presented to users, taking advantage of the logs from previous visits and 'pre-visits' to assist the current visit. In conjunction with the other features of the system, the use of this past information allows us to 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 and setup.

PREVIOUS WORK

City visiting has been a popular area for mobile information systems, in particular [7], and other PDA based systems [1, 8, 25]. Indeed, as mobile phones and other portable devices become more advanced, tourism seems to be an obvious application area. A number of phone operators have already released city guides for easy viewing that are targeted and customised for mobile phones (for example, http://www.lonelyplanet.com/mobile/). However, these and other commercial technologies have had only limited success. Generally, they 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. This type of model can often seem static and leave the user feeling that the system is not greatly interactive – that they have little input or control and that they are very much working in an isolated environment. There has been little explicit support for *collaboration* between visitors.

One notable exception in this regard was *Sotto Voce,* which allowed museum visitors to share a spoken commentary as they visited a historic house [25]. A small number of mobile systems designed for entertainment and games also specifically address collaboration. *Can You See Me Now*, for example, was a performance that employed a game format. It incorporated multiple players using wireless–enabled PDAs on city streets, who were in turn connected to online players via the Internet [9]. A recent commercial mobile game that relies on collaboration is Newt Games' *Mogi* (www.mogimogi.com)*,* which involved finding and trading objects in city streets.

Similarly, while recommendation systems usually generate recommendations by combining records of several people's past activity, collaboration has seldom been a central focus. PolyLens [19] was one recommender which worked for groups, in that it allowed two or more people to combine their movie rating profiles into one, and then create one recommendation list from this. Also, recommenders rarely use a broad set of contextual features, although the Jimminy system [21] was one temporally–specific single–user recommender that used explicitly–entered textual notes, and the names of locations and people, as contextual features to base recommendations on.

SYSTEM OVERVIEW

In the George Square system (Figure 1), each tourist can visit the physical 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 (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. 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 on a shared 'filmstrip' view, 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. This historical information is run through the Recer collaborative filtering algorithm [5] 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 own map, and sees others' recommendation lists alongside one's own (also ghosted for easy distinction). 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.

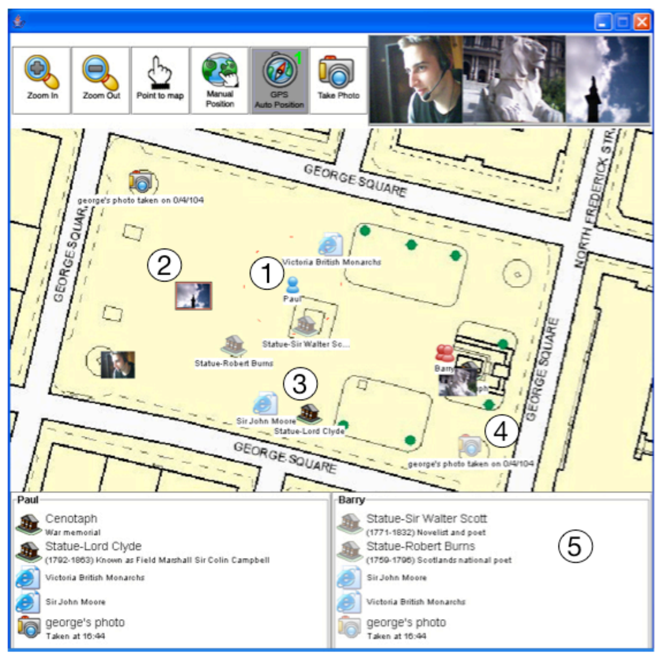

Figure 1: 'George Square' system, showing map that displays each user's location (1), thumbnail photos (2), recommended locations, web pages (3) and photos (4), and each users' recommendation list (5).

In use, the system supports a range of different scenarios. Firstly, it can support two users collaboratively co–visiting an area of a city, taking photographs and browsing web pages about that area. Secondly, it can support users physically present at the location collaboratively co–visiting with other users distant from the where they are, via the Internet. Thirdly, users who are all distant from the area but interacting via the Internet can use the system to share a purely online visit. The latter scenario is important as this is often 'pre-visiting' in which people explore photographs, web pages and attractions that are of interest before they actually arrive at a city. Vital historical data that can feed into their later activity of the actual visit may be recorded at this pre-visit stage.

In a complementary way, we support post–visit activity. The database log generated from earlier visiting is used to generate a web page: a travel weblog [2]. One can browse the web pages generated from one's visit, viewing 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. This summarises one's visit in a spatial presentation (the post-visit 'web-log' is discussed in more detail in future papers).

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 – content does not need to be produced, as it will automatically accumulate from usage of the system. Furthermore, if the system is continually run by waves of visitors then the content will always remain relatively up-to-date as users continue to generate new logs.

The implementation challenges for George Square were typical of other collaborative mobile systems, in that we needed a mix of devices that could work together as peers without relying on access to a central server. We also wanted our system to be dynamic, supporting users and devices joining and leaving at any time.

The hardware of our system consists of a lightweight Tablet PC with attached compact flash GPS unit and a USB 'stalk' camera. Headphones and microphone were plugged into the unit, and the built in WiFi was used for communications. In our trials, a temporary wireless network was bridged to a publicly available WiFi 'hotspot' to provide Internet access. This allowed users to browse and search the web, and to follow links to information provided by our system.

For our software we expanded on previous work with the EQUIP distributed tuple space systems [14], middleware which supports a peer-to-peer communication model between networks of sensors and output devices. 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 on the same system and network communication to components on other systems, 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 reconfiguration. Events describing user activity and sensor readings are recorded by logging components. These logging components also continually run algorithms comparing recent activity with historical logs, to create recommendations.

USER 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. Other papers (under review) report on more general details of interaction with George Square, but here we summarise results related to the use of logged information and recommendations.

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.

Figure 2: A co-visit with one user physically in the George Square using a tablet PC and one indoors visitor using a laptop to share the visit.

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 scenario we gave for the trial was of two friends sharing a visit to George Square, communicating via the system. For the first half of the trial, participants were asked to freely explore the square learning how to use the system. For the second half of the trial, users were given a set of tasks to carry out. This included tasks such as sharing a photograph of the square, and finding out the height of the statue in the centre of the square.

A range of data from each trial was collected: video tapes of both the indoor and outdoor visitors, 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 of into a single video stream. From the logs, we generated a 'playback' 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* [15], 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 events. In particular, situations where participants were confused revealed where the system could be improved to better support collaboration or understanding.

In use, the system presented a novel yet enjoyable experience for trial participants, with all participants exchanging photographs, and using their location and recommendations in their interactions around the square. While exactly the same software was used for both indoor and outdoor participants, differences in the visitor's situation resulted in different capabilities for 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 whilst 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 that happened out in the square. These differences in situation led to clear patterns of use and division of labor in the trial. The indoor user would search 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". However, some web pages were browsed by the outdoor user, since the recommendation system allowed browsing of recommended web pages without having to type in URLs or search terms. These results were confirmed by our analysis of the videos.

The system offered a range of different resources that visitors could use to share the visit: location (displayed on a map), voice, photographs, recommendations and web pages. These different resources supported collaboration between visitors in different ways, but 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, e.g:

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

Of all the resources provided by the system, the voice connection proved to be the most valuable for creating a sense of shared experience. Through their talk, users continually managed their shared experience, talking about what they were doing, what they had done and what they were going to do. As emphasised in similar studies [9], voice is an essential tool for repairing misunderstandings.

The recommendations of *web pages* acted as an effective way of displaying and linking together the online content available about places, with the place itself. Although we 'bootstrapped' the recommender system by browsing web pages in appropriate places, the system also recommended pages that had been browsed by users during the trials. One early trial participant browsed the 'wikipedia' pages about William Gladstone, which were then recommended to later trial participants who went to the statue of Gladstone. Recommended web pages, positioned on the map, acted as geographical 'bookmarks' in the square being visited, taken from other people's web browsing. 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.

Along with webpages, our system also recommended sets of places. Places' labels provided the names of different statues in the square, as well as those of buildings on the edge of the square. However, rather than only acting as recommendations of where to go next, these labels acted as labels 'seen in common', which could be used when talking about different parts of the square in sociable or functional ways. The indoor user, for example, could ask the outdoor user to go to a particular attraction by using its name. At times, this conflicted with recommendations' role as suggestions of where to go or what to read next. As a visitor got close to a recommended place, that label disappeared because, from an information–seeking point of view, there was no longer any need to suggest it. However, from a conversational point of view, the shared label for that place was then unavailable as a resource, causing disruption.

DESIGN IMPLICATIONS

One problem with electronic maps, and visualisations more generally, is the need to keep the display clear from irrelevant details. As each visitor in our trial navigated the square, his or her recent behaviour was used to filter the items displayed. Current behaviour was used by our recommendation algorithm to filter and select items from historical data of use. The use of these labels as conversational resources by users suggests how recommenders can be used to filter information displayed on maps in a contextually appropriate manner.

In addition, our recommender also made use of historical data to weave together online information with physical places. Photographs taken, and web pages browsed, by users were stored as an archive of information about the locations the system was used. These associations were not pre-authored but rather evolved with users' behaviour. This exemplifies how historical data can be a resource for connecting online data with different places.

Both these applications show the value of using recommender algorithms to support collaboration. While tensions exist between single user and collaborative use, we would argue that recommenders and other information seeking tools can be powerful used to support new forms of collaboration.

CONCLUSION

This paper has presented 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. It provides resources for sharing voice, photos, location and web pages. A trial of the system uncovered how, through the different resources the system provided, visitors could accomplish a shared visit. In particular, users brought together their shared location, voice, photographs and recommendations to co-ordinate and enjoy a visit.

Ubicomp technology offers the possibility of access to large bodies of information on distant servers and stores, through information-seeking tools such as search engines and recommenders. As well as access to distant information we have shown how it can provide access to distant people and past activity. Collaborative ubicomp can integrate interaction with the local context with the social context of collaborators far away, and historical context in terms of contributors from the past. We suggest that there is rich potential in combining information from near and far, from the past and the present, and from the wide range of tools and media that collaborative ubicomp employs. 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.

REFERENCES

[1] Abowd, G.D., Atkeson, C.G., et al. Cyberguide: A Mobile Context-Aware Tour Guide. *ACM Wireless Networks,* 3(5), 1997, 421-433.

[2] Brown, B. and Chalmers, M. Tourism and Mobile Technology. *Proc. ECSCW 2003*, Kluwer, 335-355.

[3] Brown, B. and Laurier, E. Maps & Journeying: An Ethnomethodological Approach. To appear in *Cartographica*.

[4] Brown, B., MacColl, I., et al. Lessons from the Lighthouse: Collaboration in a Shared Mixed Reality System. *Proc. CHI 2003,* ACM Press, 577-585.

[5] Chalmers, M. When Cookies Aren't Enough: Tracking and Enriching Web Activity with Recer. In: *Preferred Placement: Knowledge Politics on the Web*, Rogers, R. (ed.), Jan van Eyck Academie, 2000, 99-102.

[6] Chalmers, M. and Galani, A. Seamful Interweaving: Heterogeneity in the Theory and Design of Interactive Systems. *Proc. DIS 2004,* ACM Press, 243–252.

[7] Cheverst, K., Davies, N., et al. Developing a Context-Aware Electronic Tourist Guide: Some Issues and Experiences. *Proc. CHI 2000*, ACM Press, 17–24.

[8] Fesenmaier, D., Klein, S., et al. (eds.). *Information and Communication Technologies in Tourism*. Springer, 2000.

[9] Flintham, M., Anastasi, R., et al. Where On-Line Meets On-the-Streets: Experiences with Mobile Mixed Reality Games. *Proc. CHI 2003,* ACM Press, 569-576.

[10] Frohlich, D.M., Kuchinsky, A., et al. Requirements for Photoware. *Proc. CSCW 2002,* ACM Press, 166–175.

[11] Galani, A. and Chalmers, M. Production of Pace as Collaborative Activity. *Extended Abstracts of CHI 2004,* ACM Press, 1417–1420.

[12] Gaver, W. The Affordances of Media Spaces for Collaboration. *Proc. CSCW 1992,* ACM Press, 17-24.

[13] Gaver, W.W., Sellen, A., et al. One Is Not Enough: Multiple Views in a Media Space. *Proc. CHI 1993*, ACM Press, 335–341.

[14] Greenhalgh, C. Equip: A Software Platform for Distributed Interactive Systems. *Equator Technical Report 02-002,* University of Nottingham, 2002

[15] Heath, C. and Luff, P. *Technology in Action*. Cambridge University Press, Cambridge, 2000.

[16] Jefferson, G. Transcript Notation. In: Atkinson, J.M. and Heritage, J., eds. *Structures of Social Action: Studies in Conversation Analysis*, Cambridge University Press, 1984, ix-xvi.

[17] Kuzuoka, H., Kosuge, T., et al. Gesturecam: A Video Communication System for Sympathetic Remote Collaboration. *Proc. CSCW 1994*, ACM Press, 35–43.

[18] Kuzuoka, H., Oyama, S., et al. Gestureman: A Mobile Robot That Embodies a Remote Instructor's Actions. *Proc. CSCW 2000*, ACM Press, 155–162.

[19] O' Conner, M., Cosley, D., et al. Polylens: A Recommender System for Groups of Users. *Proc. ECSCW 2001*, Kluwer, 199–210.

[20] Paulos, E. and Canny, J. Ubiquitous Tele-Embodiment: Applications and Implications. *Int. J. Human-Computer Studies,* 46 (6). 861--877.

[21] Rhodes, B. Using Physical Context for Just-in-Time Information Retrieval. *IEEE Trans. Computers*, 52 (8), 1012–1014, 2003.

[22] Sacks, H. *Lectures on Conversation: Vol 1 & 2.* Basil Blackwell, Oxford, 1995.

[23] Simmel, G. and Hughes, E.C. The Sociology of Sociability. *American Journal of Sociology,* 55(3), 254-261.

[24] Urry, J. *The Tourist Gaze: Leisure and Travel in Contemporary Societies*. Sage, New York, 1990.

[25] Woodruff, A., Aoki, P.M., et al. Electronic Guidebooks and Visitor Attention. *Proc. Int'l Cultural Heritage Informatics Meeting*, Milan, 2001, 437-454.