

Shared visiting in EQUATOR City

Ian MacColl¹, Barry Brown¹, Steve Benford⁴, Matthew Chalmers¹, Ruth Conroy³, Nick Dalton³, Areti Galani¹, Chris Greenhalgh⁴, Danus Michaelides⁵, Dave Millard⁵, Cliff Randell², Anthony Steed³, Tom Rodden⁴, Ian Taylor⁴, and Mark Weal⁵

¹ The University of Glasgow, UK

² The University of Bristol, UK,

³ University College London, UK

⁴ The University of Nottingham, UK

⁵ The University of Southampton, UK

<http://www.equator.ac.uk/projects/city/>

Abstract. In this paper we describe a system and infrastructure for sharing of visiting experiences across multiple media. The prototype system supports synchronous visiting by both physical and digital visitors, with digital access via either the World Wide Web or 3-dimensional graphics, and we are extending it to support asynchronous visiting in the form of annotations and recommendations.

1 Introduction

In this paper we describe a prototype interactive system and technical infrastructure supporting shared experience by physical and digital visitors to an exhibition devoted to the artist, designer and architect Charles Rennie Mackintosh. The system supports synchronous sharing, and asynchronous sharing is under development. Synchronous sharing involves shared spatial and informational awareness, with collaboration via audio. Asynchronous sharing involves spatial and informational annotations and recommendations derived from previous visitors.

The core of the prototype is the EQUIP platform, supported by a variety of clients and services. EQUIP is a shared information service, extending previous work on tuple spaces. EQUIP has been integrated with a variety of other services, including hypermedia systems and collaborative environments, and is used as the basis for a blackboard system architecture.

The prototype has been developed within the City project, which is exploring ways to treat the city and information so as to deliberately blur the boundaries between physical and digital media. City is a project of the EQUATOR Interdisciplinary Research Centre, which is focusing on bridging physical and digital worlds through collaborative research by computer scientists, psychologists, sociologists and designers.

In the next section we describe the City project, in sections 3 and 4 we describe the City infrastructure and system, and section 5 concludes.

2 City

In this section we provide details on the City project, including its practical, technical and theoretical aims, the environment in which we are working and the development process we have used.

Practically or politically, we aim to provide an environment for collaborative research by some of the eight EQUATOR partners. University College London (UCL) and the universities of Bristol, Glasgow, Nottingham and Southampton are taking part in the City project, with Glasgow as lead. Researchers from the fields of computer science, sociology and museum studies are heavily involved, with some input from psychologists and designers.

Technically, the City project aims to explore new devices and systems supporting mobility, contextual adaptation and access to mixed information types. Physical visitors use handheld and wearable devices, while digital visitors access the exhibition via the World Wide Web or 3-dimensional graphics ('virtual reality'). The prototype provides location- and device-aware content, maintaining and presenting spatial and informational 'worlds' which provide a basis for collaboration and sharing.

Theoretically, we are working to weave digital media into the physical streets, buildings and artefacts that people use, and to do this in meaningful ways, that is, ways that fit, show and support their activity. We all tend to focus on obvious differences between physical and digital media and treat each one independently, but in the City project we are seeking a broader viewpoint that takes account of the similarities and interdependencies of different media.

The initial setting of our work is the Mackintosh Interpretation Centre in The Lighthouse (<http://www.thelighthouse.co.uk>), Scotland's Centre for Design, Architecture and the City. The Interpretation Centre is devoted to the life and work of Charles Rennie Mackintosh (1868-1928), Glasgow architect, designer and artist. The Centre is comprised of textual and graphical displays with some artifacts, and includes over 20 screens presenting video and interactive material. Our work is informed by observational studies of visitors [10], and background research on the design and use of both the Centre and The Lighthouse-Mackintosh's first major public building.

Major system development has been conducted in week-long workshops fostering collaboration between partners with diverse philosophies and approaches. These range from those interested in (for example) standalone wearable and handheld devices, emphasising power-conservation and network disconnection, to those interested in middleware for large hypermedia systems and collaborative virtual environments, with high bandwidth connections and uninterruptible power supplies. Informal demonstrations scheduled at the end of each workshop provide a strong motivation to produce a viable result. Between workshops we meet weekly on an Internet Relay Chat (IRC) channel to discuss ongoing work.

We use scenarios as the basis of our work, supporting the development of shared vocabulary and understandings. The scenarios range from very general, such as

Earlier today, while on the train to Glasgow, Vee used her mobile to take a look at tourist information about the city...

to very detailed, such as

There are two visitors: Vee, the physical visitor, is in the ReaCTor ante-room, and Ana, the digital visitor, is inside the ReaCTor. For each there are models of two Mackintosh buildings with associated textual commentaries. They are asked to inspect and discuss the two models. The two textual commentaries for each object are different they each have some unique information to impart to the other.

The characters Vee and Ana, and their companion on the World Wide Web, Dub, appear consistently in the scenarios and their names have been carried through into the code and configuration files. Scenario fragments are interwoven with explorations of technical options, informally similar to the user stories and task cards of eXtreme Programming [3].

The scenarios describe the behaviour of our prototype system in terms that participants from diverse disciplines and backgrounds can relate to and understand. In the next section we describe our infrastructure and in the following section we describe the system itself.

3 Infrastructure

In this section we describe the infrastructure supporting our prototype system: the EQUIP platform, the Auld Leaky contextual link server, the VR Juggler framework, and the Bristol ultrasonic positioning system.

3.1 EQUIP

The EQUIP platform is being developed within the EQUATOR IRC as an adaptive infrastructure to support information sharing between heterogeneous devices. The University of Nottingham is leading development of EQUIP, with contributions from the various EQUATOR projects, including City. EQUIP extends previous work on tuple spaces [7, 11], and aims to:

- provide a run-time infrastructure to address multiple domains;
- support interoperation between Java and C++ using a subset of Common Object Request Broker Architecture (CORBA) Interface Definition Language (IDL); and
- support extensibility, including dynamic loading of C++ classes.

Interoperability and extensibility are achieved by using Bamboo [18] at both build- and run-time. Bamboo, in turn, is implemented on top of the Netscape Portable Runtime (NSPR) [16] providing a platform-neutral API for system level facilities: threads, synchronisation, file and network I/O, time, memory management and shared library linking.

EQUIP is comprised of Bamboo modules that can be divided into core services and additional facilities. The core modules are:

- runtime: system facilities including memory management and base classes
- net: support for simple clients and servers, and a trader/nameservice
- Trader: the trader service
- data: information sharing service supporting tuple space operations and pattern-based matching
- Server: the data service

In the City project, EQUIP is used as the basis for a blackboard architecture. Data items representing user context, an underlying spatial model and context-aware content are stored in EQUIP for manipulation by City clients and services.

Additional EQUIP facilities support specific application domains, particularly real-time 3-dimensional graphics, for example point, quaternion and matrix classes and operations (math module), and abstract, renderable scene graph nodes (draw3d module). In addition, interfaces between EQUIP and a number of other systems have been developed, including Nottingham's MASSIVE-3 collaborative virtual environment system [12].

3.2 VR Juggler

In addition to MASSIVE-3, the University of Iowa's VR Juggler [4] has also been integrated with EQUIP, and is used as the renderer for 3-dimensional graphics in the City project. VR Juggler is described as a virtual platform for virtual reality (VR) application development. It is a high-level, cross-platform architecture supporting immersive and non-immersive presentations. Both UCL and Nottingham have immersive projection facilities, and the UCL facility has been used for development and trials in the City project.

The 3-dimensional graphics rendering provided by VR Juggler is used in the City project to provide an analogue to the space visited by physical visitors. For World Wide Web visitors, the space is represented as a 2-dimensional map. The digital visitors also require a presentation of the informational space available to physical visitors in the form of displays consisting of graphics, text, screens and artifacts. This is served by Auld Leaky.

3.3 Auld Leaky

Auld Leaky [15] is a lightweight contextual link server being developed within the EQUATOR IRC to store and serve hypermedia structures, using context to filter query results. The model used to define the structures is capable of representing a variety of hypermedia domains: navigational, spatial and taxonomic.

The University of Southampton is developing Auld Leaky, and, as part of the City project, it has been integrated with EQUIP (as described below). Auld Leaky is implemented in Perl, and has a Java API which is used in the City project. Information is encoded as an eXtended Markup Language (XML) *linkbase*, loaded into Auld Leaky and queried using the HyperText Transfer Protocol (HTTP).

Auld Leaky supports contextual queries which are used in the City project to generate location- and device-aware content to be delivered by the Apache World Wide Web server and servlet engine. For digital visitors, a location can easily be derived from the precise position of a visitor representation on a 2- or 3-dimensional rendering (described below). A system from the University of Bristol is used to provide position information for physical visitors.

3.4 Bristol ultrasonics

The University of Bristol is developing a low cost indoor positioning system using a combination of radio frequency (RF) and ultrasonics [17] as part of its contribution to the EQUATOR IRC. The system uses a single RF transmitter for synchronisation, with ceiling-mounted ultrasonic transmitters for positioning. The ultrasonics transmit at known intervals after the RF, and are received by a handheld or wearable. The variations in flight time of the ultrasonic transmissions are used to calculate the spatial position of the receiver. The receiver incorporates a magnetic compass to provide orientation information.

The City project ultrasonic system is the largest deployed to date, involving eight ultrasonic transmitters covering the Mackintosh Interpretation Centre which is approximately 10m by 20m. The Centre is a challenging environment in other ways, split into two large areas by a partial wall, with some areas set up as cubicles, covered by roofs (making ultrasonic reception virtually impossible). The Centre includes surfaces with varying acoustic absorption and, in addition, for aesthetic and coverage reasons, the ultrasonic transmissions are reflected off the ceiling. The receiver's position is calculated each second, and measurements indicate that 95% of the calculated positions are within 2.5m of the actual position.

The EQUIP platform is used as the shared space in which information from and/or for VR Juggler, Auld Leaky and the Bristol ultrasonics is stored. In the next section we detail how this occurs.

4 System

In this section we describe the City system in detail. The system supports shared visiting experiences for physical and digital visitors, with digital visitors using the World Wide Web or 3-dimensional graphics. In discussion and development scenarios we name the visitors Vee, Dub and Ana respectively, which we will also use in the remainder of this paper. The synchronous visiting components of the prototype system operate similarly for each visitor, broadly as follows:

- store spatial position and orientation in EQUIP;
- retrieve and render positions of other visitors;
- store named location in EQUIP in response to position change;
- store content from Auld Leaky in EQUIP in response to location change;
- format content for presentation and advise client program of availability in response to content change.



Fig. 1. 3-D rendering presented to Ana showing avatars for Vee (left) and Dub.

In the remainder of this section we detail how we support spatial and informational awareness for synchronous sharing of visiting experiences, and how we are implementing annotation and recommendation systems for asynchronous sharing of visiting experiences. Audio support for synchronous visiting is handled by a discrete subsystem that we will not detail in this paper.

4.1 Spatial worlds

Spatial awareness for synchronous visiting is supported by overlaying a physical space with a corresponding digital model of the physical space. We send position and orientation information for each visitor to an EQUIP data space containing the digital model and dedicated to the City system. The positions and orientations of the other visitors are retrieved from the data space and rendered for each visitor.

For Ana, position and orientation information is automatically published in the EQUIP data space by the VR Juggler client, and the positions of all visitors stored in the data space are automatically rendered as 3-dimensional avatars. The complete 3-dimensional spatial model for rendering and other services is loaded into the data space on starting the system. Figure 1 shows a non-immersive spatial awareness display for Ana, with avatars representing Vee and Dub (displaying only heads rather than complete avatars).



Fig. 2. Three physical visitors with Vee (right) wearing Bristol *CyberJacket*.

Vee uses a Jornada 568 that polls position and orientation sensors, and sends the results to a proxy for the City EQUIP data space, using a custom protocol over a TCP socket. The protocol also supports sending of visitor information to the Jornada from the proxy. The proxy is responsible for inserting Vee's information into the data space, and retrieving the information for the other visitors. If necessary, the proxy transforms Vee's spatial reference frame to that of the spatial model in the data space—for example Vee's magnetic compass orientation (clockwise) is transformed to Cartesian orientation (anti-clockwise). The positions and orientations of all visitors are presented to Vee on a 2-dimensional map. Figure 2 shows three physical visitors with Vee (on the right) wearing a University of Bristol *CyberJacket* (with a Jornada 690 rather than a Jornada 568). The figure also gives some sense of the visual richness of the physical information in the Centre.

Dub interacts with a Java applet in a World Wide Web browser frame. The applet communicates with the same proxy as Vee's client, and uses the same protocol, converting mouse clicks on a 2-dimensional map of the Interpretation Centre to position and orientation information. The applet also displays representations of all visitors. As for Vee, the proxy transforms Dub's (version of Vee's) reference frame to that of the spatial model in EQUIP. An example of Dub's map is shown in Figure 3, corresponding to Ana's 3-dimensional display



Fig. 3. 2-dimensional map presented to Dub, showing Vee and Ana in blue.

in Figure 1. The red boxes on Dub's map are trigger zones, discussed in the next section. Vee's map is similar to, but simpler than, Dub's.

4.2 Informational worlds

Shared visiting requires both a sense of spatial proximity, and some comparability of the information available to each visitor. Physical visitors to the Centre have an extremely visually rich information environment as shown in Figure 2. As outlined above, providing comparable information to digital visitors involves converting visitor positions to named locations, querying Auld Leaky using the visitor's device and location to generate informational content, and then formatting and presenting the content.

In initial trials of our prototype system we do not intend to deliver informational content to Vee, due to the high information content of the physical en-

vironment. Our implementation of asynchronous visiting experiences, described below, will deliver such content, however.

Positions are converted to locations by an EQUIP-aware collision-detection service. Axis-aligned cuboids are inserted into the data space at startup and are monitored as sensors by a server (named BBox). These represent semantically significant volumes or extents within the spatial model (shown as red outlines on Dub's 2-dimensional map in Figure 3). Also, for each visitor a target is inserted into the EQUIP data space, currently equivalent to a 10cm cube held in the hand. Detection of a collision between a target and a sensor invokes code that inserts a new item (defined as IDL) into the data space. This item maps the name of the target's user to the location name (which is also the name of the cuboid sensor).

Adding or updating a user-location item in the data space triggers a query to Auld Leaky and the results are stored in the data space. An EQUIP client registers interest in such changes using typical tuple space template-based pattern matching, and contains the code to respond to the change. The results returned by Auld Leaky are arbitrary hypermedia fragments (URLs, HTML fragments, etc), contained within City-specific structures: Links, Concepts and Annotations. Links model explanations, mapping a location name to some explanatory content, and transitions, mapping a source location to a destination location with associated content. Concepts group together multiple, contextually distinguished representations of some content. Annotations will be described below.

Adding or updating a content item in the data space triggers formatting for delivery to the visitor. The proxy mentioned above for Vee's and Dub's positions registers interest in such changes, and retrieves content fragments and combines them into an HTML page. The proxy protocol includes an optional notification of content availability. Dub's applet responds to this notification by displaying the HTML page in a separate browser frame set aside for this purpose. Ana also runs a browser and applet but the applet simply connects to the proxy and, when advised, displays HTML pages (that is, it has no map display).

The overall effect, then, is that Dub and Ana visit within spatial and informational worlds, with the informational world supdated as they move around their spatial worlds. We believe that this, with shared audio, provides sufficient commonality of experience to support a sense of shared visiting. We are planning to evaluate this prototype system in user trials in the very near future, informed by our existing studies of museum visitors [10], and tourists [5]. In parallel with these trials we are implementing asynchronous visiting experiences as outlined in the next section.

4.3 Annotations and Recommendations

There has been considerable work on generation and delivery of static, possibly contextually-selected, content to visitors of various kinds, for example [2, 8, 19]. Rather than covering the same ground, we are choosing to open the content available to visitors for annotation and appropriation by the visitors themselves. As for tour guides, there have been a number of location-aware annotation systems,

for example [9], but few in a museum context. Visitors books have, of course, long been a feature of visiting cultural institutions, and a number of museums include anonymous, electronic commenting systems. Our interest, however, is annotations that approach the authority of the displays created by curators and designers, raising issues of provenance and persistence. In a sense we are seeking to push the granularity of the visitors book down to the level of individual display or artifacts.

As mentioned above, we have developed Auld Leaky structures to model Annotations, mapping a location to an annotation with an associated author. The structures support a simple classification of annotations as questions, opinions, and so on. In addition to defining the structures, we are also adding simple checkpoint-based persistence to Auld Leaky to ensure that annotations persist. The next step is to provide location-aware user interaction to elicit and store annotations, and mechanisms to browse and search (or recommend) stored annotations.

More fundamentally, the City project aims to explore path-based recommending of spatial and informational locations in the Centre and the wider city. The Recer system [6] logs sequences of information activities, for example World Wide Web URL accesses. Recer takes recent accesses as an implicit query, and offers ranked recommendations based on closeness of match to the logged accesses. The recommendations can be filtered or selected in various ways: statistically, socially, etc.

We are extending our prototype system to incorporate Recer to provide spatial and informational recommendations. We are implementing logging of all spatial and informational location changes. We will then extend Recer to respond to location changes in the EQUIP data space by generating a list of recommendations and storing them in the data space for subsequent presentation to visitors. Recer is essentially a system for manipulating and selecting from arbitrary sequences of abstract symbols, so as we enrich the contextual and other information stored in the City data space, the Recer paths and recommendations will also be enriched. One aim of this strand of ongoing work is to move towards revealing the patterns of past activity in the various media that change a space into a place [13], by showing the patterns of use graphically in maps or 3-dimensional graphics, and using those patterns as the basis for recommendations.

5 Conclusion

In this paper we have described our system for sharing of visiting experiences across multiple media, and the various infrastructures we are developing to support and implement it. The prototype system supports synchronous visiting by physical and digital visitors, and we are extending it to support asynchronous visiting in the form of annotations and recommendations.

As we noted above, while there is much existing work on location-aware content generation in museum and gallery environments, we are working in quite

different directions. We are not aware of any other project exploring shared synchronous and asynchronous visiting experiences. The closest related work is outside the cultural heritage domain in the Cooltown [14] and Sentient Computing [1] projects.

Cooltown essentially aims to support World Wide Web presence for people, places and things. In contrast, we use shared spatial and informational worlds, supporting sharing of concrete or literal spatial experiences through to abstract or symbolic informational ones. Our focus is collaboration and sharing between users, in contrast to Cooltown's emphasis on information and services accessed or manipulated by users.

The Sentient Computing project is technically much closer to City, involving the use of sensors for context-aware services, with an underlying digital model representing physical spaces and objects. The project has been running for longer than City and is considerably more advanced. As with Cooltown, however, a major motivation is access to and manipulation of digital information and services, albeit by novel user interaction techniques. Our focus on collaboration and sharing between users is distinctive.

We have conducted limited public demonstrations of our system in the Mackintosh Interpretation Centre, and we have used spatial models of two of our laboratory spaces to simulate the room at other locations in the UK. We are preparing for larger-scale user trials in the Centre, initially focussing on synchronous visiting.

Our limited demonstrations have already provided some interesting results. We have observed that visitors in the roles of Vee and Ana, in particular, carefully manage their bodily positions with respect to each other, as in conventional collaborative virtual environments, and show little problem with mapping between physical objects and those in the 3-dimensional graphics. Talking is vital to the support of mutual reference. Our plans for shared audio take account of these observations, as well as recent work on the sharing and overlapping of audio guides [19].

The dangers of making ambitious contextual inferences have been made apparent by our initial observations. Our semantic location names and zones are derived from those identified by the designer of the Centre. However, there are several places in the room where one can stand in a position but, depending on the direction one faces, look at different exhibits and zones. Some exhibits can only be accessed from outside the zone that contains them. Orientation data helps to determine a visitor's focus of attention, but not always. A visitor may be facing one exhibit while talking about another with his or her companions. We have also observed visitors continuing to listen to an exhibit's audio description after moving on to look at other exhibits. Location, orientation and devices at hand, as well as who is where, who is reading what and who is talking to whom, are just some of the simpler features we can incorporate to enrich our model of a user's context.

In addition to the ongoing work to support asynchronous sharing of visits, we have a number of other extensions under discussion. These range from the

relatively simple, such as evaluating preferences for spatial or informational navigation and content presentation, comparing map- and radar-based awareness displays and improving collision detection, to substantial projects such as escaping into the wider city, dynamical generating location taxonomies or ontologies, and modeling the uncertainty associated with physical sensing.

We have designed and built a prototype system for shared physical and digital visiting experiences. The system provides a technical infrastructure for social awareness across different media with context-sensitive content generation and delivery. This combination of shared, multiple media visiting with context-sensitive content is unique as far as we know. The system is the work of five EQUATOR partners, and it exploits and demonstrates several EQUATOR infrastructure technologies, namely the EQUIP data service, the Auld Leaky link server and the Bristol ultrasonic positioning system.

Acknowledgments

The EQUATOR IRC is funded by the UK Engineering and Physical Sciences Research Council, and the City project is supported by a donation from the Hewlett-Packard Art & Science programme. We are also grateful to The Light-house for the opportunity to test our ideas and our systems.

References

1. Addlesee, M., Curwen, R., Hodges, S., Newman, J., Steggles, P., Ward, A., Hopper A.: Implementing a sentient computing system. *Computer*, **34.8**, 2001, 50–56.
2. Abowd, G., Atkeson, C., Hong, J., Long, S., Kooper, R., Pinkerton, M.: Cyberguide: A mobile context-aware tour guide. *Wireless Networks*, **3.5**, 1997, 421–433.
3. Beck, K.: Embracing change with Extreme Programming. *Computer*, **32.10**, 1999, 70–77.
4. Bierbaum, A., Just, C., Hartling, P., Meinert, K., Baker, A., Cruz-Neira, C.: VR Juggler: A virtual platform for virtual reality application development. In *Virtual Reality*, 2001, IEEE, 89–96.
5. Brown, B., Chalmers, M., MacColl, I.: Exploring tourism as a collaborative activity. In preparation.
6. Chalmers, M.: Paths and Contextually Specific Recommendations. In *DELOS/NSF Workshop on Personalisation and Recommender Systems*, 2001.
7. Davies, N., Wade, S., Friday, A., Blair, G.: Limbo: A tuple space based platform for adaptive mobile applications. In *Open Distributed Processing/Distributed Platforms*, 1997, 291–302.
8. Davies, N., Cheverst, K., Mitchell, K., Efrat, A.: Using and determining location in a context-sensitive tour guide. *Computer*, **34.8**, 2001, 35–41.
9. Espinoza, F., Persson, P., Sandin, A., Nyström, H., Cacciatore, E., Bylund, M.: GeoNotes: Social and Navigational Aspects of Location-Based Information Systems. In *Ubiquitous Computing*, 2001, Springer-Verlag, 2–17.
10. Galani, A., Chalmers, M.: Can you see me? Exploring co-visiting between physical and virtual visitors. In *Museums and the Web*, 2002, Archives and Museum Informatics, to appear.

11. Gelernter, D.: Generative communication in Linda. *ACM Transactions on Programming Languages and Systems*, **7.1**, 1985, 80–112.
12. Greenhalgh, C., Purbrick J., Snowden, D.: Inside MASSIVE-3: Flexible support for data Consistency and World Structuring Collaborative Virtual Environments, 2000, ACM Press, 119-127.
13. Harrison, S., P. Dourish.: RePlacing Space: The Roles of Place and Space in Collaborative Systems, in *Computer Supported Collaborative Work*, 1996, ACM Press, 67–76.
14. Kindberg, T., Barton, J., Morgan, J., Becker, G., Bedner, I., Caswell, D., Debaty, P., Gopal, G, Frid, M., Krishnan, V., Morris, H., Pering, C., Schettino, J., Serra, W., Spasojevic, M.: People, places, things: Web presence for the real world. In *Mobile Computing Systems and Applications*, 2000.
15. Michaelides, D., et al.: Auld Leaky: A Contextual Open Hypermedia Link Server. In *Open Hypermedia Systems*, 2001, Springer-Verlag.
16. Mozilla: NetScape Portable Runtime. <http://www.mozilla.org/projects/nspr/>.
17. Randell, C., Muller, H.: Low Cost Indoor Positioning System In *Ubiquitous Computing*, 2001, Springer-Verlag, 42–48.
18. Watsen, K.: Bamboo. <http://watsen.net/bamboo/>.
19. Woodruff, A., Grinter, B.: Sotto Voce: Exploring the Interplay of Conversation and Mobile Audio Spaces. In *Human Factors in Computing Systems*, 2002, ACM Press, to appear.