

Sharing the square: collaborative leisure in the city streets

Abstract. Sharing events with others is an important part of many enjoyable experiences. While most existing co-presence systems focus on work tasks, in this paper we describe a lightweight mobile system designed for sharing leisure. This system allows city visitors to share their experiences with others both far and near, through tablet computers that share photographs, voice and location. A collaborative filtering algorithm uses historical data of previous visits to recommend photos, web pages and places to visitors, bringing together online media with the city's streets. In an extensive user trial we explored how these resources were used to collaborate around physical places. The trial demonstrates the value of technological support for *sociability* - enjoyable shared social experiences. We also discuss support for collaborative photography and the role history can play in collaborative systems.

Introduction

Supporting co-presence, collaboration and shared experiences between distant individuals are long-standing goals of CSCW research (Gaver 1992). The many limitations of current collaborative technologies, such as telephones and video conferencing, have prompted researchers to explore new ways of sharing space, objects and presence. Techniques such as moving cameras (Kuzuoka, et al. 1994), laser pointers (Kuzuoka, et al. 1994), multiple screens (Gaver, et al. 1993) and mobile robots (Paulos and Canny 1997) have all been used to support shared

interactions, with some limited success. However, despite some success current systems require considerable set-up and configuration and are predominantly designed for use in stable office or work settings. In this paper, we present a more lightweight approach to sharing space at a distance, focusing specifically on *mobile* users and collaboration as part of *leisure*. Building on ethnographic studies of tourism (Brown and Chalmers 2003), previous systems (Brown, et al. 2003), and conceptual work on weaving media together (Chalmers and Galani 2004) the *George Square* system uses a small, portable tablet PC to allow a mobile visitor to explore a city while sharing his or her location and activity with distant others. The tablet is connected via the Internet to other users running the same software on other tablet or desktop PCs. This allows users to share the place being visited in both its online and on-site physical aspects. The scenario we explore in this paper is of a tourist visiting a city square, sharing that visit with a companion who is at home, although our system is generally applicable to sharing of places at a distance when mobile.

The system provides four collaborative resources for sharing places. First, users' locations are tracked using GPS and displayed on a map, with non-mobile users able to move an equivalent avatar by clicking on the map. This supports a shared sense of context in terms of location. Second, users can share photographs taken from an attached camera. Third, users' behaviour is recorded and compared to the history of others' past behaviour, producing a focused set of recommendations of places, web pages and photos to be displayed on the map. Lastly, the system uses voice-over-IP to support talk between participants.

In an extensive trial of the system we studied how the system could support a shared visit to a city square across the Internet. In particular, we describe how the system supported the enjoyable aspects of shared visiting, in particular *sociability* (Simmel and Hughes 1949), the experience and enjoyment of shared experiences with others. In use, the *shared conversational resources* of the system proved to be of primary importance in that photos and webpages provided visitors with topics to discuss as they visited together. Along with collaborating around the viewing of photographs, participants also shared the *taking* of photographs—collaboratively creating and arranging photographs of the square. Lastly, the system's use of history, through the recommendation system, worked to bring together online aspects of the visited square with the physical site. Broadly, this paper underlines the potential of mobile CSCW systems designed for leisure and sociability.

Previous work

City visiting has been a popular area for research in mobile information systems, in particular (Cheverst, et al. 2000) and other PDA-based systems (Abowd, et al. 1997; Fesenmaier, et al. 2000). Indeed, as mobile phones and other portable devices become more advanced, tourism is one obvious application area. A number of phone

operators have released city guides that can be accessed on one's phone (e.g. <http://www.lonelyplanet.com/mobile/>), however, these and other commercial technologies have had only limited success. Generally, these systems are based around a 'walk-up, pop-up' model where information, such as text and pre-recorded speech, is pushed at a user based on his or her current location. There has been little explicit support for *collaboration* between visitors (with rare exceptions such as *Sotto Voce* (Grinter, et al. 2002)). As studies of tourism underline (Brown and Chalmers 2003), collaboration is a key part of visiting and of leisure more generally. Indeed, online collaboration has recently developed into a surprisingly popular form of leisure in its own right, in the form of online games. This success suggests the value of experimenting with new forms of shared leisure experiences.

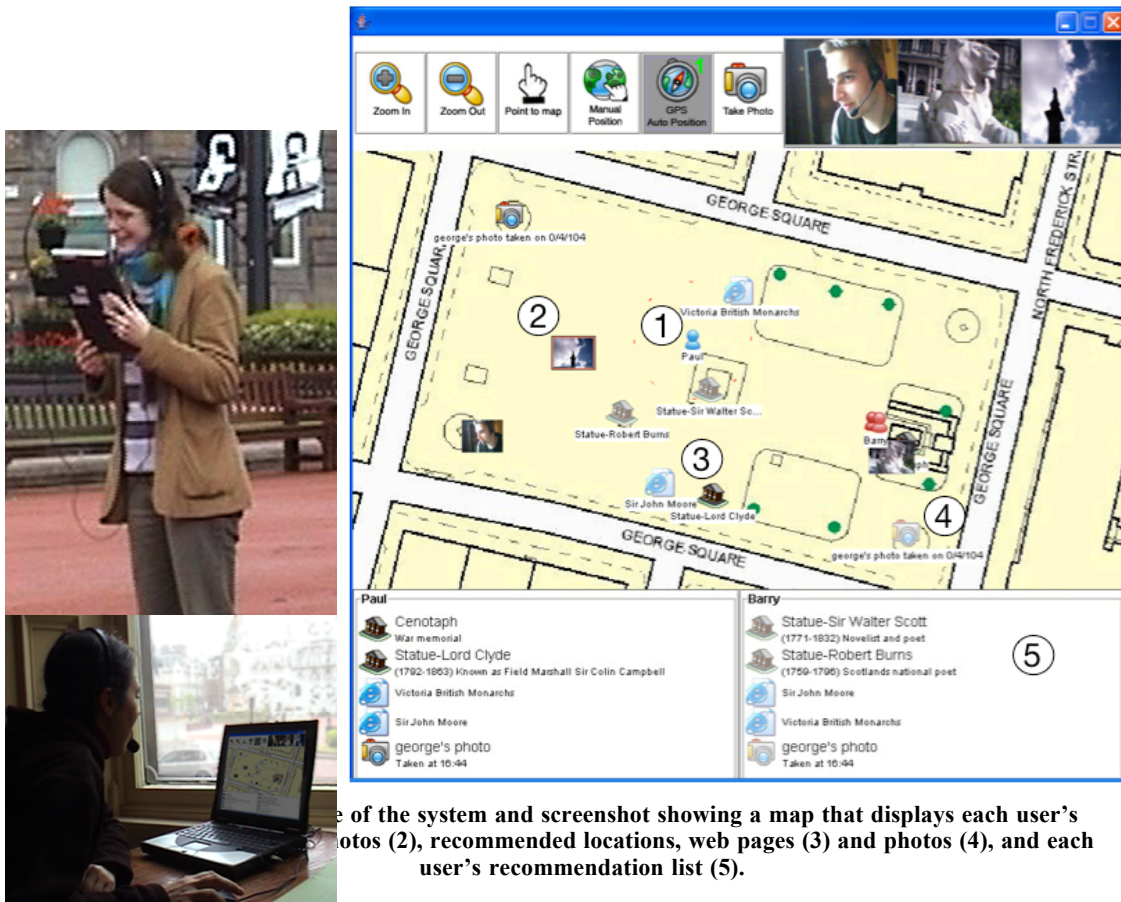
The Lighthouse system (Brown, et al. 2003) attempted to support collaborative museum visiting by connecting online interactions with traditional visit experiences. This system supported collaboration between online museum visitors and those visiting the site of a museum. A group of visitors, each in a different location, used VR, 2D maps, the web and audio links to share a museum visit. The on-site visitor in the museum used a PDA with an ultrasonic tracking system to collaborate with two online co-visitors using a virtual reality (VR) and web version of the museum that we had created. In a trial of this system, we highlighted how users could bring together digital and physical exhibits through their interactions. Visitors would share and interact around exhibits which were physically in the museum for some visitors, and were presented as web pages to others, forming what we called 'hybrid exhibits'. These were exhibits that linked places and electronic information about them, through visitors' collaboration.

Although this system demonstrated the feasibility of collaborative leisure experiences, the system itself had a number of limitations. The system was restricted with regard to mobility, since its use was fixed to one location, and not easily scalable beyond the Lighthouse, since web and VR versions of each new setting would have to be produced at some cost and effort. The differences between users were also artificial in that the design imposed a restrictive model of distinct PDA, web and VR users, rather than allowing users to choose their own configuration of devices and tools. Lastly, visitors were also to a large extent passive consumers of the content of the museum, rather than having any opportunity to leave lasting comments or contributions themselves.

System Overview

To address these limitations, we have developed support for collaborative leisure further with the *George Square* system. This system breaks with our previous work in a number of ways. It is designed for use outdoors in the city streets, working anywhere a network connection is available. In this less constrained setting, content

is much harder to produce, so its use is centred on the sharing of the existing digital resources that are available about places, such as maps and web pages, and content that users themselves produce, such as photographs. The system also reverted to the more common practice of users—mobile or stationary—having the same software, rather than the strong asymmetry explored in the Lighthouse.



In use, each tourist can visit the city much as they would in a normal city visit. On each tablet PC, the visitor's location is tracked using a GPS and shown (e.g. 1 in Figure 1) on a map of the city. Maps are automatically downloaded over the Internet from a map server, allowing the system to be run anywhere we have map data available for. As an alternative to specifying location via GPS, visitors can select a 'manual position' mode, and then click on the map to specify their position. As a visitor moves around the square, he or she can take photographs of attractions using an attached camera. The pictures are geo-referenced and shown on all users' maps at the location where the picture was taken (2). These pictures are also shown in temporal order on a shared 'filmstrip' view at the top, alongside buttons to control the map's zoom level, briefly highlight a position on the map, change positioning mode and take a photo.

User context and activity is logged in a database, recording the attractions in the square each user encountered, web pages browsed and photographs taken. The last few minutes' log entries are used to find periods of time with similar context and

activity data in the logs of all previous visitors. This is the first stage of the Recer collaborative filtering algorithm (Chalmers, et al. 1998), which we use here to find attractions and web pages (3) accessed by previous visitors in similar contexts. Pictures taken by visitors in similar contexts are also recommended (4). These recommendations are displayed on each user's map, and in a legend below each map (5). In order to support sharing and discussion, one sees others' recommendations 'ghosted' on one's map, and sees others' recommendation lists alongside one's own. Map icons for web pages and photos can be clicked to view the related content in detail. Lastly, a voice-over-IP subsystem allows visitors to talk as they visit together.

In use, the system supports a range of different scenarios. First, it can support two users physically co-visiting an area of a city, taking photographs and browsing web pages about that area. Second, it can support users collaboratively co-visiting with other distant from the where they are, via the Internet. This is the scenario shown in figure one and described in more detail below. Third, it can support users who are all distant from the area but interacting via the Internet. The latter scenario is important since our observational studies of city visitors emphasised that the visit itself is only one part of a visitor's experience; the 'pre-visit' and 'post-visit' have an important role for both planning and sharing. Our design therefore supports users in planning their visit in advance, and in reviewing their visit afterwards. The database log generated from earlier visiting is used to generate a web page: a travel weblog. One can browse the web pages generated from one's visit, look at a temporally-ordered list of all the pictures, web pages and places that one has visited, and explore a map based on the one used during the visit but summarising one's visit in a spatial presentation. (The post-visit 'blog' is discussed in more detail in a forthcoming paper.)

Our use of past activity to build up content in the form of webpages and photographs gives the system considerable flexibility. It can be run in a new city with the minimum of reconfiguration—new content does not need to be produced in advance, since it will automatically accumulate from usage of the system. Together these features develop further the concept of co-visiting, in the form of a lightweight mobile system that can be run almost anywhere with the minimum of configuration, pre-authoring and setup.

Technical overview

The implementation challenges for George Square were typical of collaborative mobile systems, in that we needed a dynamically changing set of devices working together as peers without continuous reliance on a central server, supporting users and devices joining or leaving at any time during the visit.

The hardware of our system consists of a lightweight tablet PC with attached compact flash GPS unit and a USB 'stalk' camera. Headphones and mic were

plugged into the unit, and the built-in wi-fi was used for communications. In our trials, to provide Internet access, a temporary wireless network was bridged to a publicly available wi-fi 'hotspot'. This allowed users to browse and search the web, and to follow links to information provided by our system.

For our software we used the EQUIP distributed tuple space system (Greenhalgh 2002), middleware which supports a peer-to-peer communication model between networks of sensors and output devices via stores (or 'spaces') of records (or 'tuples'). EQUIP is used to send data both between the different devices, and system components. Tuple space events are used both for data sharing between components and network communication, supporting the flexible combination of system components. By using a peer-to-peer architecture, each component can also be used without reliance on a central server. The event-based architecture allows devices and users to leave or join at any time, with dynamic and automatic reconfiguration. Events describing user activity and sensor readings (e.g. GPS) are recorded by logging components, and entered into a database run by each user. These logging components also continually run algorithms comparing recent activity with historical logs, to create anonymous recommendations. After each visit, we aggregate individuals' databases, so that their logs add to the shared source of recommendations for future users.

The Trial

We ran an extensive user trial of the George Square system in the city streets of Glasgow. In evaluating the system we were sensitive to how it could support enjoyable interactions around place, rather than an optimal, yet potentially sterile, experience. Our focus was thus on the lessons we could learn for designing for enjoyment, as much as evaluating how well our specific system performed.

Procedure

We ran a trial with 20 participants, in pairs of two, recruited as pairs of friends. We chose a mix of locals (10) and visitors (10) to the city, recruiting participants through the city's tourist information centre, language schools and our university. Ages ranged from 19 to 35, with 13 female and 7 male participants. Participants were paid for their time at the end of the visit. Each trial lasted between 35 and 60 minutes, with a post-trial debriefing of 10 minutes.

Each pair of users was taken to George Square, an open city square (125 meters by 90 meters) in the centre of Glasgow. This square is a focus for tourists in the city, has a number of statues, monuments and gardens in it, and is surrounded by several major civic buildings. One user was taken to an indoor venue on the corner of the square (the *indoor visitor*), and one visitor was taken out to the square itself (the *outdoor visitor*). The outdoor visitor was given the tablet computer as described

previously, while the indoor visitor sat at a conventional laptop PC, equipped with a USB camera (Figure 2). The indoor visitor had limited visual access to the square, however a frosted window and the seating arrangement obscured this view.

The scenario for the trial was of two friends sharing a visit to George Square, communicating via the system. For the trial, participants were asked to freely explore the square, learning how to use the system and sharing their visit to the square. Since we were specifically interested in how the system could support social interaction, for the last ten minutes of the trial the visitors were given a short list of tasks, such as sharing a photograph of the square, and finding out the height of the statue in the centre of the square.

Analysis

While we had designed our system for a specific application area, we were also charting a new social experience with many differences from existing city visiting. The popularity of online social experiences (such as games and chat) underlines the potential value of new social experiences, although we were unsure if this would transfer to the outdoors. Our approach was therefore to explore what worked, or did not, rather than compare the experience strictly to what went before. A range of data from each trial was collected: video tapes of the outdoor visitor, video of the indoor visitor, audio recording of the participants' communication, and log data of the system and users' behaviour. For analysis we combined the shared audio channel and the video images into a single video stream. From the logs, we generated a 'replay' visualisation of the system as seen by the trial participants, and this was superimposed onto the video stream. We also analysed transcripts of the post-trial debriefings, and our general observations of the use of the system. We were interested in exploring how the system was used, to inform our future designs. Accordingly, we chose a technique known as *interactional analysis* [17], based on paying close attention to the details of how users interact with each other and with technology, usually through the analysis of video. We paid special attention to where the participants used the resources provided by the system, such as location awareness. Having a visualisation of the system's behaviour allowed us to better interpret users' reactions to system events. In particular, situations where participants were confused showed something of where the system could be improved to better support collaboration or understanding.

Results

In use, the system presented a novel yet seemingly enjoyable experience for trial participants, with all participants exchanging photographs, and using their location and recommendations in their interactions around the square.

Division of leisure

A key concept in the study of work situations has been the working division of labour that develops as individuals collaborate around tasks and activities (Anderson, et al. 1989). In our trials, a complementary *division of leisure* developed. While exactly the same software was used for both indoor and outdoor participants, differences in the visitor's situation gave different capabilities to each user. The indoor visitor used a laptop with a larger screen, keyboard and mouse. He or she could type URLs and interact with multiple web pages more easily. However, this user was stationary. The outdoor user, through their presence in the square itself, could move around to different statues and attractions, taking photographs of statues and of other events out in the square. These differences in situation led to a clear division in what users did:

	Outdoor Users	Indoor Users	Total
Pictures Taken	214 (78%)	61 (22%)	275 (100%)
Web Pages Browsed	25 (20%)	98 (80%)	123 (100%)

Table 1: Number of actions by type of user (N=20)

The indoor user predominantly searched the web for information about particular statues, whereas the outdoor user would take pictures and relay information about the different statues and their plaques. As one of our outdoor participants put it: “if you can't type, you can't surf the web”. These results were confirmed by our analysis of the videos. In these, participants form a division of activities by asking the other visitor to either take photographs or find out information, overcoming the limitations of their own situation. For example, one outdoor visitor takes a picture of a statue of William Gladstone and then asks the indoor visitor to look up information on the web:

Out: (Takes photo) Oh There it goes
In: Did you take it? (.) Ah yeah got it
Out: Can you look up something about William Gladstone?



As with a division of labour, this division between visitors was not fixed or total, e.g. some web pages were browsed by the outdoor user. Indeed, the recommendation system would suggest pages to browse, placing icons on the map that were then clicked to open the page. In this way, the outdoor visitor viewed some relevant web pages without having to type in a URL or search terms. The system allowed the visitors to form a flexible division of leisure, with participants collaborating so to make the most of their different situations.

Sociability

Of all the resources provided by the system, the voice connection proved to be the most valuable for creating a sense of shared experience. Users continually talked

about what they were doing, what they had done and what they were going to do. Yet conversation during the trial was not only functional; it ranged widely covering different aspects of the square, such as the different buildings, people in the square, shared past events, and even the wildlife in the square such as the ever-present pigeons. As a historic square in the centre of Glasgow, George Square contains many statues of the 'great and the good' from British and Scottish history, e.g. Prince Albert and Robert Burns. Much of the talk of the visitors centred around trying to place these figures in history, and understanding a little of who those figures were beyond the somewhat opaque information given on each statue's plaque:

Out: On the statue to James Oswald it says given by a few good friends
In: (Browsing a web page) Eh (.) I think he's a MP it says he was one of the first Glasgow MPs. He was elected in 1841 and the statue was erected in Charing Cross and then moved to George Square
Out: No way
In: So they obviously thought he was good enough

As could be expected with a tourist visit, the more factual 'high cultural' aspects of the square were combined with the more playful. Here the tourists move quickly from talking about statues in the square, to chatting about the pigeons in the square, taking photographs of the pigeons and joking around these pictures:

Out: I'll take some more. I'll take one of a horse
In: No I don't want to see that I've got one of those. Take one of a pigeon
Out: Pigeon? What about will I take one of Barry
In: With a pigeon
Out: K wants to see a picture of you with a pigeon
In: Feeding a pigeon
Out: What? You can't feed the GPS system to a pigeon (laughs)

These sorts of conversations are notable for their playful and non-goal focused nature. Earlier ethnographic work has emphasised that tourism is as much about shared enjoyable experiences as it is about the specific place being visited. Research in CSCW often ignores these experiences, and more broadly the importance of enjoyment in companionship with others (Brown and Bell 2004). The sociologist Simmel (1949) argued for the importance of this kind of shared experience of enjoyable conversation—experience he called 'sociability', where the purpose is not external to that experience but rather is that experience itself. As he put it, when we engage in the company of others (in its purest form) we engage for the company itself. For pleasurable and enjoyable city visiting, many of these aspects of life come to the fore, for example the enjoyment of shared experiences and conversation with family and friends, in a setting that supports these experiences and this conversation.

As can be seen in the extracts above, the George Square system had some success in supporting sociable interactions. Shared photographs, recommendations and web pages acted as 'local resources' (Sacks 1995, vII p96) for conversation, in that objects that were seen in common could be used as topics to talk about. For example, in the following extract the outdoor visitor has just sent a picture of the Cenotaph (pictured) to her co-visitor. This prompts a discussion about a previous shared experience:



In: Reminds me of the other day when we went to the cemetery.
Out: (laughs) Why?
In: Eh I don't know. Because of the crosses and all that.

Alternatively, local resources could lead to photo requests:

Out: There's a shopping centre nearby there I went to a Scottish party with Scottish dancing... It was really really nice and the people either they were drunk or they were crazy. I think it was half and half
In: Can you take a picture?
Out: From the shopping center, yes if you like (takes picture)

The different features of the system did thus not just support narrowly defined tasks, but were resources for sociability. The visitors could use what the system provided to talk and share their experiences. Rather than focus on particular clearly defined goals (such as making sure they saw a list of attractions) the visitors could enjoy each other's company, talking about and around the different aspects of the square and the city.

In particular, photographs acted as rich 'seen in common' objects that could be used as topical resources for conversation in this way. They gave users something to talk about, in a similar way to the events and attractions of traditional tourist visits. Shared photographs could be used to comment on and talk around objects in the square, either through direct talk about the features of an object (such as the height of a column or statue), or by referring to events or memories evoked by the photograph. In this way the relatively simple ability to send a photograph and talk as that photograph was taken helped to produce the sociability of the shared visit.

Collaborative photography

There was a clear distinction between the use of photos in our trial, and the "photo talk" around photographs discussed in previous studies (Frohlich, et al. 2002) . While digital cameras add the ability to share photos the instant after they are taken, using the camera's LCD display, photography is predominately directed towards viewing and sharing much later on, in a different location. The acts of production and consumption are more disconnected from each other in terms of time and location. In the George Square system, the ability to share the photographs taken brought these two acts together temporally, even though it involved people in two different locations. Rather than taking photographs to be kept for later, photographs would be taken for sharing in that instant, serving as a conversation topic at that point in time:

Out: I'll take a picture of the stone lions for you they're very majestic
In: Get him to take a photo of you and the lions. Get a photo of him [videoing you]
Out: [Hang on]. Oh. Lion lion back up I can't really see. Ok come on. There you are can you see how bi(h)g they arr. (2s)
Out: Have you got it?
In: Noooh. Oh there it is (.) Very Elvis.



Here a photograph of the stone lions in the square is taken by one visitor and automatically sent to the other. After some joking talk about the size of the lions (“can you see how big they are”), the indoor visitor receives the picture and gives a humorous description of the lions: “Very Elvis”. Both the taking of the photograph and its reception are tied together in time and integrated into the conversation. Indeed, photographs often had a sequential organisation, in that the meaning of a particular photograph would depend on previous photos and the discussion around them. For example, a photo would be taken as an elaboration of a point made around a previous photograph. While we had told participants that we would give them access to their photographs after their trial, the instant sharing of photos allowed photos to act more like exchanges in conversation than as a way of capturing memories for later.

Furthermore, the features for sharing of photos produced collaboration not only around receiving photographs but also around *taking* photos. Visitors talked about the photograph they were taking as they framed the photo and waited for it to be sent to the other visitor. This talk invited requests from the other visitor for particular photographs, asking each other to take photographs of particular attractions, sometimes rejecting the photos taken and asking for a different one. The choice of what to photograph thus developed over time as visitors talked about different statues in the square, photographing and re-photographing them as they talked.

The sending of photographs had a very different use when compared with sending a continuous video stream, in that taking a photograph was an explicit action. Participants would work around framing photos, taking photos, and waiting for comments or confirmation from the other visitor. The taking and sending of a photo became a shared event, and the photographs taken by our trial participants were chosen specifically to elaborate certain points being discussed, or as pictures of specific features of the square. An undifferentiated video stream would have far less of this deliberate identification. The shared persistence of photographs in the system allowed visitors to talk about a picture, rather than having to potentially hold a camera steady to talk about a particular object.

Working around Maps

As one of the most familiar of tourist artefacts, most of us have used maps while visiting a new place. In our system, much of users’ interaction and collaboration took place around a map overlaid with information, photographs and web pages—information they needed to bring together to use the system. The outdoor visitor also had to relate the map information to what he or she could see in the square. Use of the map required that the visitors connect that information with what they wanted to do.

Since the map in our system was shared between visitors, it served both *collaborative* and *informative* purposes, and these two aspects were often combined but at times conflicted. The map displayed where the users were in the square, and

the photographs they had taken. In addition, recommendations of web pages, places in the square and photographs taken by others were generated by running the previous history of visits to the square through a collaborative filtering algorithm (Figure 3). Each visitor also saw the co-visitor's recommendations ghosted on the map, aiding their use in collaboration.



Figure 3: A screenshot of the map showing recommended places, photos and webpages. One's own recommendations shown in full colour while one's co-visitor's are shown 'ghosted'.

One example of an informative use of the map involved an early trial participant browsing the 'Wikipedia' pages about William Gladstone—pages that were then recommended to later trial participants who went to that statue. The recommended web pages were positioned on the map and acted as geographically specific web bookmarks taken from others' past behaviour. These recommendations proved particularly useful to the outdoor visitors, since they could view these recommended web pages by clicking on them, without having to navigate the web. In addition, due to the collaborative filtering algorithm picking places in the square to recommend, statues in the square visited frequently by users when previously in similar contexts were recommended more often than others.

These informative uses of the map worked in combination with collaborative uses of the map. Web pages and recommended places on the map were used by visitors when they talked about different places in the square. The indoor user, for example, would ask the outdoor user to go to a particular attraction by using the name listed under the icon on the map. The recommendation algorithm, as well as being informative, took on a role in collaboration by labelling the square. Yet this conflicted with the recommendation's role as suggestions of where to go or what to read next. As an outdoor visitor got close to a recommended place, that label often disappeared because (for our collaborative filtering algorithm) there was no longer any need to suggest that place—because he or she already there. Rather frustratingly for our users, this meant that the shared labels disappeared just at the point where (or when) they wanted to talk about them.

This tension between informative and collaborative uses caused further dilemmas in our design. Much of the literature on electronic maps argues that maps should be rotated as a user turns (Montello 2003). This overcomes what are known as ‘alignment effects’, i.e. the problems that users have in reading a map to find what places are in their visual range. By automatically rotating the map around the visitor as he or she turns, the map always faces the ‘right way’ up. Unfortunately, as experience with the Lighthouse system showed, users frequently used relational terms to refer to objects on the map, e.g. “I’m right above you now”. These terms, which would have an ambiguous meaning if the system rotated the map, causing problems for collaboration.

However, collaborative aspects of the map could work to improve its effectiveness. By talking around places on the map, visitors could help each other to find places and objects:

- In: Or see if you can get that----->
picture of that human rights plaque
that’s just near you just now.
There’s a human rights plaque that’s
just near you I’ve circled it don’t
know if you [can see]
- Out: [Human rights] plaque. Well I can’t
but I’m walking (.) Oh wait a minute
Hold on (laughs) hhh I’ve got a
thing (.) no I can’t. Put your icon
where it is
- In: Like here
- Out: Alright well I don’t know where that
is but I’ll walk towards you.



Here the participants are trying to find the ‘human rights plaque’ in the square; a particularly difficult attraction to find as the plaque is small and embedded into the ground. In just this short conversation, the visitors work at understanding each other’s divergent viewpoints, refer in conversation to the different parts of the square and of the system and finally decide where to go next. Only by talking about the map, the square, so they find the plaque and start moving towards it.

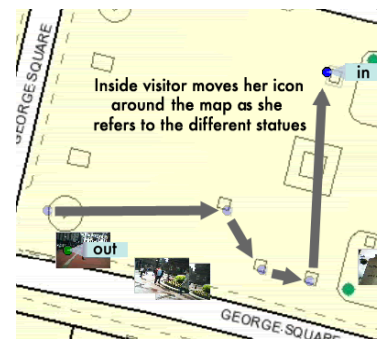
A final issue we will discuss concerning the use of the map was the use of the location system which placed users’ icons on the map. This was important in the way it helped visitors manage each others’ context. Location could be set either via a GPS unit, or via a self-reported location given by clicking on the map. With the exception of one trial, when the GPS unit had poor accuracy, the indoor visitors used manual positioning and the outdoor visitors used GPS positioning. The map proved to be a focal point of collaboration for both the indoor and the outdoor visitor. The indoor visitors made use of the outdoor visitors’ location to access the local context of the outdoor visitor. For example, by referring to a statue that the indoor visitor could see was ‘right next to you’:

- In: Take a picture of the Robert Burns statue. It’s right next to you.



When an outdoor visitor talked about a particular statue, his or her location could be used to find what statue he or she was actually talking about. The outdoor visitor's *implicit* location could be used to find out what they were looking at when it was automatically updated from the GPS readings. However, the indoor visitor had to move his or her position explicitly. This led to the indoor visitor's position having a different meaning to that of the outdoor user. Instead, the indoor position was used *explicitly* for deixis, i.e. to point to parts of the map. This need to explicitly move thus made position more a gestural resource than an implicit indicator of context. For example, in this extract the indoor visitor moves her icon around the map to point to the different squares on the map, asking if they are all statues:

In: The circles in the corner?
 Out: Yeah that's all statues again
 In: There's lots of little squares like **this** one **this** one **this** one **this** one.
 That's all statues?
 Out: Yeah all of them.



The indoor visitor uses her position to refer to points on the map. While the system had specific telepointer support for gesturing on the map, users simply moved their avatar around to gesture at different points on the map.

Design implications

Although George Square was designed to support city visiting, the synchronous aspect of its collaboration are similar to a wide range of collaborative applications, from instant messaging to picture messaging. Therefore, we will attempt to draw implications more broadly for the design of CSCW systems, in particular for systems with an application to leisure. We draw three main points: firstly we argue that CSCW systems should provide resources not only for completion of narrow tasks, but also for sociable interactions between users. Second, we argue that collaboration around photos is not just about photo sharing but also shared photo taking. Lastly we discuss how history can be used both to manage the information displayed to users and also as support for collaboration.

Designing for sociable interaction

The first lesson we draw from our trials is that in designing CSCW systems we should provide resources not only for specific tasks but also for sociable interactions. The George Square system was not used only for specific tasks in the square, but more broadly for sociable interactions around the square. The key distinction here is between designing for goal centred activities or activities with less focused benefits, where we obtain pleasure in the experience itself. Earlier studies of tourism have emphasised to us the value of these interactions as part of tourism. Sociable interaction, while not necessarily goal-oriented, is not without purpose. Rather it is part of the flow of an enjoyable experience; ‘talk for the sake of talk’ is the very stuff of sociability. Interactions of this kind are an often-neglected aspect of collaborative systems, since there is a traditional focus in design on the task being carried out, rather than talk around that task. However in creating enjoyable collaborative environments, particularly those used for leisure, this sort of talk is important in how it can help to create enjoyable sharing experiences. We enjoy doing things with others not only because of the activity itself but also because we can naturally talk around that activity. Systems should support users’ interactions around the tools they use, rather than simply the use of those tools.

Key to supporting such interactions were the ‘local resources’ that the system provided. These were topical resources in the environment, which the system shared and made available to others. In particular, our trial participants used shared photographs in this way, to tell stories inspired by photos and to express their opinions on different aspects of places. The photographs, when being taken and after being taken, acted as resources for conversation about their experiences and for more general chat. This underlines how systems can support sociable interaction through relatively simple mechanisms.

Photos are one powerful resource, but other collaborative features could be used to support sociable chat. For example, in an instant messaging client we have recently been experimenting with sharing screenshots and a history of system commands used. Although in some ways quite different to the sharing of photographs in George Square, sharing selected parts of ongoing and recent activity also provides local resources that can be used for interaction.

Collaboration around photo taking

While the popularity of camera phones shows the value of the sharing of photographs between mobile devices, camera phones currently support only the asynchronous sending of photographs (via MMS) or mobile video-conferencing (on 3G phones). The coupling of both taking and sharing photographs was a particular powerful use of photographs in our trials. This suggests that mixed modes of interaction around photos on phones may provide more value to users. For example, phones or other mobile devices could support augmenting phone calls with the

ability to take photographs and then talk around those photos. Photographs in this context are potentially more valuable for collaboration than a video stream, since the static shared display of a photograph supports talk around specifically aspects of that photograph.

Taking a photo can also be an event in itself, differentiating one moment over others, whereas a video stream would show an undifferentiated flow. We found other features helped support this richer interaction, for example the georeferencing of pictures on a map and the persistent display of photographs on the map or in a timeline or ‘filmstrip’. These allowed users to relate the current image to previous photographs and to the people, artefacts and buildings in the areas that the photographs were taken in, and hence to request particular photographs, to rejecting the photos taken and asking for a different ones, and so forth. Timestamping of images is already common in digital cameras and phones. We suggest that designers can use map-based and time-based displays of photographs and other aspects of recent activity to share with photographs on mobile phones, PDAs or even cameras. These features could support collaboration around both photo viewing, and photo taking.

Using the past for filtering and collaboration

A last lesson we draw from George Square is that the past can be used to filter information to prevent information overload, and to connect digital information with specific places. A key problem with maps, and visualisations more generally, is the need to keep the display clear from irrelevant details. This is particularly true in mobile contexts where devices may have a small or restricted display. Through collaborative filtering algorithms, systems can use statistically significant patterns in previous users’ actions to automatically prioritise relevant information. Future applications could emphasise or prioritise objects used by users with similar histories to a particular user, with similar periods within their histories (as in George Square), or with particular significance to the current user, e.g. friends or relatives explicitly chosen by the current user. Of course, users might also use more explicit means to select relevant information, via queries for example, but ‘implicit queries’ made by tracking ongoing activity can offer useful information in a lightweight way. One additional lesson we draw is that collaborative filtering of information is more than simply a means of recommendation—it is also a resource for ongoing collaboration. In our trial the displayed map items were used both as a source of information and in the collaboration between users.

In addition, our recommender also made use of historical data to weave together online information with urban locations. Photographs taken and web pages browsed by users, such as the Wikipedia page on William Gladstone used near his statue, were stored as an archive of information about particular locations. Without this gradual adaptation to users’ behaviour, a large amount of context would have had to be manually entered placed onto the map in the locations that we judged to be

appropriate for visitors (such as in . Instead, our system made use of patterns of co-occurrence of location and browsing, to place information in contextually relevant locations on the map. This suggests a broad method for making use of people's behaviour to connect together information and locations, complementary to the common pre-authored static designs.

Ongoing and future work

The George Square application takes many of the features of synchronous groupware systems such as instant messaging and conferencing software, and moves them to a mobile context. In our future work we are experimenting with both expanding the mobile features of our system, as well as integrating aspects of the system into desktop applications. The functionality from George Square is being integrated into instant messaging (IM) software, expanding the use of IM to more mobile contexts. In particular, we are exploring the use of photographs within instant messaging software and awareness of location. This is to support not only discussions around current photos, but the taking of new photos during interaction. In a very different context, we have also been experimenting with version of George Square tailored for scientists collaborating with remote colleagues in monitoring urban pollution, where tablet PCs each have an attached carbon monoxide sensor and relate ongoing pollution readings to databases and simulations of urban pollution.

Currently, the George Square system is limited by hardware constraints and by reliance on one wireless network. The tablet PCs we used have a relatively limited battery life (under two hours) and while lightweight it is still beyond a size desirable for carrying around longer than a few hours. We have ported the bulk of George Square to both smartphones and PDAs, and are piloting recommender and map systems that use ad hoc networks to spread information in a peer-to-peer way between PDAs. These different form factors enable more longer-term and more wide-ranging use of the system, in that users can choose to use the system over longer periods such as a whole day, and need not rely on one local wireless network for communications or on one particular server for storage. In future trials, we plan to experiment with more asynchronous forms of collaboration for visitors, collecting photographs and comments over days of travelling in a wide area of the city.

Conclusion

This paper has presented a study of the George Square co-visiting system. The main goal of this system was to support geo-spatial collaboration around a place as well as the information about that place, with a particular focus on support for leisure. The system supports city visitors sharing their visit with those at a distance. A trial

uncovered how, through the different resources the system provided, visitors could accomplish a shared visit. In particular we discussed how users achieved a division of leisure, used the local resources that the system provided, and collaborated around photography and maps.

Mobile technologies offer the possibility of access to large bodies of information on distant servers and stores, through information-seeking tools such as search engines and recommenders. More importantly, perhaps, they allow new forms of interaction with other people, both mobile and stationary. In George Square we have focused on these collaborative applications with a lightweight system that supports rich interactions between users. This paper shows how we can design to support interaction that weaves these apparently disparate places, times, and media into a coherent, manageable and even pleasurable whole.

Acknowledgements

To be added

References

- Abowd, Gregory D., Christopher G. Atkeson, Jason Hong, Sue Long, Rob Kooper and Mike Pinkerton (1997): *Cyberguide: A mobile context-aware tour guide*. *ACM Wireless Networks*, Vol. 3, pp. 421-433.
- Anderson, R. J., J. A. Hughes and W. Sharrock (1989): *Working for profit: The social organisation of calculation in an entrepreneurial firm*. Aldershot:Gower.
- Brown, B. and M. Bell (2004): *CSCW at play: there as a collaborative virtual environment*. In: *Proceedings of CSCW 2004, Chicago, IL*. New York: ACM Press, pp. 350-359.
- Brown, Barry and Matthew Chalmers (2003): *Tourism and mobile technology*. In: Kuutti, K., Karsten, E.H., Fitzpatrick, G., Dourish, P. and Schmidt, K. (eds): *Proceedings of the european conference on computer supported collaborative work (CSCW) 2003, Helsinki, Finland*. Dordrecht: Kluwer Academic Press, pp. 335-355.
- Brown, Barry, Ian MacColl, Matthew Chalmers, Areti Galani, Cliff Randell and Anthony Steed (2003): *Lessons from the lighthouse: Collaboration in a shared mixed reality system*. In: *Proceedings of CHI 2003, Ft. Lauderdale*. New York: ACM Press, pp. 577-585.
- Chalmers, Matthew and Areti Galani (2004): *Seamful Interweaving: Heterogeneity in the Theory and Design of Interactive Systems*. In: *Proceedings of DIS 2004, Boston, Mass*. New York: ACM Press.
- Chalmers, Matthew, Kerry Rodden and D. Brodbeck (1998): *The order of things: activity-centered information access*. In: *Proceedings of WWW 1998*. pp. 359-367.
- Cheverst, Keith, Nigel Davies, Keith Mitchell, Adrian Friday and Christos Efstratiou (2000): *Developing a Context-aware Electronic Tourist Guide: Some Issues and Experiences*. In: *Proceedings of CHI '2000*. The Hague, Netherlands: ACM Press, pp. 17-24.

- Fesenmaier, Daniel, Stefan Klein and Dimitrios Buhalis (eds.) (2000): *Information and communication technologies in tourism*. Springer, Vienna, Austria.
- Frohlich, D.M., A. Kuchinsky, Pering C., Don A. and Ariss S. (2002): Requirements for photoware. In: *Proceedings of CSCW '02*. New York: ACM Press.
- Gaver, W. (1992): The Affordances of Media Spaces for Collaboration. In: *Proceedings of Computer Supported Cooperative Work, CSCW'92. Toronto, Canada*. ACM Press, pp. 17-24.
- Gaver, William W., Abigail Sellen, Christian Heath and Paul Luff (1993): One is not enough: multiple views in a media space. In: *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM Press, pp. 335--341.
- Greenhalgh, C. (2002): EQUIP: a Software Platform for Distributed Interactive Systems. In: *Equator Technical Report 02-002*. Nottingham: University of Nottingham.
- Grinter, Rebecca E., Paul M. Aoki, Margaret H. Szymanski, James D. Thornton, Allison Woodruff and Amy Hurst (2002): Revisiting the visit: understanding how technology can shape the museum visit. In: *Proceedings of ACM Conference on Computer Supported Cooperative Work (CSCW)*. ACM Press, pp. 146-155.
- Kuzuoka, Hideaki, Toshio Kosuge and Masatomo Tanaka (1994): GestureCam: a video communication system for sympathetic remote collaboration. In: *Proceedings of the 1994 ACM conference on Computer supported cooperative work*. ACM Press, pp. 35--43.
- Montello, Daniel (2003): Navigation. In: Shah, P. and Miyake, A. (eds): *Handbook of visuospatial cognition*. Cambridge: Cambridge University Press.
- Paulos, Eric and John Canny (1997): Ubiquitous tele-embodiment: applications and implications. *Int. J. Hum.-Comput. Stud.*, Vol. 46, no. 6, pp. 861--877.
- Sacks, Harvey (1995): *Lectures on conversation: vol 1 & 2*. Oxford: Basil Blackwell.
- Simmel, Georg and Everett C. Hughes (1949): The sociology of sociability. *The American Journal of Sociology*, Vol. 55, no. 3, pp. 254-261.