15], we included a Den, an area of the classroom where children are briefed about missions, access research materials (books, videos and websites), plan their strategy and can reflect on how they performed in previous levels. Reflection is supported by an interface for replaying a system-recording of each completed level. This offers a birds' eye view of the virtual savannah across which recorded movements and actions can be replayed, scrolling forwards and backwards through time, and is displayed on an interactive whiteboard. An extended version of this interface has supported our own analysis of the experience and is described more fully later in the paper.

Realization of location-based triggering

Our realization of location-based triggering is typical of the straight-forward approach adopted by many location-based applications, in that players enter discretely bounded spatial locales in order to trigger events. Specifically, players move around the field with GPS reporting their positions which are mapped onto their positions on the virtual savannah. Whenever a player enters a new locale they get a new combination of sound and image and possibly an action button if it contains an active target. Importantly, groups of players have to be in the same locale to share common information and so be able to act together.

Each player's GPS unit is mounted at the top of a backpack just below the player's neck and is used to track their position as they move around the game space.

Figure 1: The game interface

Every second, a player's PDA transmits her or his location to a central game server via WiFi. The server replies by triggering the display of sound (through headphones), static images and status information on the PDA, revealing the virtual savannah and its inhabitants. A player invokes an action such as an attack by touching a soft button on the PDA's touch-screen (figure 1, inset), where the appropriate action is determined by the player's current location and game level. A notification of this action is sent to the game server over WiFi which responds with a notification of success or failure. The PDA disables further user actions while waiting for this response from the game server.

Locale-maps with foreground and background content

The game server implements the game rules. Each level is defined by its own locale-map and corresponding lookup table. A locale-map is created by coloring in different locales within a rectangle that represents the football field (see figure 2). Defining locales by coloring enables them to be given different shapes and sizes so as to represent different terrains, and provides designers with a quick and convenient way of authoring. The nature, locations and size of locales varied between levels. For each color used in the locale-map (corresponding to one or more locales), an associated lookup table specifies the following:

- A background image and sound to be displayed whenever this locale is empty, i.e. contains no foreground content. Examples of backgrounds include 'marsh', 'woodland' and 'raging river'.
- Details of the foreground content (e.g., an animal) that can appear in this locale including: its image and sound, an associated action that appears on the soft button on the

PDA interface (e.g., 'attack' and 'drink'), and the number of minutes into the game after which this content will appear for the first time.

Common to all levels is a further lookup table that specifies the characteristics of the animals that can appear as foreground content including how many lions are required to successfully attack them, how much health is gained by eating them, how much health is lost by unsuccessfully attacking them, and how many minutes elapse before the animal re-spawns after being eaten.

The locale map for level 3 can be seen inset in figure 2. As with the maps for levels 1and 2, we defined border locales around each map containing hazards such as fire, other lions and Masai villages in order to prevent the players from straying out of the game area. Inner locales varied greatly in terms of shape and size, from large background locales such as 'long grass' that spanned approximately thirty by twenty five meters down to smaller locales of roughly four by three meters containing specific animals.

Synchronization

There can only be one ongoing attack in a given locale at any time. When a player attacks they therefore either initiate a new attack or join in an ongoing one. When a new attack is initiated, other players have ten seconds to join in before the attack is resolved. Whenever they are attacking, a player's PDA shows a special attack image (a leaping lion) and does not update other sounds and images even if they subsequently wander out of the locale and into a new one. This combination of a ten second attack window and freezing the display during an attack is intended to provide an opportunity for players to be able to successfully coordinate their actions. Once an attack is resolved, all of the players involved receive text messages notifying them of success or failure at approximately the same time. They then see any changes to their health, hunger and/or thirst. Finally, the prey (foreground content) disappears to be replaced by the background images and sounds until such times as it re-spawns.

STUDYING PLAYER BEHAVIOR

We studied the game being played by groups of school children throughout three days of trials. Each trial session lasted for half a day, which provided sufficient time for each group to play all three levels, including repeating one level. The ages of the children ranged from nine to twelve years and each group contained three girls and three boys.

For each of the twenty-four levels played, we captured two video views of the live action in the field, supplemented with field-notes. Each camera followed an individual child throughout the level. However, because the children had a strong tendency to work together in gender sub-groups, a strategy of following one girl and one boy in each group enabled us to capture most of the action in the field using only the two cameras. Each camera was mounted on a tripod and located just off the field of play, with pan and zoom being used to keep the target child in shot. The two selected children each wore a radio mike enabling us to record their conversations with their colleagues.

However, these video recordings alone are insufficient for understanding players' experiences. Whereas video can inform us about players' physical actions – movements, gestures and conversations – it does not give us the system's view of events. In particular, it will not tell us when players

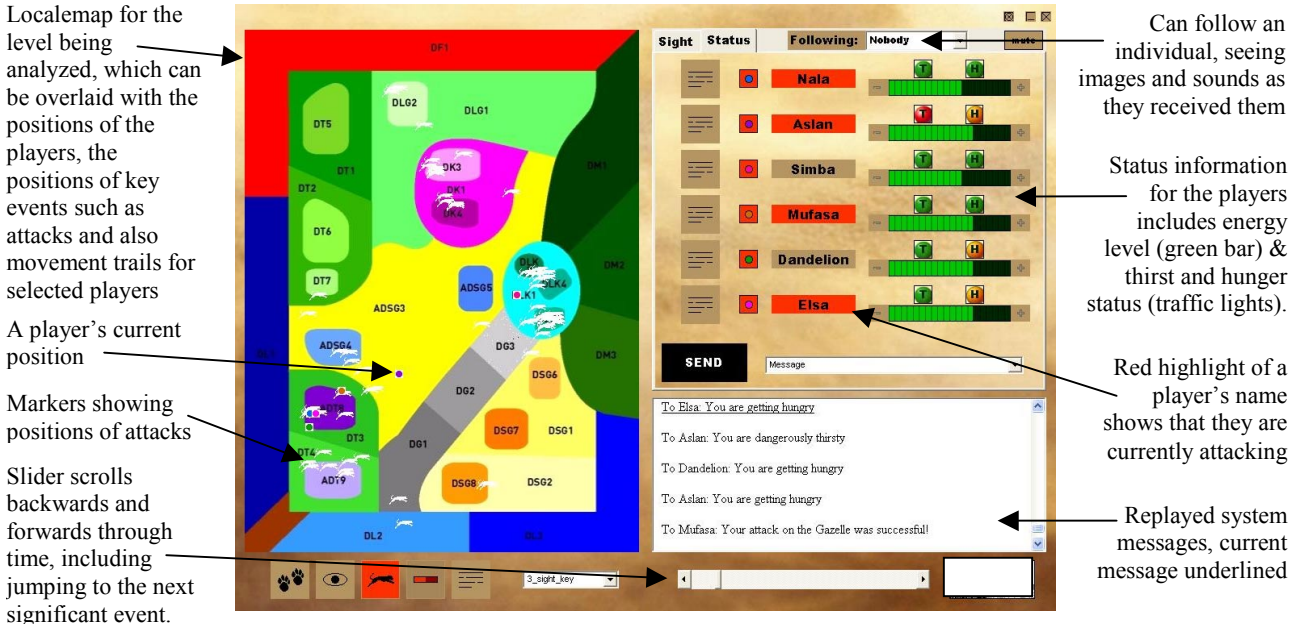

Figure 2: Interface for replaying system recordings during analysis

entered different locales, what each player was seeing or hearing via their PDA, or what buttons they were pressing. We therefore used an extended version of the Den interface (figure 2) to replay system recordings of players' interactions with the game and our analysis involved replaying these system recordings alongside the two video recordings, swapping between them in order to construct a detailed picture of players' physical activities in tandem with their system interactions.

EXAMPLES OF COORDINATION DIFFICULTIES

Overall, we believe the game to have been an engaging learning experience for children; the game was clearly understood and generally playable and we saw many examples of groups successfully hunting together. Playing the game was a highly physical experience: players would frequently run, shout, shriek, gesture, jump and even dance, and some were red-faced and panting for breath at the end of a level. Not only was this physicality enjoyable, but it also led to a highly dynamic form of collaboration in which players would rapidly form subgroups, act and disband over the course of a few tens of seconds, and in which group membership could be extremely fluid.

However, our observations also revealed significant difficulties with coordinating activity and these define the focus for the remainder of this paper. We now present four sequences of action that illustrate the range of issues involved. In the following, Elsa, Nala and Dandelion are always girls and Mufasa, Aslan and Simba are always boys.

Sequence #1

Nala, Dandelion and Elsa are hunting as a group on level three when Elsa suddenly stops and calls to the other two who are just ahead:

Elsa: Ohh look! OK ready? One ...
Nala: [indistinct] we haven't got it
Dandelion: we haven't

Nala and Dandelion turn back to Elsa and the three form a tight circle facing inwards, a formation that they maintain throughout the sequence. Indeed, players would often stop straight away on encountering a target and stand shoulder to shoulder to compare PDAs, especially when systematically searching. Elsa wishes to immediately launch an attack and has begun to count out loud in order to help synchronize with the others, also a common behavior. However, as we now see, Nala and Dandelion are unable to join in:

Elsa: ... two ...
Dandelion: I haven't got it
Elsa: One ...
Dandelion: Ready?
Nala: Wait wait wait wait wait wait wait

Elsa's counting out loud gives Nala the opportunity to prevent the attack and the three lions compare PDAs again.

Nala: I'm going to wait until it just changes. It'll change in a sec. [then to herself] Please change.
Dandelion: No mine's changed back again
Elsa: Err, one ...
Nala: Hold on wait wait wait
Nala: [Holds her PDA out from her body and reorients it in a searching gesture] Please wait. Come on. It's got to work
Dandelion: Mine's not worked
Elsa: Mine is
Dandelion: [Looking at Elsa's PDA] Yeah ' cos youre standing ...
Elsa: One ...
Dandelion: I can't even see now

Figure 3. The three lions stand in a tight circle to compare PDAs

Even though they are standing right next to each other throughout, as shown in figure 3, the three are clearly struggling to acquire a common image of the impala and the associated attack button, preventing them from attacking together. Replay of the system recordings reveals why this is so. Elsa has entered locale ADT8 (the impala), skirted its western edge and stopped close to its boundary. When the other two lions gather around her, the group ends up spanning the boundary, receiving different images and sounds. Early on in the sequence (figure 4 left), Dandelion manages to just enter the locale, while Nala remains outside. Towards the end (figure 4 right), Nala has entered, while Dandelion has drifted out, perhaps due to GPS jitter.

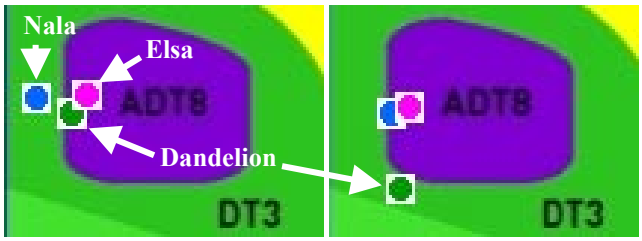

Figure 4: shifting virtual positions during the attack

Nala's behavior is revealing here. Having already played three previous levels, she seems to be aware of the possible effects of jitter and/or system latency, as she explicitly waits on the spot for the image to appear commenting "It'll change in a sec". In addition, the distinctive gesture of moving her PDA a little while holding it away from her body appears to be one of trying to 'catch' the target; we see other players undertake such gestures when trying to acquire content that they believe to be close by, although sometimes at a larger scale (e.g., by walking about). Similar sweeping gestures have been observed in previous location-based experiences [9].

Elsa now receives a message from the game that she is hungry and this prompts her to finally launch an attack:

Elsa: Two ... . You're getting hungry. I'll attack. OK
Nala: OK
Elsa: One, two three

Dandelion: Oh no mines not …
Dandelion: Don't do it. Don't do it
Elsa: I've attacked
Dandelion: Ahhh flipping hell .. go [indistinct]
Elsa: I've attacked. I can't help it.
Nala: Sorry
Elsa: Well I'll do it again
Dandelion: Who had it? Attack?
[They pause for a few seconds before being notified of success]
Elsa: I can't actually help it
Nala: Mine was successful
Elsa: Mine was successful
Dandelion: No. Why couldn't you wait for me for me you guys?

This sequence ends badly for Dandelion as she misses out on the kill. Ironically, it transpires that Elsa was not strictly to blame. The system replay reveals that while the three lions have been trying to coordinate their attack, two other lions Mufasa and Aslan, have been independently scouring the area, and it was Mufasa who initiated the attack just before Elsa. Furthermore, Aslan manages to join in while quickly passing through the locale; although by the time it resolves he is a long way distant (see figure 5).

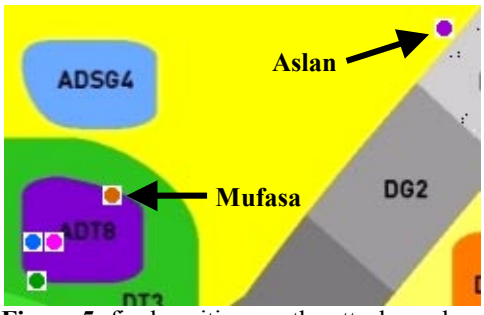

Figure 5: final positions as the attack resolves

Sequence #2

Our second sequence occurs later in the same level. The players have come up with the idea of taking physical markers and a paper map into field in order to more easily locate and subsequently remember targets (a technique that was only tried by these players for this particular level). We join them when Mufasa is explaining the secret of his success – that he regularly revisits the catfish – to Simba.

Mufasa: I just eat the catfish all the time
Simba: Do you?

Dandelion, Nala and Aslan have rushed ahead towards a pole that has been stuck in the ground to mark the location of the catfish. Simba is notified that he is getting hungry:

Simba to himself: [Reading PDA] I'm getting hungry here. Hmmm. Is there any catfish over here? I'm hungry.

He runs over to join the others by the catfish pole and is the last to arrive. Five players now bunch around the pole, again forming a circular attack formation looking inwards.

Mufasa: just keep walking around. You'll get it.
Simba: ahhhh I couldn't find any

Simba wanders round the back of the group, keeping close. The system replay shows that he is actually the first to enter locale DLK5 where he finds and attacks the catfish as shown in figure 6, right.

Simba: ahh. Yes. I've got some catfish.

Nala has wandered away from the group, but now turns round and moves back towards them so that the group is tightly bunched around the pole, as shown in figure 6, left.

Nala: Where's the catfish?
Someone: Where?
Simba: I got some

At this point Aslan performs a distinctive side to side sweeping motion with his PDA, reminiscent of Nala's from sequence #1. Simba then issues the following warning:

Simba: Don't walk too close you'll scare them away. That's what happened last time.
Aslan: oh look. I've got them now
Simba: yeah

Simba appears to interpret an idiosyncrasy of the game as a deliberate feature of gameplay. We noted other examples of players believing that animals could move around, both during play and in interviews afterwards. Indeed, such misconceptions may have added to the children's experience of the game, rendering it significantly more complex and engaging than it might otherwise have been.

Figure 6: Nala: "where's the catfish?" (left) and the virtual positions as Simba attacks (right)

As illustrated by this sequence, players were concerned with remembering the locations of targets, often pointing at or otherwise referencing particular areas of the playing field. They were also encouraged to develop shared maps of the savannah during the debriefing sessions in the Den and we have seen how this group carried this over to the actual game-play, taking a paper map into the field and using physical objects to pin down potential targets.

However, while physical markers clearly helped them establish the general locations of targets, coordination difficulties persisted. The logs show that Aslan briefly enters the locale and fleetingly sees the catfish, but then leaves again, losing them before attacking. Physically, however, he remains standing still, so this may be due to GPS jitter. In the meantime, Simba has been informed that his attack has succeeded and he walks away from the group.

Sequence #3

Our third sequence involves a different group of lions, this time playing level two. Mufasa has moved ahead of the pride. He looks at his PDA, stops, turns around and shouts to the others, who then rush over to join him.

Mufasa: I've got an elephant. Elephant! [Raises his hand] Four people. [Shows four fingers] Four people.
Boy: Elephant. Elephant.
Boy: I'm getting hungry.
Other Boy: Err, where's the elephant?
Boy: Four people are here. I'm attacking the elephant.
Other boy: You're getting hungry?

Figure 7: Mufasa: "four people"

Once again, the players form and maintain a tight circular huddle facing inwards. The system logs show them to be tightly clustered on the virtual savannah, with their virtual positions remaining stable throughout. However, the logs also reveal that, yet again, the group spans the boundary of the target locale, T5, containing an elephant. Specifically, Dandelion is stranded outside the locale (figure 8, right).

Boy: Oh shit
Mufasa: I'm attacking it. I'm attacking it.
Girl:I'm not getting hungry
Girl:I'm attacking it as well.
Girl: [Elsa?]: We're all attacking it.
Girl: [Elsa?]: Are we all attacking it?
Other girl: Attack it because we need help.
Mufasa: Oh come on baby
Girl: Attack
Girl: Everybody
Girl: Come on everybody
Girl: I can't. I can't see it. Where is it?

Figure 8: physical and virtual positions while attacking

The four attacking lions are now notified of their failure.

Mufasa: No
Boy: Ahhhh
Mufasa: You're attack on the elephant fai …
Mufasa: No!! [Raises his hand]
Girl: It's because we should have all attacked it at the same time

Here, and also in the previous sequences, we see, a notable tendency to verbalize interaction with the system, for example saying 'attack' when pressing the attack button or reading feedback messages aloud (we suspect for the benefit of other players, although possibly for the camera).

The lions now launch a second attack, this time counting in order to synchronize their actions, but their problems persist and only three manage to join in.

Boy: OK attack again
Aslan?: One, two, ….
Mufasa: you have to see the elephant
Girl: yeah, exactly and I can't see it
Boy: OK, ush ush everyone
Girl: I've got it, attack
Other girl: I haven't got it. I haven't got it.
Girl: why is everyone pressing because I haven't got it?
Mufasa: Come on you stupid elephant
Mufasa: you're going down punk

The attacking lions are now notified that this second attack has also failed and Mufasa quickly walks away from the group, followed by the other two boys. However, Dandelion now seems to have finally acquired an image of the elephant on her PDA and she manages to initiate a solitary, but ultimately unsuccessful attack.

Mufasa: No [Indistinct] your attack failed.
Aslan: Naah
Mufasa: Alright, let's leave the elephant alone. The dude's had enough
Dandelion: Mine hasn't failed
Girl: [looking at Dandelion's PDA] oh you're still attacking see
Girl: You're still attacking
Mufasa: [Looks back and shouts] No. leave the elephant alone.

Sequence #4

For our final sequence, we briefly join a third group of lions, this time playing level 3. Elsa has just been resurrected and has reentered the game from the start point. While rushing out on the field to catch up with her colleague, pursued by Dandelion who has also been resurrected, she notices an impala on her PDA, immediately stops and opportunistically launches an attack.

Elsa: You've been resurr…ected. Yes!
Elsa: Sam. I'm Alive [Laughs].
Elsa: Ooohh.
Elsa: Attack!
Elsa: [Shouts] Sam, Stef, come on!
Nala: [Indistinct in the distance]
Dandelion: I can't hear anything
Dandelion: attack.
Nala: 'tack.
Dandelion: What you got?
Elsa: Laughs [then indistinct]
Nala: Err I don't care come on.
Dandelion: Attack.
Elsa: Yesss!

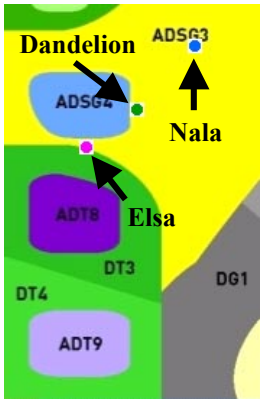

Figure 9: Elsa stops to attack far beyond the target locale

Elsa's attack is successful, but neither Dandelion nor Nala can join in. As shown in figure 9, Elsa's speed, perhaps

combined with system latency, have carried her right through and well beyond locale ADT8, the impala, before coming to a halt. Consequently, when the other players gather around her they are far away from the target. Furthermore, it turns out that Nala is not transmitting GPS updates anyway due to a problem with her PDA, although she is unaware of this.

This sequence shows that locating and attacking are not always systematic. Elsa is attacking opportunistically while moving quickly, as was Aslan in sequence #1.

THE CAUSES OF COORDINATION DIFFICULTIES

Our four sequences show that coordinating attacks was difficult and frustrating at times. Perhaps the most obvious problem was that of establishing a shared context – players who were standing next to each other, deliberately trying to coordinate their actions, could not acquire the same information. Understanding these coordination difficulties requires us to unpack the complex relationship between player behavior and the characteristics of the technology. We summarize the most significant observed behaviors, as illustrated by our sequences, as follows:

- Rapid movement and dynamic group formation;
- A tendency to halt on encountering new information;
- Encouraging others to gather by shouting and gesturing;
- Often attacking systematically, standing still in circular formations and facing inwards so as to compare PDAs;
- Sometimes attacking opportunistically on the move;
- Counting in order to synchronize attacks, also providing opportunities for others to prevent or delay them;
- Sweeping gestures with PDAs when acquiring a target;
- Remembering the locations of targets and, in one session, even marking them with physical objects;
- Verbalizing interactions and system messages; and
- Sometimes explaining difficulties in terms of game play.

Attacks spanning boundaries

Perhaps the most striking aspect of coordination difficulties is that attacking groups often spanned locale boundaries. Even when players systematically adopted tight-knit circular formations, it was common for some to fall inside a locale while others remained outside (sequences #1, #2, #3). Given that locale boundaries were invisible, it was difficult for players to understand or avoid this.

Even though we might naturally expect groups to span locale boundaries from time to time, we suggest that a combination of behaviors may have exacerbated the problem in this case. First, the tendency to physically halt and act as soon as a target is found, rather than pushing on into a locale, may lead to attacks being initiated near to boundaries, at least when players are moving slowly and systematically. Second, other players often gather in a circle around the initiating player, which is not surprising given that players are the only visible landmarks on the field. Third, some of our locales were so small that large groups could not easily fit into them.

However, sometimes the converse happens, with the system grouping players together because they happen to be the same locale, even though they are not collaborating, as we saw with Mufasa and Aslan's involvement in sequence #1.

A further, less frequent, problem resulted from attacking on the move (sequence #4), where a player's momentum carries them through and beyond a target locale before an attack registers and they come to a halt. Others then try to form an attacking group far away from the initial locale.

The impact of GPS

Of course the technology was also a factor in coordination difficulties. GPS has well known limitations as a location sensing technology with a variety of error types including drift, jitter, lag and unavailability. Although GPS worked relatively well as anticipated given the use of an open football field, our observations do reveal some important subtleties in its impact on coordination.

First, even small inaccuracies and jitters of just one or two meters assume much greater significance when players are close to a locale boundary. Second, the Garmin unit that we used incorporates a Kalman filter tuned to reduce jitter for a moving user. However, this leads to jitter apparently increasing when they come to a halt. The use of dead reckoning may also have introduced an overshoot, as for a short while their position continues to be interpolated as if still moving before being corrected. As a result, players may robustly discover the boundary of a locale when walking, but then appear to move in and out of the locale randomly while standing still and attacking. Furthermore, while the initiating player may move slowly, subsequent players often run over after being summoned and so may be more affected by these problems.

Third, facing inwards in circular attack formations may slightly broaden the apparent radius of groups during attacks because the GPS receivers are on the players' backs and therefore further from the centre of the group then their PDAs. Moreover, the position of the GPS unit high on the player's back means that the player's back, neck and head tends to occlude parts of the sky and hence some of the available satellites. When multiple players form an inward looking circle, their bodies will tend to occlude different parts of the sky and each GPS unit will compute its location from overlapping but different sets of satellites, leading to differential errors. Consequently, not only may players in a circle experience greater uncertainty in their perceived locations, but these uncertainties are unlikely to be correlated. As a further note, we observe that in carrying out sweeping gestures players, quite naturally, assume that the PDA is the sensing object, when in fact the sensor is maybe half a meter away on their back and therefore rotates in the opposite direction to the PDA. While probably not a factor in our case,

this could be an issue for sensing technologies that work with centimeter accuracy.

It is important to note that all of these GPS-related problems are amplified precisely through common attacking behaviors – halting, being near a boundary and facing inwards in a circle.

System latency was also a factor in players' difficulties; there could be a few seconds delay between a player's PDA sending a position update, receiving new information from the gameserver and the player reacting and coming to a halt. While a slowly moving player might not move very far into a locale in this time, a quickly running player could easily pass through it and beyond as we saw in sequence #4.

EXTENDING LOCALES

We now propose some extensions to our basic locale mechanism intended to address the observed coordination difficulties. In general, this might be done at several levels within the overall structure of the experience.

The design of the content – designers might affect the experience by specifying locale size. A large locale can accommodate a large group of players, will be easier to find and will be less affected by jitter (there is less chance that a player will jump right out of the locale). In contrast, a small locale will be relatively difficult to find, will only accommodate a few players and will reveal the effects of GPS jitter, although this may result in the perception that the target is moving. As a general rule of thumb, a large ponderous target such as an elephant should be placed in a large locale whereas a small skittish target such as a gazelle should be placed in a small locale. Peppering an area with many small locales might create the effect of a herd of fast moving targets, each of which is difficult to find but can be brought down by an individual player when flushed out.

The deployment of sensors – physical design choices appear to affect the quality of the GPS data. Positioning receivers on players' backs or fronts may cause problems if they adopt circular formations. Head-worn sensors may lead to the most mutually consistent GPS data, although may not be convenient. Mounting sensors on the PDA may best match player expectations, but may lead to poorer performance. We should also be cautious about use of Kalman filters, dead reckoning and related techniques, that may not be tuned to smoothing out sensor data at the most relevant moments, e.g., when players are standing still.

The interpretation of sensed positions – we can also redesign the way in which the software interprets the GPS sensor data in order to decide when a group of players has managed to gather together in a common locale in order to launch an attack. We might extend the current locale mechanism in which target locales are discrete bounded regions with fixed, absolute locations and sizes to be more flexible with regard to supporting a range of behaviors. The remainder of this section explores this third option in greater depth.

Extending locales to deal with boundary effects

Inspired in part by previous research into managing group formation in collaborative virtual environments [4], we propose two incremental extensions to the current locale mechanism, both of which temporarily extend the boundaries of a locale to more effectively enclose a group of players, but in different ways.

Two level locales: a locale is divided into two levels, an inner 'trigger zone' and an outer 'capture zone'. In order to initiate a collaborative action the first player involved has to enter the inner trigger zone. Once initiated, all players in the capture zone then have the opportunity to join in. In effect, the entire locale temporarily expands to capture players who are in its general vicinity as shown in figure 10 (a).

Personal auras: each player has a personal aura which projects their presence into the space around them. If a player initiates an attack in a locale, any other player who is within their aura is able to join in, even if they are not directly within the trigger locale. In effect, the locale temporarily bulges out to encompass players who have gathered near to the initiating player as in figure 10 (b).

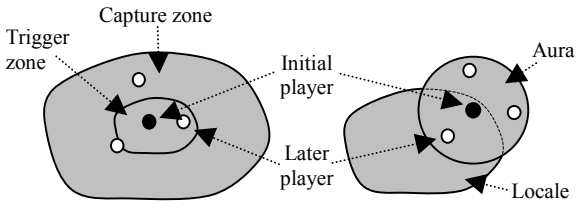

Figure 10: two ways of temporarily extending locales

In comparing these approaches, we consider two important criteria that emerged from our observations:

- How flexible are they at dealing with different behaviors, ranging from systematic attacking in circular formations to opportunistic attacking on the move?
- To what extent do they do support players remembering where targets are likely to be?

Two-level locales favor systematic behaviour, where a player finds the centre of a locale and stops there, providing a target for others who then gather around them. In such cases, they will tend to steer the action towards locale centers, encouraging players to better remember the locations of targets for future attacks. Put another way, two level locales encourage players to find the centers of locales rather than their boundaries.

Personal auras act in the reverse way, grouping players together, but leaving them at the boundaries of locales. We therefore anticipate that they might be less supportive of players remembering the locations of targets. On the other hand, personal auras may be adaptable to less systematic player behavior. For example, the initiating player's aura rather than the original locale might assume the role of the

target for subsequent players as shown in figure 11 (a), potentially dealing with the behavior of attacking while on the move. Metaphorically, it is as if the initiating player is dragging the prey with them. Another potential problem is chaining, where a player's tendency to stop on encountering a target overrides their urge to group with other players, so that later players may still halt on the edge of an expanded locale. Personal auras might be applied transitively to overcome this problem, as in figure 11 (b). However, these additional flexibilities come at the cost of enabling attacks to occur even further away from the original target locations, potentially further hindering spatial memory.

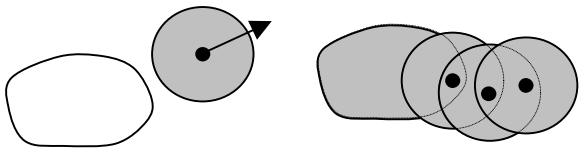

(a) A player carries their still active aura with them as they move off

(b) Three players and a locale form an aura chain

Figure 11: mobile triggers and chains with personal auras

We can employ a third criteria to compare these approaches; the flexibility that they offer designers for creating different kinds of content. The key difference here is whether they enable the experience to be tailored on a per-locale or per-player basis. With two level locales, the sizes of target zones and capture zones can be tailored to reflect different content types. A large camouflaged animal such as a zebra might be placed in a locale with a small trigger zone (initially hard to find) but with a large capture zone (easy for others to see once discovered). In contrast, personal auras can be tailored for individual players, for example a group leader can be given a large aura that more easily groups together other players.

In summary, both mechanisms introduce some additional flexibility into locales with the intention of providing a more consistent player experience. However, our comparison suggests that choosing the right solution will involve complex trade-offs between likely behaviors (systematic or opportunistic) and also balancing the need for an immediately consistent experience against the ability to remember target locations at a later time. We also noted further options in terms of designing different kinds of content and deploying sensors. This is clearly complex design space and so in the final section, we step back from the problem to reflect on the more general issues involved.

BROADER DESIGN ISSUES

In comparison to previous studies of location-based experiences, this study has directly focused on the issues surrounding tightly coordinated action among groups of players. The most significant issue raised by our study concerns group formation. Specifically, there appears to be a considerable divergence between the players' view of forming a group and the underlying system's view, and this demonstrably can lead to serious confusion and frustration when players are unable to establish a shared context and so act together.

Considering the player view first, we suggest that the users of collaborative location-based applications engage in the following general activities:

- **Locating resources** – generally exploring the environment or specifically searching for the resources that they require for the experience.
- **Congregating** – attracting the attention of colleagues and sharing information with them, both at a distance by shouting and gesturing and also by gathering together. A key concern here is establishing a shared context as the basis for possible action.
- **Acting as a group** – deciding whether or not to join a group action, initiating the action, experiencing the resolution of the action and then disbanding.

However, it is important to note that these processes are often highly dynamic and also interleaved, both temporally and spatially, for example some players may be acting, while others are congregating while others may be locating resources on their own, even though they are close by.

The system view of grouping, in direct contrast, is far more rigid. The system interprets multiple sources of data, which are often clouded by uncertainty and latency, and makes concrete and discrete decisions about when groups have formed. In doing this, the system may grossly simplify the nature of groups, for example assuming that group membership is determined by spatial proximity, is a discrete property and that players remain proximate once they have initiated an action.

We therefore argue for a more flexible approach to supporting group membership in collaborative location-based applications. In particular, we encourage designers to carefully consider the following issues.

First, could different aspects of group activity in fact be supported by different sensing systems? While GPS may be broadly acceptable for individuals locating resources on a football field sized area, it is perhaps less appropriate for supporting congregation and acting as a group, which operate at a finer scale, with players grouping together in an area of just a few meters radius, in which case even small inaccuracies become significant.

Related to this, when is it appropriate to work with absolute position and when relative? Using absolute position to determine whether a player is in a locale may be appropriate for locating resources, but perhaps comparing players' relative positions to each other would be more appropriate for congregating and acting as a group? The answers to these first two questions may sometimes suggest the additional use of more localized and relative approaches to detecting proximity, for example using Bluetooth and RFID to support

congregation and acting as a group while using GPS for locating resources.

A further design question is to reconsider whether and how different activities should be spatially constrained. Some previous mobile applications have demonstrated interesting non-spatial approaches to sharing context. The Sotto Voce audio guide for museum visiting enables visitors to selectively eavesdrop on others information [2] whereas the Schminky audio-based game enables a player to invite others into an experience [12]. Like shouting and gesturing, such mechanisms might usefully extend the activity of congregation away from the immediate locale.

Designers need to consider whether locale boundaries could be fuzzy or flexible. We have already proposed some simple ways of temporarily extending locale boundaries. Another option might be to use orientation as well as positional information, for example by incorporating a digital compass, to support a spatialized media presentation in which players locale boundaries would appear more fuzzy and players would be guided towards their centers.

Our final questions concern the temporal aspects of events. How should the activities of locating resources, congregating and acting as a group be ordered, can they be interleaved and at what point does the system formally recognize that a group has formed? In our current game, the system forms or expands a group only when players press the attack button. Other possibilities would be to form the group as they congregate, in which case acting as a group might involve quite different mechanisms (e.g., voting) or even form groups right at the beginning of the experience, for example sending out two predetermined groups of three lions from the start. Conversely, when do groups disband? For example, can players leave a group before an action resolves (e.g., by walking away)?

SUMMARY

Our study has revealed some of the complexities involved in designing collaborative location-based experiences. It appears that (what to us at least) was the most obvious approach – using discrete, bounded locales as trigger zones for content – is fraught with difficulties that arise from interactions between player behavior, content design and the operation of the underlying technologies. Although we have proposed some specific extensions to the locale mechanism that we believe may introduce useful additional flexibility, it seems unlikely that there is a universal solution that will suit all applications. Rather, designers need to carefully consider the different activities involved, possibly introducing different mechanisms and constraints for each, and being clear about how they are interleaved.

In terms of future work, it will be interesting to see to what extent these issues arise in other applications, especially those in which virtual content is attached to existing physical landmarks rather than being overlaid on an empty space or is itself moving through the game environment. Studies of such applications will shed further light on the complexities of locating resources, congregating and acting as a group and the mechanisms that might support them.

ACKNOWLEDGEMENTS

We are very grateful to Mark Jacobs, Tim Scoones and colleagues at the BBC's Natural History Unit for their collaboration. We would also like to thank Kathy Fawcett and Neil Dennison for curriculum advice and facilitation, and the I-City Learning Centre at Monks Park School, Bristol for their participation. We are also grateful for additional support from the EPSRC through the Equator Interdisciplinary Research Collaboration.

REFERENCES

1. Abowd, G., Atkeson, C., Hong, J., Long, S., Kooper R., Pinkerton, M., Cyberguide: A mobile context-aware tour guide. *Wireless Networks*, 3, 421-423, 1997.
2. Aoki, P., Grinter, R., Hurst et al., Group Spaces: Sotto voce: exploring the interplay of conversation and mobile audio spaces, *CHI 2002*, 431-438, Minneapolis, 2002.
3. Bellotti, V., Back, M., Edwards, W. K., Grinter, R. E., Henderson, A. and Lopes, C, Making Sense of Sensing Systems: Five Questions for Designers and Researchers, *CHI 2002*, 415-422, Minneapolis, 20-25 2002, ACM.
4. Benford, S., Greenhalgh, C and Lloyd, D., Crowded Collaborative Virtual Environments, *CHI 97*, 59-66, Atlanta, USA, 1997.
5. Benford, S., Seagar, W., Flintham, M., et al, The Error of our Ways: The Experience of Self-Reported Position in a Location-Based Game, Proc. Ubicomp 2004, Nottingham, September 2004, Springer.
6. Bjork, S., Falk, J., Hansson, R., Ljungstrand, P. (2001). "Pirates! Using the Physical World as a Game Board", Proc. Interact 2001, 2001, IFIP.
7. Brown, P. & Jones, G., Context aware retrieval. *Personal and Ubiquitous Computing*, 5, 253-263, 2001.
8. Cheverst, K., Davies, N., et al, "Developing a Context-Aware Electronic Tourist Guide: Some Issues and Experiences", *CHI 2000*, 17-24, The Hague, 2000.
9. Crabtree, A., Benford, S., Rodden, T., et al, A. 2004 Orchestrating a mixed reality game 'on the ground', *CHI 2004*, Vienna, April 2004,.
10. Imielinski, T., and Navas, J., GPS-based geographic addressing, routing, and resource discovery. *CACM*, 42 (4), 86-92, 1999.
11. Piekarski, W. and Thomas, B., ARQuake: the outdoors augmented reality system. *CACM*, 45 (1), 36–38, 2002.
12. Reid, J., Lipson, M., Hyams, J., Shaw, K., Fancy a Schmink: A novel networked game in a café, *Proc. Advanced Computer Entertainment, ACE 2004*, Singapore, June 2004, ACM Press.
13. Want, R., Hopper, A., Falcao, V, & Gibbons, J., The active badge location system. (1992) *ACM Trans Information Systems*, 10 (1), 91-102.
14. Want, R., Schilit, B., Adams, N., Gold, R., Petersen, K., Ellis, J., Goldberg, D., & Weiser, M., The PARCTab ubiquitous computing experiment. *Tec. report CSL-95-1*, Xerox PARC, 1995.
15. Weal, M., Michaelides, D., Thompson, M., and De Roure, D., The Ambient Wood Journals - Replaying the Experience**..***Hypertext'03*, Nottingham, 2003, ACM.