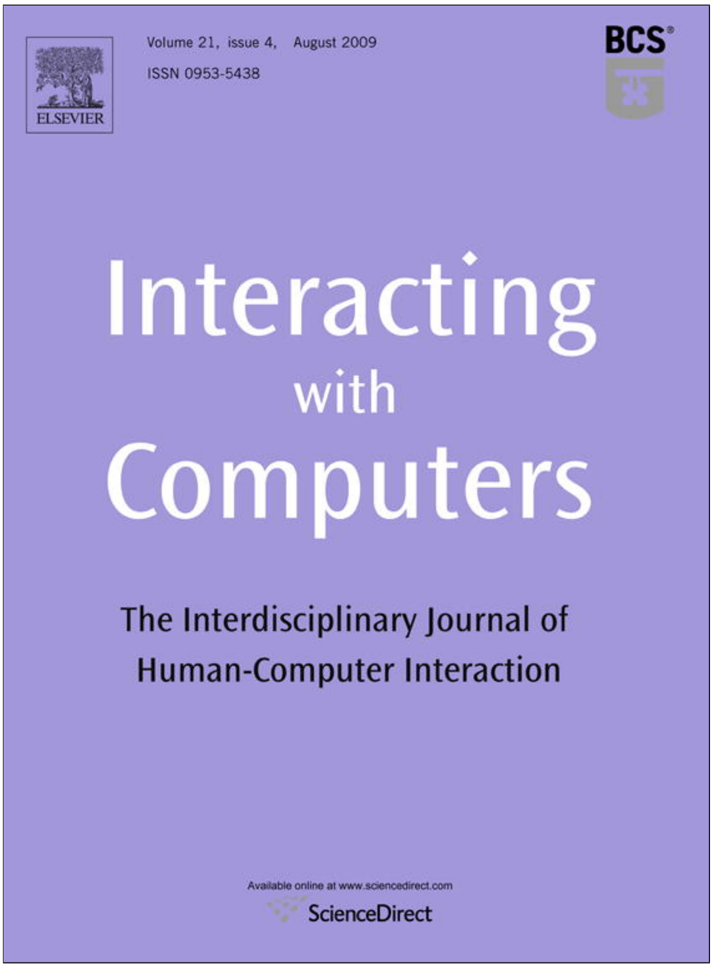
Volume 21, issue 4, August 2009
ISSN 0953-5438
ELSEVIER
BCS
Interacting with Computers
The Interdisciplinary Journal of Human-Computer Interaction
Available online at www.sciencedirect.com
ScienceDirect

Interacting with Computers 21 (2009) 288–303

Contents lists available at ScienceDirect

Interacting with Computers

journal homepage: www.elsevier.com/locate/intcom

UbiCicero: A location-aware, multi-device museum guide

Giuseppe Ghiani, Fabio Paternò *, Carmen Santoro, Lucio Davide Spano

ISTI-CNR, Via G. Moruzzi, 1 56124 Pisa, Italy

ARTICLE INFO

Article history:
Received 1 August 2008
Received in revised form 26 May 2009
Accepted 3 June 2009
Available online 8 June 2009

Keywords:
Mobile guides
Multi-device adaptation
Location-awareness
Interactive games in museums
User interface software and technology

ABSTRACT

In this paper, we propose UbiCicero, a multi-device, location-aware museum guide able to opportunistically exploit large screens when users are nearby. Various types of games are included in addition to the museum and artwork descriptions. The mobile guide is equipped with an RFID reader, which detects nearby tagged artworks. By taking into account context-dependent information, including the current user position and behaviour history, as well as the type of device available, more personalised and relevant information is provided to the user, enabling a richer overall experience. We also present example applications of this solution and then discuss the results of first empirical tests performed to evaluate the usefulness and usability of the enhanced multi-device guide.

1. Introduction

Recent technological advances, including increasing availability of various types of interactive devices, sensors and communication technology, enable novel interactive software environments to support users in different contexts with different objectives. In the edutainment area, museums are an interesting domain due to the large amount of digital information available and the increasingly technological resources adopted in such environments. Indeed the diffusion of the Web together with the possibility to search and visit a virtual museum already stimulated efforts in digitalising museum material in order to promote the museum collections by surpassing physical barriers and then offering to virtual users an experience as much similar as possible to that of the real visitors. The result is the massive amount of digital data now available in this field. Thus, museums are a particularly suitable context in which to experiment with new interaction techniques for guiding mobile users and improving their experience. However, at the same time, such a wealth of both information and devices might become a potential source of disorientation for users, if not adequately supported.

Traditionally, support for museum visits is limited to audio guides and interactive kiosks, which suffer from various drawbacks. It is important to exploit new technologies to identify new solutions also able to promote social interaction (Leikas et al., 2006) and enhance user experience. Many studies have underlined that the interaction with exhibitions, as well as communication and interaction among visitors are important for a successful learning environment (see for example, Hindmarsh et al., 2001; Leinhardt and Crowley, 2002). Indeed, museum settings generally imply a number of visitors at the same time, and this can be exploited to improve the user experience (regarding the learning and infotainment aspects) both at an individual and a cooperative level. If the mobile guides also provide some forms of games (both individual and collaborative) connected to the museum content physically visited, the visitors would be encouraged to challenge themselves and other visitors, and this could provide an interesting and amusing way to promote user interaction and learning in such settings. In this regard, the use of different types of devices within the museum (e.g., mobile devices and large screens) can be seen as a means to enrich user experience by enabling enhanced functionality, such as displaying the visitor's position and representing individual and cooperative games on the large screen while the user visits the museum using the mobile guide. In addition, intelligent user interfaces, by means of providing adaptivity and personalisation, are well suited for augmenting educational museum visits. Indeed, understanding and learning is facilitated when the systems use concepts familiar to the users (considering their interests and knowledge level). Thus, museum guides can automatically adapt the content presentation using user data stored in a user profile.

Here we present a new environment (UbiCicero), which allows for both individual and collaborative games together with location-aware support in multi-device environments. The goal of this study is to show the relationships between the use of mobile devices and stationary screens in the area of museum guides and provide indications of the kinds of information delivery that draws on their

* Corresponding author.
E-mail addresses: Giuseppe.Ghiani@isti.cnr.it (G. Ghiani), Fabio.Paterno@isti.cnr.it (F. Paternò), Carmen.Santoro@isti.cnr.it (C. Santoro), Lucio.Davide.Spano@isti.cnr.it (L.D. Spano).

0953-5438/$ - see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.intcom.2009.06.001

respective strengths, in order to support the user experience of museum visitors. In this context we also consider the provisions of location-based services in the museum.

In the paper, we first discuss related work, including our previous experiences with mobile guides. Then, we present our novel approach for mobile guides exploiting RFID technology to enable location-aware features, and games together with large shared screens for supporting social interaction and learning. We also discuss the associated adaptivity, which has been included to provide more personalised support. Then, we describe the underlying software architecture, and the various modules making up the museum guide. Lastly, we report on some first evaluations carried out on the solution proposed, and provide some concluding remarks and indications for future work.

2. Related work

In recent years, there has been growing interest in the development of mobile guides, thanks to the spread and increase in the performance of PDAs (Personal Digital Assistants) and mobile phones, the progress of wireless communications and improvements in localisation technologies. The main advantage of these technologies is the possibility, by following the context-aware computing paradigm, to provide users with context-dependent services, exploiting information such as the user's location, time, nearby people and devices, and current task.

One of the first examples of mobile guides with context-aware features was Cyberguide (Abowd et al., 1996), which supported navigation inside buildings or outdoors, visualizing a schematic map of the area, automatically updated according to the user's position, which was determined by means of infrared sensors (indoors) or GPS (outdoors). Research work has also investigated the possibility of dynamically generating and delivering personalised comments to the user. In this regard, HyperAudio (Petrelli et al., 1999) is an example of an adaptive mobile guide that generates audio comments for a palmtop taking into consideration the physical location of the users and the amount of time they spend in a certain location as an indication of user interest in certain artworks. More recently, the HyperAudio mobile guide has been the object of further consideration from some of its own authors (Petrelli and Not, 2005), in order to analyse whether the results that were gathered at that time are still valid in light of more recent experiences and findings. The GUIDE project (Cheverst et al., 2000) has been developed in order to provide city visitors with an intelligent, context-aware tourist guide. Information on user position is used to provide an automatic information delivery service: each point of interest is associated with a geographic area, each one supported by a Wireless LAN. When the user enters these areas the corresponding information is automatically provided as text descriptions, images and audio comments. The notion of context considered in GUIDE includes information related not only to users' location but also to their personal profile as well as external conditions (e.g., opening times of city attractions). Another example of a mobile guide that exploits various kinds of information to adapt its services is Hippie (Opperman et al., 1999), which combines the user's location with other information to provide additional details on the exhibits. It furnishes comments on the artworks in sight to the visitors, adapting the information to user's location, interests and knowledge, which has been derived from prior interaction. However, solutions based on automatic generation of comments on the closest artwork may sometimes be judged annoying. CRUMPET (Poslad et al., 2001) personalises its services, both on PDAs and mobile phones, not only according to the position of the users and their interests, but also according to previous interaction with the system. Lol@ (Pospischil et al., 2003) is a tourist guide for the city of Wien, using GPS as localisation technique and a GIS support for generating the maps. It adapts the information to the device, but not to the user's characteristics. VeGame (Bellotti et al., 2003) is another project that uses mobile technology to explore the city of Venice and learn about its history and architecture through games based on observation, reflection and action (e.g., video games). The system enables wireless communication but, due to limited bandwidth, communication between two peer PDAs is implemented for real-time exchanges as in video games. Visitors may play in teams and the only goal is to achieve the best score. In addition, each team can have multiple members but only one PDA. While such approaches to mobile guides have provided useful and important contributions, none of them has exploited the possibility of using large screens that can dynamically be available for the visitor.

The City project (Brown et al., 2003), part of the Equator project, was carried out in Glasgow at the Lighthouse Museum (http://www.thelighthouse.co.uk), which is dedicated to the work of the designer Mackintosh. The system exploits three kinds of technology: (i) for the real visit, the visitor uses a PDA equipped with headphones, a microphone and an ultrasonic location system; (ii) for the virtual reality visit, a visitor navigates in a 3D representation of the museum; and (iii) for the Web visit, a visitor navigates using a standard browser with Java applets. With this system, visitors are able to share their museum experience, and navigate jointly through mixed realities: the Web, the virtual and physical reality. Information is provided about each visitor location and orientation. In addition, users may communicate through audio channels. The authors have observed that voice interaction, location and orientation awareness, and mutual visibility are essential to the success of museum co-visiting between remote users. While the Equator City project provides for mixed visits combining the real and a virtual representation of the museum, through a Web site in a 3D manner, our work is instead aimed at supporting "physical" visitors moving about a real museum.

Furthermore, there is increasing interest in the use of RFID (RadioFrequency IDentification) technology in order to supplement physical artworks with associated digital information. In this regard, examples can be found in (Mäntyjärvi et al., 2006) and in (Bellotti et al., 2006). In (Mäntyjärvi et al., 2006) we proposed the scan and tilt interaction paradigm. Physical selection is obtained by scanning RFID tags associated with the artworks, and single handed tilt gestures are used to control and navigate the user interface and multimedia information. By pointing at the artwork of interest and controlling audio information with small single hand gestures, we avoided overloading of the visual channel, resulting in a less intrusive interaction technique. We performed a first empirical evaluation of this prototype. The test showed an overall good acceptance among users but, at the same time, highlighted some limitations. For example, the passive RFID tags used in this prototype forced the users to stand in very close proximity to the artworks, which is not realistic in museum environments and, even when possible, it would force the user to behave uncharacteristically (standing very close to artworks). One of the benefits of our new solution is overcoming such a limitation. As for (Bellotti et al., 2006), their mobile guide was specifically designed to support blind people while visiting an exhibition. Users were provided with multimedia descriptions about points of interest in an event-driven manner: when a user entered an area associated to an RFID tag the software proposed the corresponding description. However, some zones of the exhibition were not covered by any tag. Instead, in our guide there are no uncovered zones, because our application exploits the RFID technology both to provide the user with his/her position inside a section and to monitor user movements. Data on the path covered by the user during the visit are then compared with the artwork-related information explicitly accessed through the PDA interface. Such strategy of user modelling has an adaptation purpose that will be explained in later sections. Another example of a museum application exploiting the use of RFID technology to improve the social interaction between users is the History Hunt project (Fraser et al., 2003), which has been implemented in the Nottingham castle museum. In this project, groups of visitors (generally families), on arrival, collect a set of clues written on pieces of paper electronically tagged with RFID tags that lead them in search of a specific historical figure. Visitors can interact with some display installations around the museum in order to receive further information about the clue described in the RFID tag. Collaboration occurs only when visitors are close to the interactive displays. Instead, our approach is to support collaboration throughout the whole visit (see Section 3 for more details). We found developing games interesting in order to enhance interactivity in museums and also to offer visitors the opportunity of improving their learning experience by challenging themselves and/or competing with others. The idea of games has been extensively explored mainly through more conventional media such as paper-based treasure hunts. However, little research has been reported on the role of digital adaptive games inside galleries, even if similar research in other contexts has been carried out (Bell et al., 2006).

The development of a museum guide prototype tested on the Canadian Museum of Nature in Ottawa is presented in (Hatala and Wakkary, 2005). The guide, called Ec(h)o, is an augmented audio reality system with location-aware (RFID and optical position tracking) capabilities. The authors found that, exploiting ontologies and monitoring user interaction, an accurate dynamic user model and recommender system can be built in order to augment museum visits. (Sparacino, 2002) describes a wearable museum guide that gathers visitor's interests from the peculiarities of the physical path covered (e.g., length of stops). A progressive refinement of the user model allows the delivery of personalised multimedia descriptions to enrich the museum visit. Differently from such works, in our approach we take into consideration the multi-device opportunities that museum environments nowadays are able to offer to visitors (e.g., large screens, given their low cost), not only for improving the individual visitor experience, but also as a means for stimulating socialisation and interaction among users. More recently, (Stock et al., 2007) proposed an interesting multi-device support to improve museum visitors experience in the PEACH (Personal Experience with Active Cultural Heritage) project. Both mobile and stationary devices are exploited and a virtual presentation agent transfers from the handheld device to the *Virtual Windows* (large displays) when such devices are in sight, and assists the visitor. Adaptivity is achieved by considering the elements that seemed of major interest to the visitor and tailoring multimedia presentations. One of the claimed principles is the unobtrusiveness of the visit support: the system should never intrude between the exhibit and the visitors, letting them focus their attention on real things. In our work, we provide novel solutions to exploit interactivity with mobile and stationary devices, for example by providing the possibility of dynamically migrating part of the game interfaces (including social games) from one device to another. The next sections will provide a detailed discussion of such features.

3. The mobile guide

Our interactive environment for museum visitors has been applied to the guides of two museums: The Marble Museum (Carrara, Italy) and the Museum of Natural History (Calci, Italy). We started with a previously existing application for mobile devices: Cicero (Ciavarella and Paternò, 2004), which was a digital museum guide developed for a PDA platform and freely available to the visitors of the Marble Museum located in Carrara. Such experience allowed us to understand specific requirements of mobile museum guide users. First, users are not willing to spend too much time learning how to use it. In addition, users rarely visit the same museum more than once,. Thus, the mobile guide should be intuitive and easy to use, even for novice/occasional users. Lastly, the mobile application should support/integrate the users' experience in the real museum. Therefore, it should not be too intrusive or hinder their appreciation of the real artworks, and at the same time it can represent a useful learning/edutainment opportunity for visitors. In order to fulfil such requirements, a number of features have been included in the mobile guide. We will describe them in the following sub-sections. More in detail, in order to offer an intuitive support to users visiting a museum, the mobile guide provides them with a rich variety of multimedia information (graphical, video, audio, etc.) regarding the available artworks. Most information is provided mainly vocally in order to allow visitors to freely look around. However, the visual interface can also show related videos, maps at different levels (museum, sections, rooms), and specific pieces of information.

In addition, the guide. is able to support other location-aware services such as showing the path to a specific artwork from the current location.

The location-aware features of the mobile guide will be discussed in Section 3.1. Moreover, in order to provide information that is considered interesting for the user, personalisation/adaptation capabilities have been included (they will be described in Section 3.2). In addition, in order to enable more user involvement and interaction, different types of games are available as a form of learning activity, both as single-user and multiple user games (see Section 3.3). Lastly, the entire Section 4 is dedicated to the description of the features for exploiting the multi-device capabilities that museums can offer.

3.1. Location-awareness

Location-awareness is supported in the mobile guide, as a means for improving its flexibility and efficacy. Indeed, by providing users with information and services that depend on the user's context (in this case: the location), it is possible to offer more appropriate support, since it is possible to infer what the most relevant information is and deliver it. In this case, the location-aware mobile guide allows for navigating both a physical space and a related digital information space at the same time, thereby offering a more complete experience to the user. To this end, we have added some location-awareness in the mobile guide by means of long range active RFID tags for automatically detecting nearby artworks. Active RFID technology allows a less intrusive localisation approach since the user does not need to be in very close proximity to an artwork to detect it, which is a useful feature in real museum settings. Fig. 1 shows the museum guide implemented on the PDA device equipped with the Compact Flash RFID reader, which is a small add-on that can be plugged into a standard PDA interface.

In our system museum artworks are fitted with physical tags, and each tag has a unique ID number. However, a single tag may be associated with more than one nearby artwork, when they are very close. This is due to the difficulty of distinguishing two or more tags that are very close to each other. Indeed, if two tags were placed in a very small area, the reader would detect both of them with the same RSSI (Received Signal Strength Indication). The use of RFID technology for localisation and the problems related to tag density are also discussed in (Bellotti et al., 2006). When the user enters a new room, a short report summarizing the visit of the previous section, (e.g., visited artworks) is generated. Then, the guide activates the museum map highlighting the new room and the corresponding vocal comments. When the user is close

Fig. 1. The museum guide in a mobile device with RFID reader.

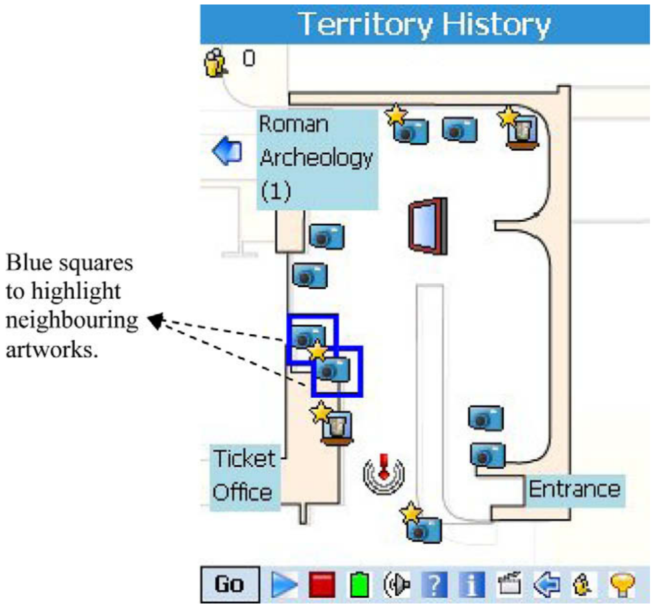

Fig. 2. The artworks closest to the user are highlighted by blue squares. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

to an artwork the guide vocally asks whether the user wants additional information; given an affirmative response, this is provided vocally with further related graphical content (such as videos, when available).

Continuous monitoring of the tags' signals allows the guide to calculate the artworks closest to the user. When a new tag is detected, i.e., a new area is being visited, an audio clip is played to alert the user, and a vocal message indicating the number of nearby artworks is generated. Depending on user preferences, the guide may ask confirmation before describing any artwork or even it can describe them automatically. Fig. 2 shows how the artworks detected via their tags are highlighted in the map through a blue frame around the corresponding icon.

Fig. 3 (left part and right part) shows how the guide dynamically updates map information as the user moves about. The three artworks displayed in a different colour on the left screenshot in

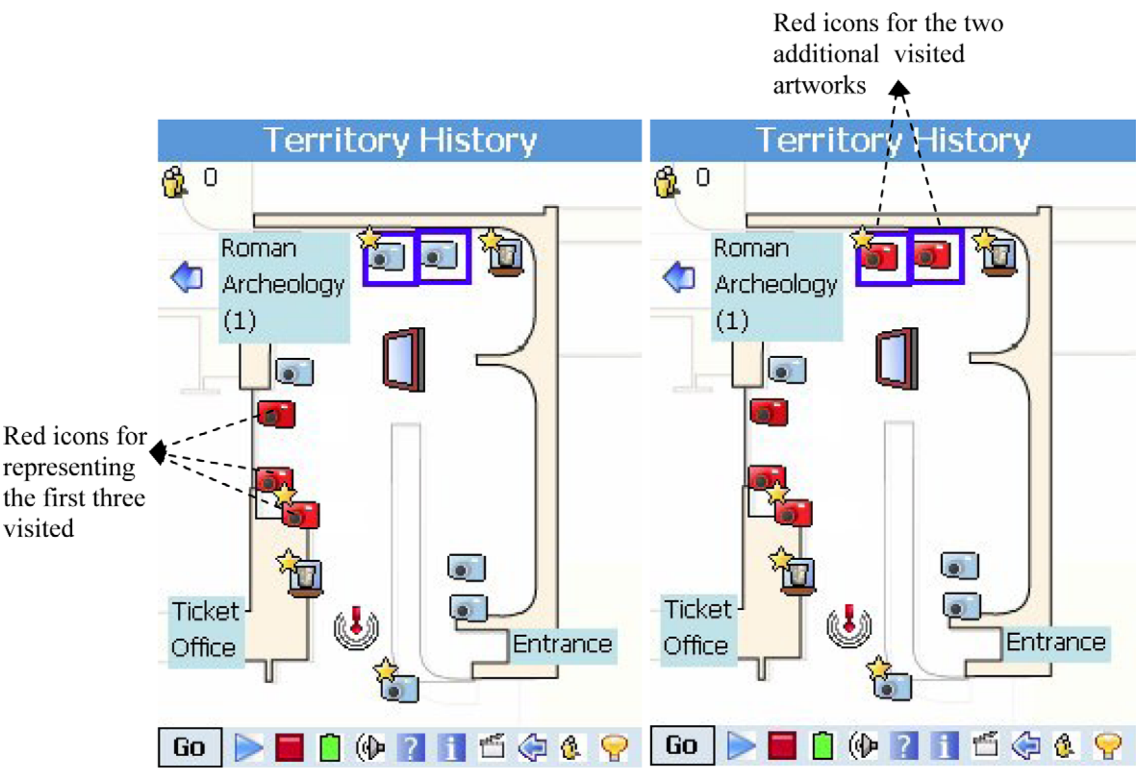

Fig. 3. Dynamic update of location-aware map information.

Fig. 3 represent those already visited. Then (right part of the figure), the user accesses the descriptions of two different artworks and the colour of the associated icons changes to signal the fact that the related artworks have been viewed.

3.2. Adaptive support

The mobile guide has a number of adaptive features, which enable a degree of personalisation in the user interface. In order to provide such features, the adaptive support is based on a *user model* which is designed to contain information on the user (interests, preferences, etc.). The user model is dynamically updated by considering recorded information of the events occurring during the visit (the user is tracked) and it also infers information regarding user's level of interest towards the various museum items in order to provide personalised information. The application logs any interesting event, such as section and area changes (in terms of nearby artworks) and how much time the user has lingered within such areas. It is then possible to know the path covered by the visitor through the museum. The time spent by the user listening to an artwork description is also logged. Artwork and room descriptions are vocally reproduced by means of a TTS (Text To Speech) engine. The Loquendo Embedded TTS (http://www.loquendo.com) has been used as the vocal engine. The main benefit of a dynamically generated audio description is the limited memory required: the audio files do not need to be physically stored on the PDA but they are dynamically created by the TTS module. Additional (specific) artwork information can then be easily managed because the texts do not have to be expressly read by a human to create the audio files but are dynamically generated by the TTS module. This is particularly significant when the descriptions have to be available in different languages.

The analysis of the collected data allows the guide to estimate in real time the user's currently favourite type of artwork and assess his/her preferences. Logged information is used to keep the user model updated. At each section change, the user model calculates which artwork in the new section would best match the user preferences. The suggested artwork is then highlighted by a light-bulb icon next to the corresponding artwork (see Fig. 4). Details on how user's interest estimation is presented on the large screen device are provided in Section 4.

The application uses several sources of information to estimate user preferences. The data analysed to infer such preferences includes the physical and the virtual visit the user performed. Indeed, since one of the requirements of the mobile guide was to be non intrusive, we decided to derive data for the user model by only observing the visitor's behaviour.

To this end, each museum item accessed in the mobile guide is interpreted as an indicator of interest/preference of the user towards the characteristics of that item. In particular, author, chronology, material and artwork category (e.g., sculpture, painting, photo/picture) are considered in the computation of the preferences. Also, the time spent on listening to the description of an artwork is interpreted as an implicit sign of interest and therefore affects the specific weight of the values of the various attributes associated with a particular artwork in the user model. Moreover, the playing of any game also influences the user model: whenever a user selects a game, this is taken as an indication of interest, which is further augmented if the user correctly solves it. In addition, if the user also physically lingers in front of the artwork (and consequently the application has logged that the user has entered the corresponding artwork's RFID tag area) then a higher level of interest is recorded for this user with respect to that kind of artwork.

Local suggestions are also available on request: when the light-bulb button is selected in the toolbar, the artworks in the current area that have not yet been visited are considered as possible next candidates and a light-bulb icon besides an artwork indicates the most interesting artwork among the nearest ones.

The user model also exploits the possibility to generate audio descriptions on the fly. The insertion or deletion of some parts of the automatically read text is based on the estimated user preference, tailoring the length according to the calculated interest level.

The *UserModel* module consists of two main parts, one maintains the data and the other one infers user preferences and updates the model. These parts are:

1. *Directory component*, storing information about the user, which can be of three types:

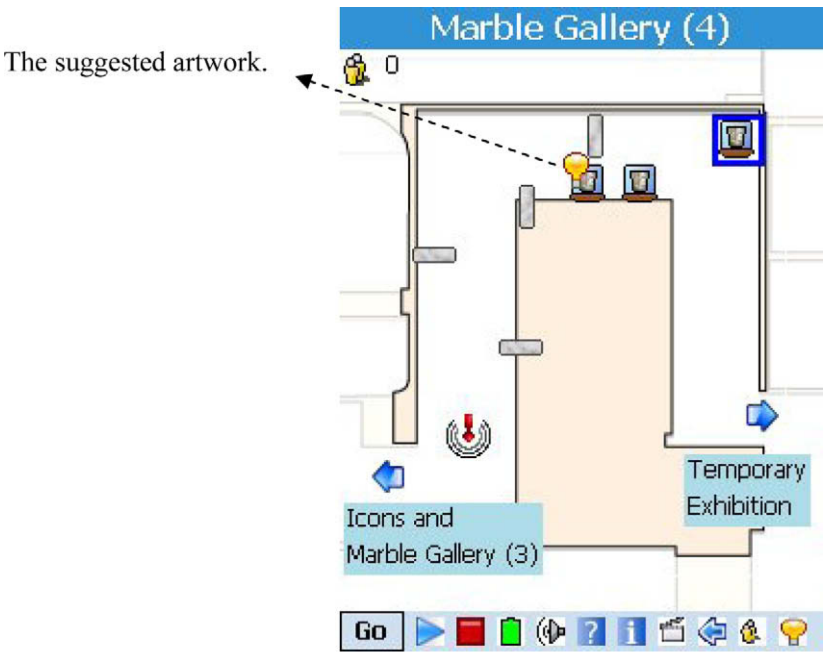

Fig. 4. Highlighting the suggested artwork by a bulb icon.

   a. *User data*: related to user preferences and knowledge;
   b. *Usage data*: related to the interaction between user and system, and implicit indications of interest, such as time spent nearby the artwork;
   c. *Environment*: aspects external to the user (device, collaboration through games, etc.).
2. *User modelling component*, reads logged data and updates the *Directory component*. It also infers user preferences and knowledge by exploiting three sub-components:
   a. *User learning component* updates *user data* by analysing the direct interaction with the system;
   b. *Mentor learning component* exploits previous visitor models and predefined stereotypes to estimate the possible level of interest for a topic, even if the guide has no direct data on the user (heuristic reuse);
   c. *Domain inference* is a representation of the museum domain, for establishing correlations between topics and artworks.

The updating of the user model begins with the communication of an interaction by the user interface to the logger, which then notifies the appropriate event. Each user model component that is subscribed to this event infers data with its own algorithm. Then, when the presentation layer reads the status of the user model it can adapt the user interface accordingly.

Visitors generally go to museums in order to acquire knowledge (e.g., about museum collections) and gain new interests. In addition, while they visit the museum, they should become more familiar with the collections and then progressively tend to focus their attention on specific types of items. This is true especially when the exhibition is very large. For this reason, our user model also adopts a recentness-weighted approach attributing higher weights to the ratings of the most recently accessed artworks. Such ratings will be used by the guide to provide the user with suggestions regarding the next artwork(s) to visit (in the user interface, such suggestions will be rendered through light-bulb icons).

Whenever the user accesses a certain artwork, the algorithm has to update the vector that contains the current ratings. The calculation of the rating of the current artwork is performed by taking into account the ratings of the previously accessed artworks that are *correlated* with the current one. In case of the first access, the rating is calculated by considering static data provided by the above mentioned *User modelling component* module (by using the *Mentor learning component*, which analyses predefined stereotypes to estimate level of interests for a certain topic/artwork). Then, initially, each artwork in the museum is assigned a default rating value, which is statically calculated; afterwards, the values are dynamically calculated.

3.3. Games

Games are meant to be a tool for improving the visitors' educational experience. They are related to the real visit, and are essentially aimed at stimulating more active participation, by enabling the visitors to challenge themselves and others about the information they received, also to check what they have actually retained. It is worth pointing out that the latter point is not trivial, because especially in large museums the information conveyed can be quite overwhelming. Thus, games are a way to check the knowledge they should have actually gained during the visit and a useful learning aid.

In addition, individual games are a way to support learning at users' own pace (a game can be repeated multiple times if users are unable to successfully solve it at first, a more interested user can play many more games if s/he is interested in deepening/broadening her/his knowledge of the museum collection, ...). Moreover, collaborative games are meant to stimulate talking and discussion about the subject of the game, and, in this regard, the use of a large screen is meant to facilitate discussion among a large group of people.

We introduced some preliminary ideas regarding the use of games in mobile guides in previous work, though we did not use any localisation (Dini et al., 2007) or the multi-user support (Ghiani et al., 2008). In such systems, the inclusion of games within

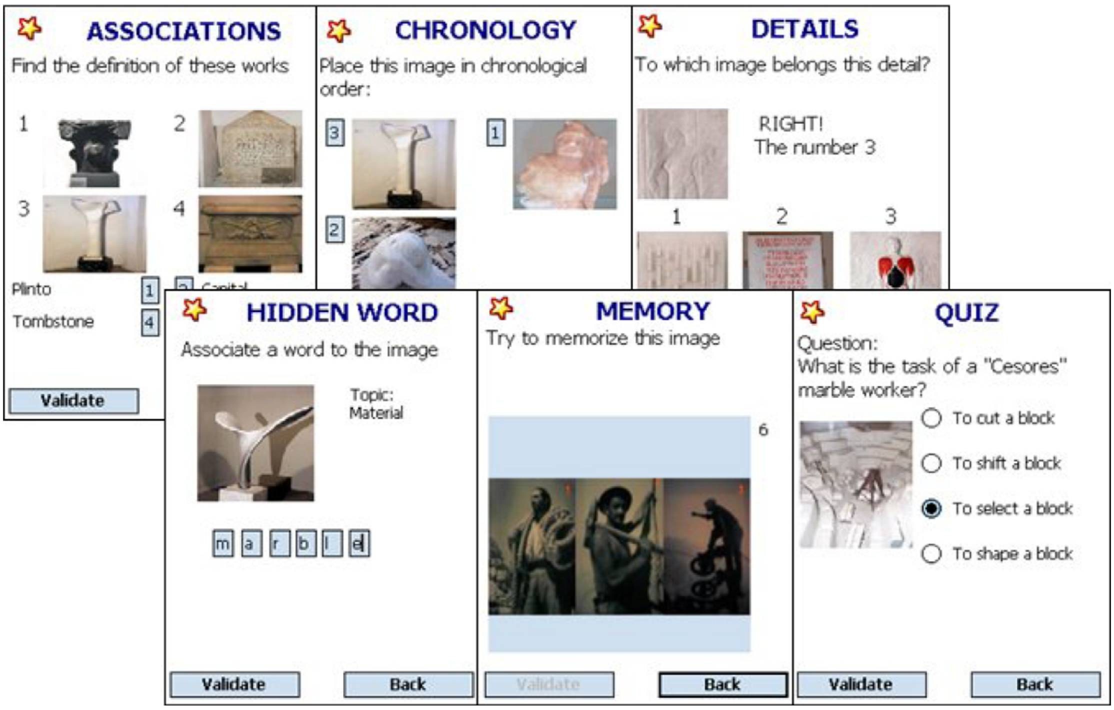

Fig. 5. The six types of individual games.

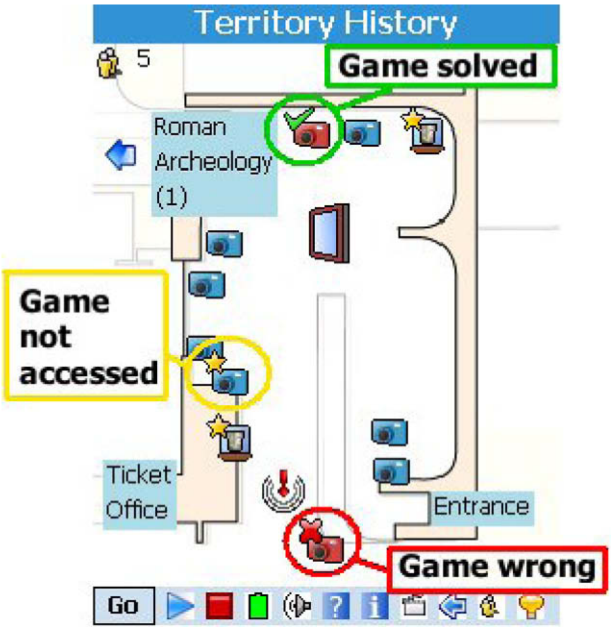

Fig. 6. Representations of the game state.

the mobile guide was appreciated as a good way to test what was learned during the visit and also to make the museum visit more interactive.

Here we present a comprehensive solution, tested into two museums, with a wider set of games and an associated game editor. Such games are integrated with the user model able to support some adaptive features. Playing is carried out by interacting through radio buttons or text boxes and just requires simple clicks or the insertion of some text. Then, interaction is quick, games may be solved in a few seconds and thus they are suitable to visitors of a wide age range. Even children can play the games included in the system to enhance their experience and improve learning.

3.3.1. Individual games

In order to increase the learning experience, we defined six types of individual games related to the museum content (see Fig. 5):

- The *associations* game requires linking images to words (e.g., the image of the artwork to its author or material).
- In the *details* game an enlargement of a small detail of an image is shown. The player has to guess which of the artwork images the detail belongs to.
- The *chronology* game requires the user to order the images of the artworks shown chronologically according to the creation date.
- In the *hidden word* game, the user has to guess a word associated with a particular attribute of an artwork: the number of characters composing the word is shown as help.
- In the *memory* game, the user has to look at an image (which disappears after a while) and s/he has to answer to a question associated with the previously shown image.
- The *quiz* is a multiple-choice question with a single correct answer.

As can be seen in Fig. 6, the artworks which have an associated game show an additional star-shaped icon (beside the icon representing the artwork), through which the related game can be accessed. If the game is correctly solved, the icon turns into a green check sign, otherwise it becomes a red cross.

3.3.2. Cooperative games

Social interaction is acknowledged as a relevant tool for improving learning (see for example, Hindmarsh et al., 2001; Leinhardt and Crowley, 2002; Grinter et al., 2002; Cosley et al., 2008). Therefore, in order to stimulate cooperation among visitors while heightening social interaction, besides individual games, we have also included group games (in which users are organised in teams). One example is the shared enigma, which we introduced in (Ciavarella and Paternò, 2004). It is composed of a series of questions on a topic associated with an image hidden by a jigsaw puzzle. Each player must solve an individual game to reveal one piece in the puzzle, which is shared by the players of the same team, thus helping the team to answer the question associated with the puzzle. As soon as an individual game is solved, the group earns seven points. The goal of the game is to stimulate interaction and cooperation among players in order to find the answer to the shared enigma. Indeed, when a member of the group answers the corresponding question, such a question is no longer available to the other players in the same team. This stimulates interaction among visitors so that they can first discuss the answer and agree on it. However, this mechanism might favour groups with a larger number of members since they are likely to earn points more easily: for this reason, the maximum number of people in a team was fixed at five. The PDA interface of the shared enigma has two parts (see Fig. 7): the first one shows the current players' scores and the hidden puzzle image, the second one shows the questions (with the possible answers).

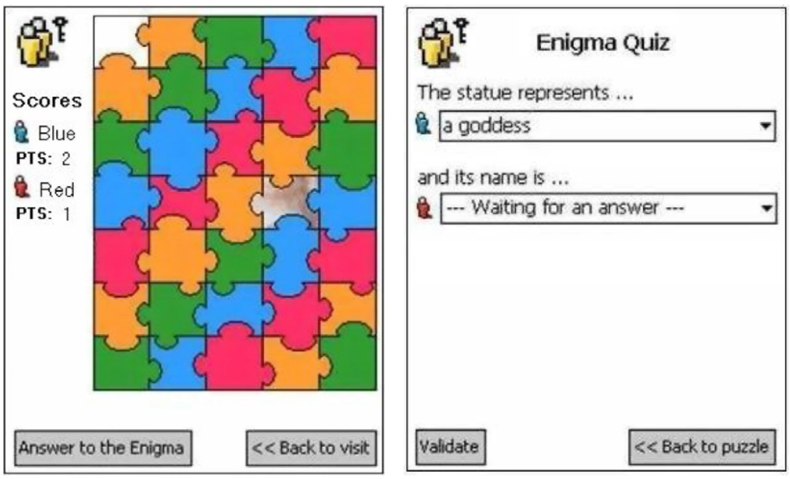

Fig. 7. The shared enigma on the PDA.

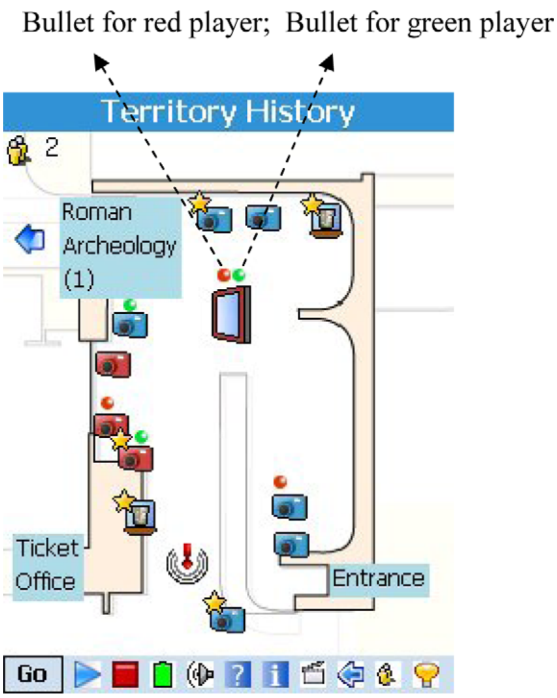

Fig. 8. PDA interface for the cooperative version.

Fig. 8 shows how the various players in the same team can have mutual awareness of each other. Indeed, coloured bullets are placed beside the presentation of each visited artwork: each coloured bullet is associated to a different player in the same team, so that a player can see which artworks have already been accessed by the other players. Differently from Fig. 8, in Fig. 7 it can be seen how many points each player (identified by a colour) has gained up to the current moment.

As we will better describe in Section 7, the current games - though not particularly complex- received, on average, positive feedback by the users. We plan to add further games, especially social ones in order to further improve the educational experience.

4. Exploiting multi-device environments

The main feature of our solution is to support visit and game applications exploiting both mobile and stationary devices equipped with large screens. The typical scenario is users freely moving and interacting through the mobile device, who can also exploit a larger, shared screen of a stationary device (which can be considered a situated display) when it is nearby. (Brignull and Rogers, 2003) reported on the "honey pot effect" arising from the interaction with public situated displays: shared screens connected to stationary systems promote social interaction and improve user experience, otherwise limited to individual interaction with a mobile device. They also stimulate social interaction and communication with other visitors, though they may not know each other. Regarding our domain, a larger shared screen extends the functionality of the mobile application enabling the possibility to present individual games differently, to share social game representations, to show the positions of the other players in the group and also to perform a virtual pre-visit of the entire museum.

Each shared display can be in two basic states:

- STANDALONE: the screen has its own input devices (keyboard and mouse) and it may be used for a virtual visit of the museum. It can be exploited by visitors who do not have the PDA guide.
- SPLIT: indicates that one visitor has taken control of the display (through the PDA), which shows the name and group of the current controller.

The communication between mobile device and large screen is described in Section 5.

Since a shared display can support various types of accesses, the structure of both its layout and some parts of the interface remain unchanged and are similar to the PDA interface, in order to avoid disorienting users. This permanent part of the user interface provides information such as the current section map and its location within the museum, an explanation of the icons used to represent the artworks and the state of the shared enigma. In standalone mode users can select from three kinds of views of the section map, using the toolbar on its top-left corner. These views are:

1. ICONS: the artworks are represented by an icon indicating the type.
2. THUMBNAILS: the icons are replaced by a small photo of the artwork.
3. THUMBNAILS AND ICONS: small artwork photos are accompanied by icons on their bottom-left corner (see Fig. 9).

By exploiting the user model data available from the connected PDA, the large screen application generates an "interest" rating for each artwork in the selected room. The user may look up the ratings, which are expressed by "LED light bars" on a scale of 0–5 (see Fig. 10).

When the shared screen is in "standalone" or "split" mode and a user selects an artwork or a game, the screen interface changes its layout by adapting the focus: it magnifies the correspondent panel and shows the artwork details (see Fig. 11) or the game interface. The representations provided take into account the screen size. Thus, when information on an artwork is provided on the large screen the user interface shows the following information: how to reach it by highlighting the shortest path in the museum to get to it; the map of the associated room highlighting the artwork location in it; the state of the last game played; a large resolution image annotated with pieces of information that cannot fit in a single presentation of the PDA user interface (which shows them in different cards).

The screen changes its state to "split" when a player selects the connection through the PDA interface. In this case the large screen is used both to show additional information and also to focus the attention of multiple users on a given game exploiting the screen size. When a player is connected to the large screen, its section map view is automatically changed to thumbnails, while the artworks types are shown on the PDA screen (see Fig. 12). The artwork presentation uses a higher resolution image on the large display, adding more description information.

The game representation on the large screen is used to share and discuss it among multiple users. Differently from the mobile representation (see Fig. 13 A) , in the "split" mode the interactive controls are available only on the PDA interface of the current controller (see Fig. 13 B1), while the question, the higher resolution images, and the game results are shown on the larger screen (see Fig. 13 B2).

Even the visualisation of the shared enigma differs depending on the availability of the large screen. If just the PDA is used, it is composed of two presentations visualised sequentially on the PDA: the hidden image and the associated questions. If the larger shared screen is available, then the hidden image is shown on the large display, while the questions and the possible answers are presented on the PDA user interface.

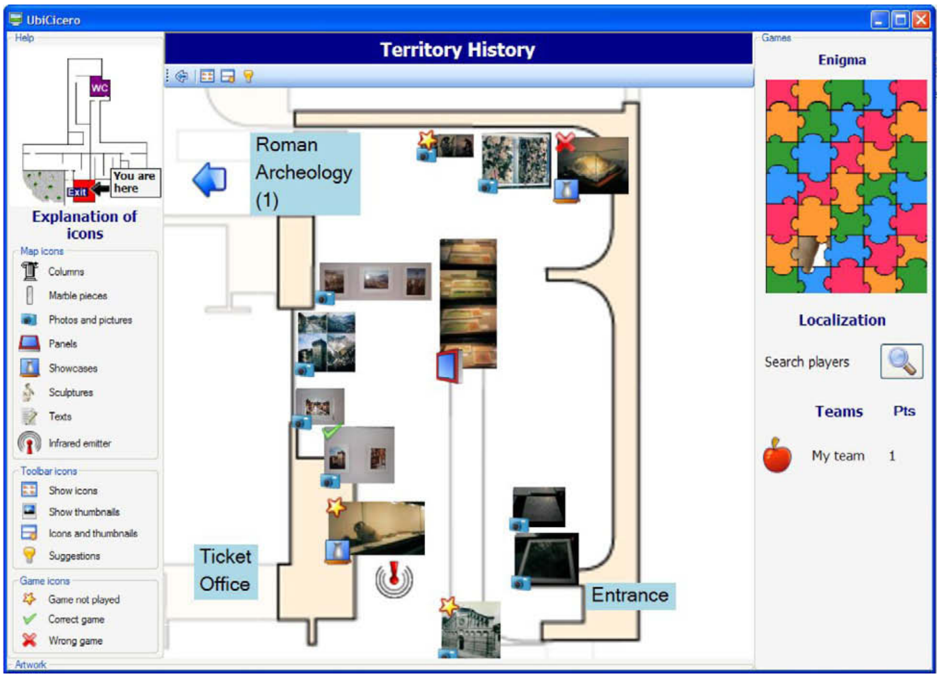

Fig. 9. A museum room displayed on the large screen (thumbnails plus icons).

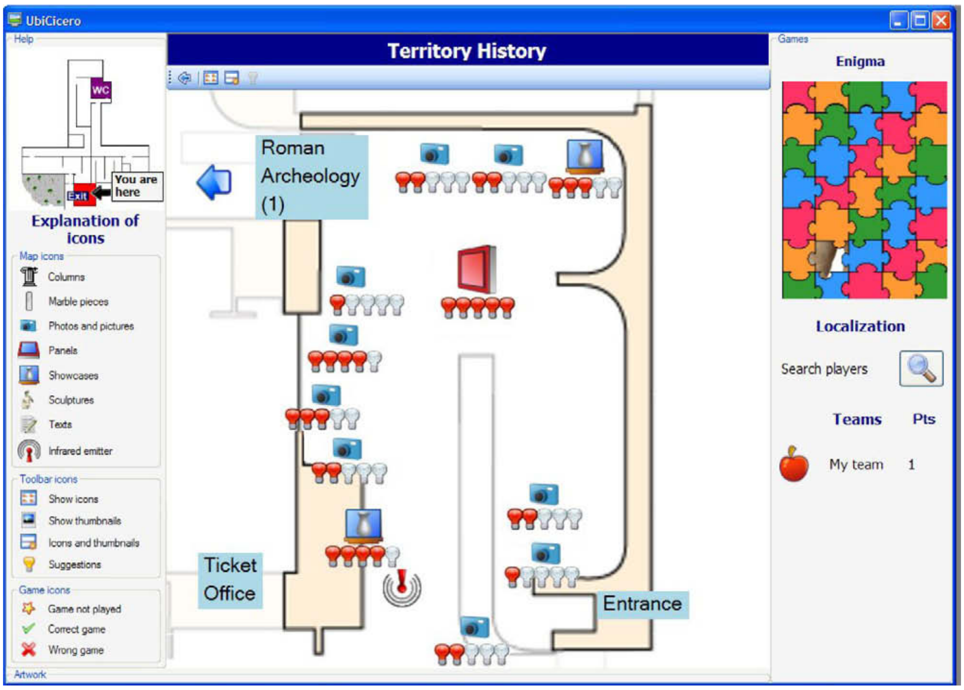

Fig. 10. Presentation of different levels of user interest for various artworks.

Providing an effective representation of players' position on the PDA is very difficult because of the small screen area available on that device, especially when the players are in different rooms. Thus, the large shared screen can be better used for this purpose: it is divided into sections, one for each player. Each section shows the name and the room where the player is located and the artworks close to the player are highlighted by rectangles (see Fig. 14). The location of the player is automatically identified through the detection of the RFID tags near each visitor, this information is passed wirelessly by the PDAs to the computer controlling the large screen.

5. Software architecture

In this section we provide a description of the software architecture of our solution, which is composed of different modules allocated on several devices, and detail some of the techniques used. The main elements of the software architecture are the software components in the PDA, in the stationary device and the communication protocol of the environment. The PDA module is composed of five layers, each of therm provides the others with services (see Fig. 15). Going up from the bottom they are:

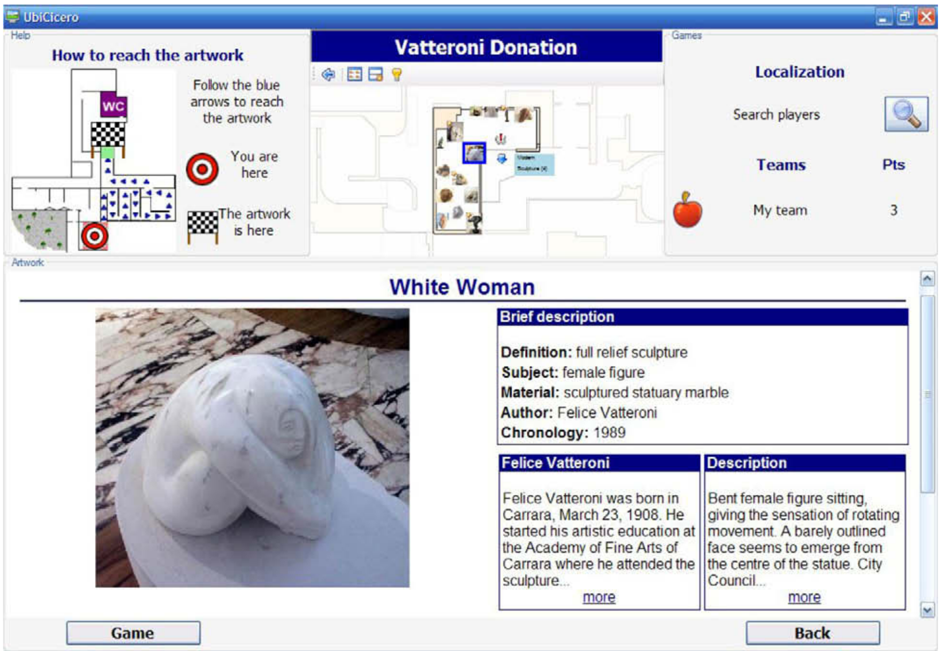

Fig. 11. Artwork presentation on large screen.

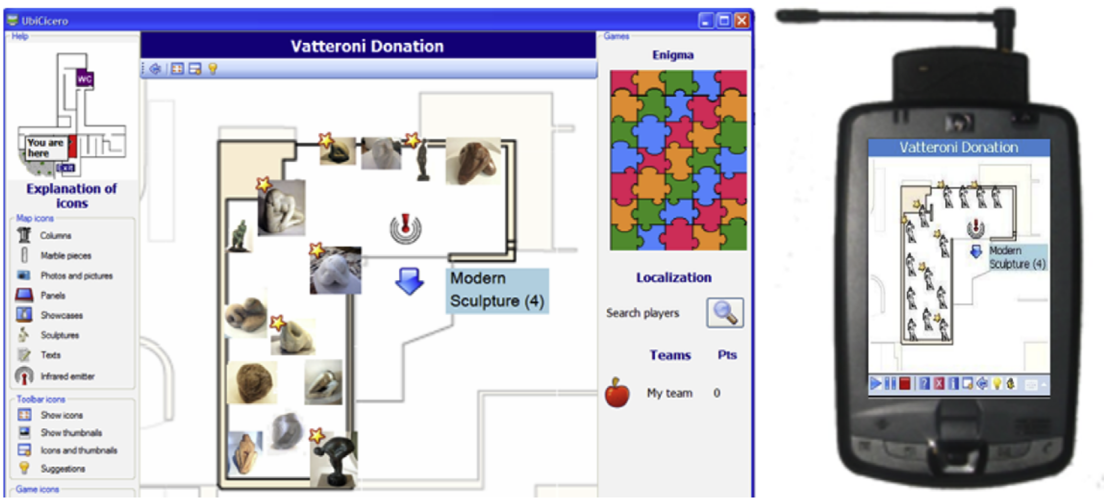

Fig. 12. Different presentations of the same section map on the large and small screens.

- *Museum definition*, which provides the specifications of the whole museum:
  - XML database.
    - Sections' layout (room sizes and shapes, and item positions).
    - Description of artworks and authors.
  - Multimedia resources such as photos and videos concerning museum items.
- *Localisation*, enabled by the RFID detection function.
- *Core*, defining routines for accessing the database from the upper layers, for communicating with the other devices and for logging relevant events.
- *User model*, which contains the level of preferences regarding various information attributes.
- *Visit*, supporting interactive access to museum info.
- *Games*, providing educational games.

Fig. 15 shows the architectural blocks of the configuration for our guide application. It is worth pointing out that not all the functionalities must necessarily be present at run time, since several modules are "detachable" by means of conditional compilation. Therefore for example, if a certain museum is not going to adopt a localisation infrastructure (i.e., RFID tags), the application version for that particular environment would be re-compiled omitting the RFID detector, and this will be done without the need of restructuring the whole source code. In the following sub-sections we provide further details regarding such modules.

5.1. Museum definition

The museum environment is defined through an XML database and a set of multimedia resources. The database specifies the shape, size and layout for each section of the exhibition. The

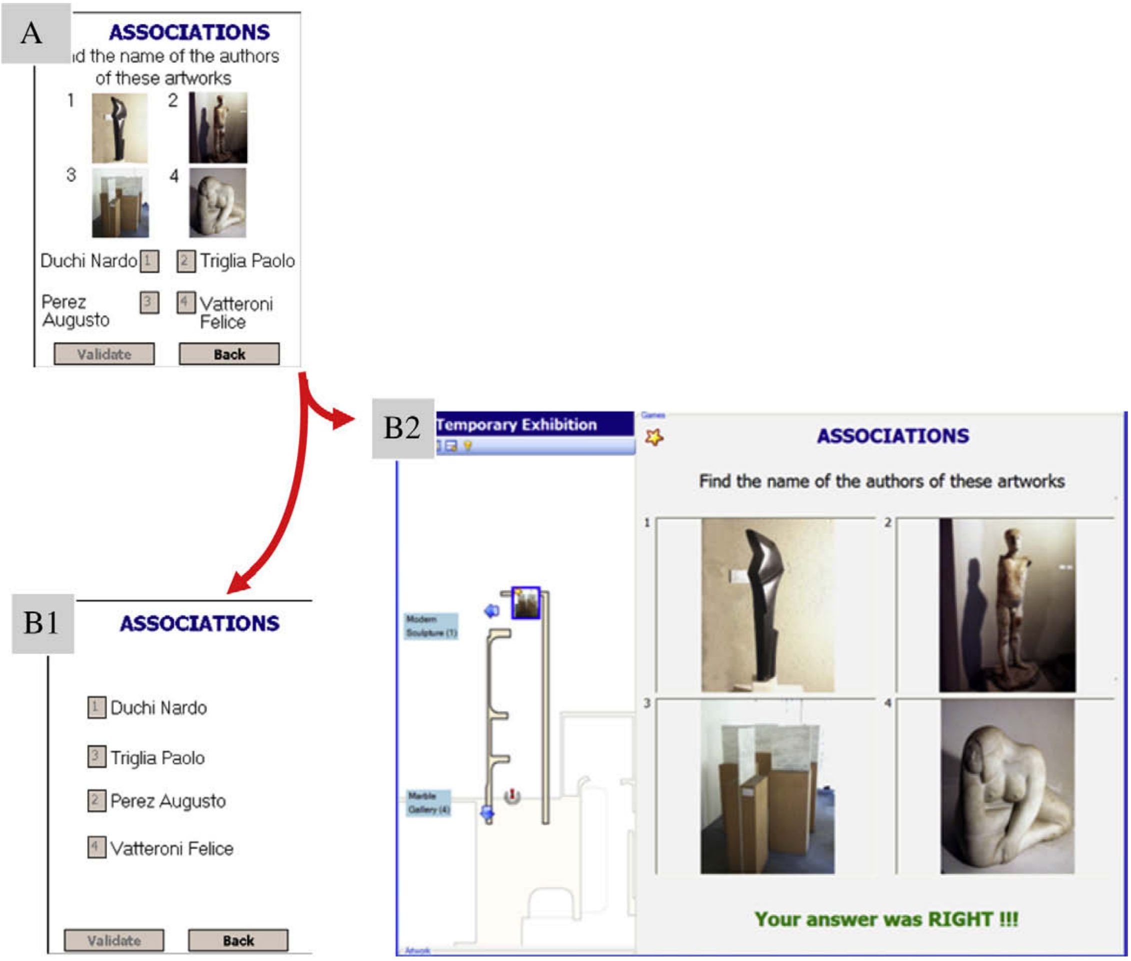

Fig. 13. Example of game in mobile mode (A) and in distributed mode (B: large screen (upper part); PDA (lower part)).

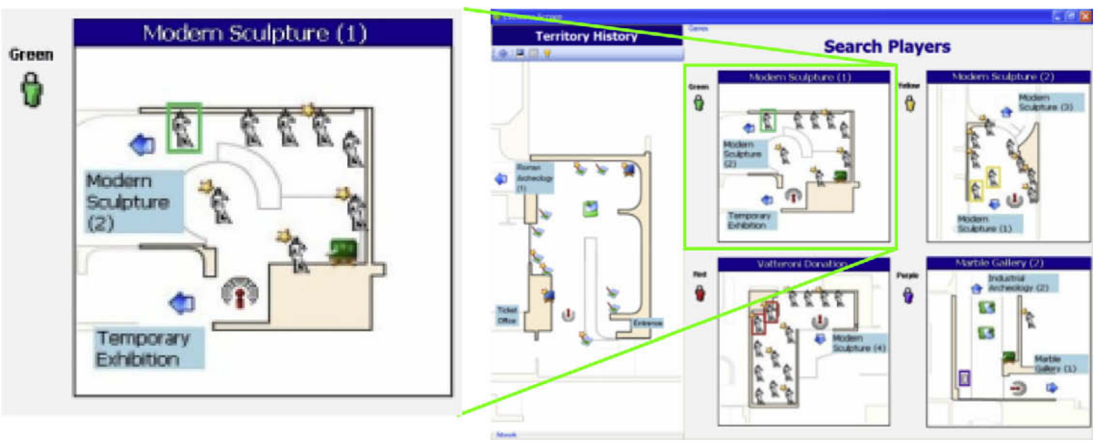

Fig. 14. Showing players' positions with respect to the visited artworks on large screen.

artwork-RFID tag correspondence is also stored in the database. Such an association is fundamental to enable highlighting the currently visited item(s) at run time (see Fig. 2). The stored multimedia files represent additional information about artworks/sections/authors that the user may want to view. Pictures or videos showing the creation of a contemporary sculpture are an example of such resources.

5.2. Localisation

User's position is showed on the PDA by highlighting the icons of the artwork(s) nearby. For instance, Fig. 2 shows the artworks closest to the user highlighted by means of blue frames around the corresponding icons. This position is identified by detecting the corresponding RFID tags.

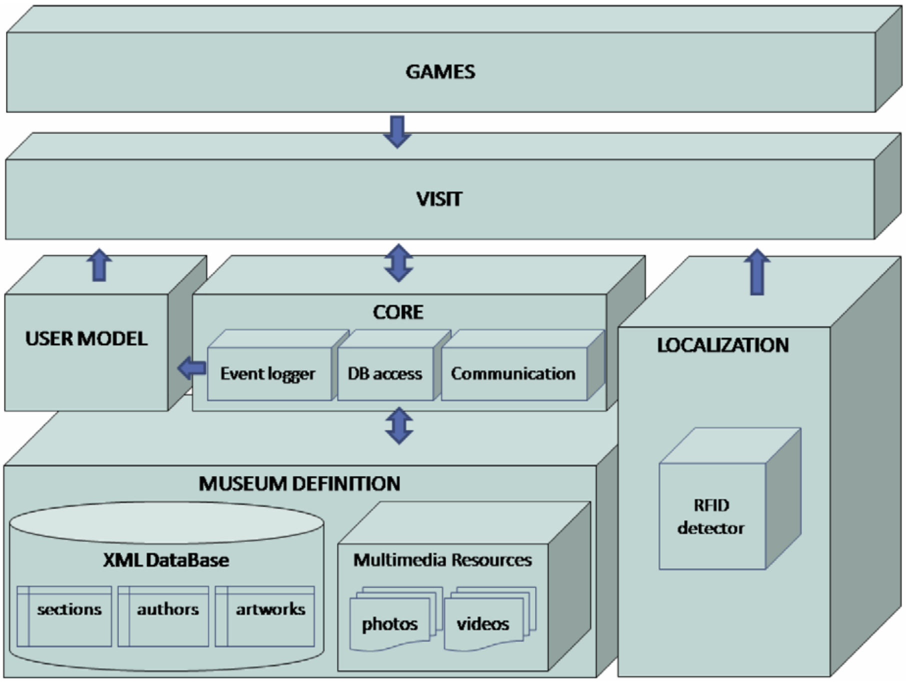

Fig. 15. The software architecture of the mobile guide.

5.2.1. RFID technologies

RFID-based solutions are composed of two main parts: the set of tags, or radio transponders, and the tag reader, or transceiver. Tags basically have a static identification number (ID), but may also store different types of information such as sensed data (e.g., environmental temperature). The reader scans for available tags, and, depending on their features, may query them for additional information stored in their embedded memory.

RFID technology can be applied by using *passive* or *active* tags.

- *Passive* RFID tags do not have any internal battery and exploit the energy electromagnetically inducted by a neighbouring antenna. Thus, passive tags can respond with their IDs when and only when a reader within a few centimetres' range interrogates them. Thanks to their simple internal structure, such devices can be embedded in a credit-card form factor and are extremely inexpensive. As already mentioned, passive tags on the 13.56 MHz frequency were used in a first version of our museum guide application.
- *Active* RFID tags are equipped with an internal power source and are able to transmit autonomously RF signals at any time. Thanks to the internal power source, active tags offer more flexible solutions also from a technological point of view. There are two different types of active tags, depending on the way they work: beacon/broadcast and respond.
    - *Beacon/broadcast* tags have a radio transmitter and they continuously send their data at certain intervals of time. The reader tunes into the proper radio frequency and listens to the tag information. Also, beacon technology uses read-only tags;
    - *Respond* tags (which, in contrast to the previous ones, are read/write tags) wait for a reader request before responding. Due to the need for both transmitter and receiver modules, respond tags involve more complex architectures, larger dimensions and higher costs than beacon tags. However, while the battery life of beacon tags (2–4 years) depends on how frequently they transmit, the battery life of respond tags depends on how often they are interrogated. Generally, the typical lifetime of a respond tag battery is more than 6 years.

5.2.2. Our RFID configuration

In our case we have used active beacon/broadcast tags with 868 MHz European frequency because they can cover wider areas than the passive ones. Our reader triggers tag detection at about 5 m distance. We preferred the beacon/broadcast solution, which is more scalable than the respond solution, which experiences performance decay when the number of readers increases. Indeed, the main benefit of the beacon mode is that the detection is not affected by the number of users, since the readers do not send any request to the tags (and therefore request collisions do not need to be handled). Actually, the reader just tunes into the proper radio frequency and "listens" to the tag(s) information, reporting the list of visible tags together with their RSSI (Received Signal Strength Indication). However, when multiple tags are sending their IDs, overlap may occur and data may be lost. For this reason, the detection of all the visible tags may take a few seconds (depending on the number of tags in range).

To make our mobile guide as small and as light as possible, we opted for a totally handheld-based solution consisting of Compact Flash (CF) RFID reader with small-sized antenna. The PDA does not need any additional expansion or adapter because the reader plugs directly into the CF slot.

Interfacing to the RFID hardware is achieved via the libraries provided by the hardware supplier (http://www.identecsolutions.com/) whose functions allow the application to configure and interrogate the RFID reader. For each query (queries are performed with the frequency of two per second) from the localisation support, the RFID reader provides the list of visible tags. Every list element contains data related to a visible tag, such as ID and RSSI. The application localisation support keeps a list of all the tags that have been detected at least once. In the application each tag is associated with its last reported RSSI and the time elapsed since its last detection. Only tags that have been detected recently are considered in the computation. The best tag is always the one with the highest RSSI. However, a "new tag event" is generated only if the new tag is the best candidate for n consecutive queries. The value of n is specified in the application settings to achieve a trade-off between reliability and speed of the localisation: the higher the value, the more reliable the localisation will be, but also the more time it will take the identification of the closest artworks. Therefore, the value of *n* must be chosen carefully, especially when tag density is high (i.e., artworks are very close), in order to avoid erroneous detections. To facilitate localisation an RSSI threshold is used to adjust the reader sensitivity. Lower sensitivity makes the reader report only the nearest tags and simplifies the computation.

5.3. Core

The core implements data structures useful for the upper layers (e.g., support for configuration and help) and the XML parsers. It also provides functionality used to update the information regarding the state of the players, to connect to shared stationary displays and to exchange information among PDAs, and therefore implements algorithms for managing sockets, messages, and group organisation. The Core also contains a set of methods for interfacing with the speech synthesizer engine, which is a package provided by a third part as it has been previously mentioned. The *Core* has three parts: the logging facility, the database interface and the communication module.

5.3.1. Event logger

Event Logger manages the recording of relevant actions (e.g., location update, game solved) on a log file. It also raises events that are caught from the *User model* to update the user profile.

5.3.2. DB access interface

This module provides a suite of functions for accessing the museum definition layer. XML parsers and methods for getting the resources facilitate the access to the stored information.

5.3.3. Communication functions

The state of the application is shared among all the mobile devices in use. For keeping the state up to date, real-time communication is fundamental. When a new device is joining the session (i.e., the mobile application is being initialized) it multicasts a discovery message and the first answering device (mobile or stationary) provides the current global state.

Communication may also involve two specific entities such as a mobile device and one of the stationary devices of the environment. The mobile device explicitly requests connection to the large screen, and data exchange is possible because the former device is aware of the latter's IP address after the discovery phase. After establishing the connection, the mobile device sends some messages to the stationary device containing the shared state and the state of the mobile application (current section, visited artworks, solved games). Every relevant event is then sent by means of update messages (score changed, artwork accessed, ...) to the large screen, which keeps its internal state up to date. The uploading capability is also fundamental for the stationary device to furnish a real-time feedback when the user is previewing the museum sections through the large screen. Fig. 12 shows how the map of a room is presented on the mobile and on the stationary device.

5.4. User model

The user model indicates the level of preference about information attributes and related values contained in the museum database. Information regarding user movements, visited artworks and games result are used to update the user model. Accessing a game associated with an artwork is considered as an expression of interest towards the corresponding knowledge, and, if the game is correctly solved, the level of interest is considered even higher. Likewise, if the user moves close to an artwork, this is considered another type of expression of interest in it.

5.5. Visit

The visit layer supports the presentation of the current room map and a set of interactive elements. Each artwork is associated with an icon identifying its type (sculpture, painting, picture, ...) and it is positioned in the map according to its physical location. By selecting such icon, users can receive detailed information on the corresponding artwork. In addition, during the visit, users can access games, receive help, access videos, change audio parameters and obtain other info.

The whole guide application has been developed in C# using the Microsoft .NET Compact Framework. A custom graphical component, which we called *MuseumMap*, is embedded on the main form in order to present the visit and enable an interactive access to its features. *MuseumMap* waits for location update events produced by *Localisation*, queries the database for the room layout (type and coordinates of the items) and uploads the map representation highlighting the neighbouring artworks. *MuseumMap* also manages graphic-oriented computations such as scaling (for adapting the drawing area to the device display), and solves selection correlation between the touched area and the actual object.

5.6. Games

Museum visit is extended by the games layer. The games are defined through XML-based representations associated with the various types of games, to allow possible modifications or additions through any text editor. The game layer accesses the communication services (to inform every player of the team about the score updating) and invokes the XML parser.

6. Customisation facilities

Museums are dynamic entities and often they change their contents. Thus, it is important to allow their curators, who might not have any programming experience, to be able to (re)configure the mobile guide. In order to facilitate the creation of content for a new museum or changes for an existing one, we have developed a specific editor. This tool accesses an XML-based description of the museum, which defines rooms and artworks position as well as additional information. Starting from this data, which includes the photos and a database of the artworks, the editor allows to:

- Create museum rooms or sections just drawing polygons on the overall museum map.
- Create links between rooms using icons (e.g., arrows or stairs-shaped) or text boxes.
- Add, remove or change artwork icons and select the associate photo, database information, video and text files (used by the TTS engine to create the vocal comments on the fly).
- Create help sections.
- Insert on the map of the overall museum interactive parts for quick room selection (allowing the user to manually change room by clicking the interactive part).

Fig. 16 shows the interface for editing the virtual environment. The rooms and the associated items are listed on the left side on a tree-manner (elements can be expanded for editing). The same strategy is used on the right panel for listing the resources (e.g., photos). The central part is dedicated to the currently edited room. New elements can be added to the room by just selecting the corresponding icons from the toolbox and locating them in the museum map through drag-and-drop.

The games can be created and modified with the game editor, which has an interface similar to the mobile guide one. There is

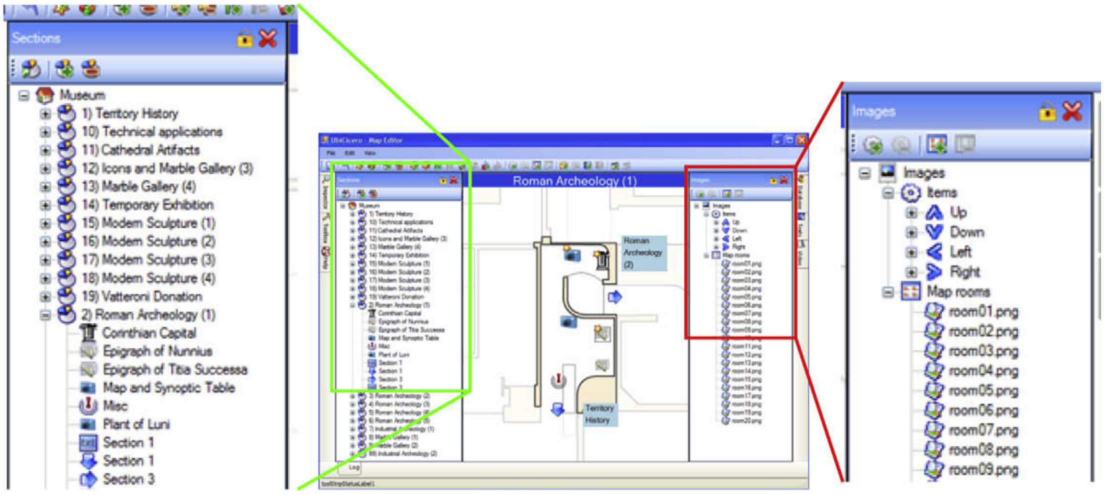

Fig. 16. The sections editor interface.

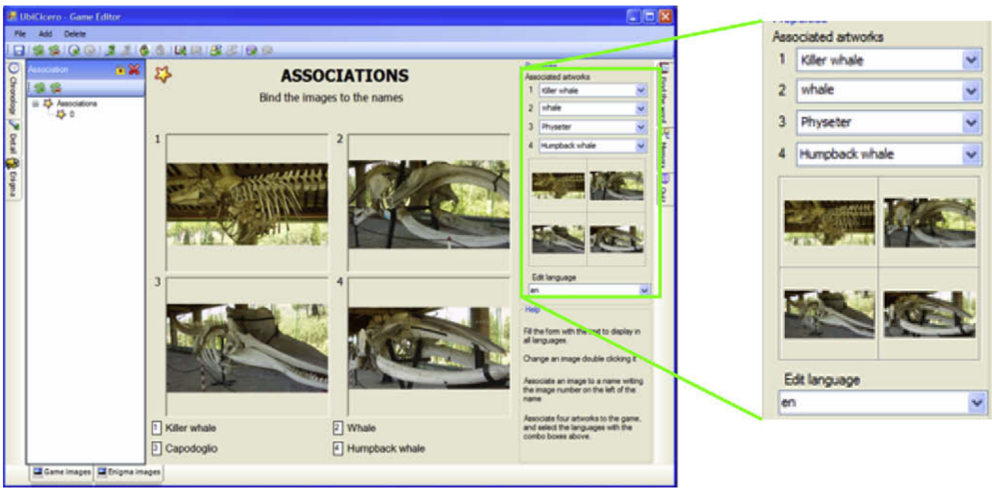

Fig. 17. The games editor interface.

one template for each type of game available: the user can modify the texts and provide a solution for the game (Fig. 17 shows an example). After completing the game, a star on the corresponding artwork icon will automatically appear.

After saving the configuration, the tool generates a collection of XML files that define the database and that can simply be deployed on the PDAs and on the stationary devices.

This high configurability of the environment has made it possible for one collaborator of our laboratory to create a guide for the Natural History Museum of Calci in about 1 week. Providing a museum editor is judged important especially after the application deployment, since the layout of a museum can often change. The museum curators can thus directly update the data of the guide, without knowing the underlying implementation languages.

7. Evaluation

We first performed some early, quite informal user tests, which were useful for providing some basic suggestions on the guide user interface and determining the most suitable type of RFID technology. Indeed, after such early tests we decided to discard the use of active RFID tags based on “Respond” technology since they cover too large distances and are more difficult to manage.

Once we obtained a more consolidated version, we performed a first evaluation of the multi-device guide involving seven users in the Marble Museum and five in the Natural History Museum (8 males and 4 females), with an average age of 36.4 years old; three of them with secondary school education, the others university graduates.

Users were requested to read a short introduction about the application, explaining its objectives and main features, and they were also instructed about the tasks they were expected to carry out. They were invited to test two versions of the guide. One version was equipped with the multi-device support, the RFID module for detecting proximity of artworks within a museum room, and the adaptive features triggered by the user model. The other version was a *basic* mobile one without the multi-device support, no adaptation capabilities and no RFID support. The user interactions were automatically logged in the mobile devices.

Users were asked to visit some sections of the museum using the two versions. In order to test the additional features provided by the enhanced prototype, they were also requested to access and solve some games and perform at least one splitting between PDA and a large screen available in the museum. Half of the users were requested to start the visit using the enhanced version and then continue with the basic one, while for the others the opposite sequence was followed. Afterwards, the users had to fill in a questionnaire.

Regarding users' background information, almost all test participants reported not having had much familiarity with PDAs (a few reported having used one for navigational applications or calendars, etc.). A few of them had previous experience in using digital museum guides, even fewer reported having used digital games in museum settings.

Regarding the more technical aspects of the mobile guide(s). in the evaluation questionnaire we asked questions that were aimed at evaluating the specific features of the two guides (the vocal part of the guide, the games, how the splitting was supported, etc.), and also questions that were aimed at comparing the two versions of the guide. Generally, the questions asked for a rating on a 1–5 scale, in which 1 is the worst case and 5 is the best.

The enhanced museum guide feature for detecting the current position of the user and presenting information dependent on such position (e.g., by surrounding the icon representing the artwork(s) closest to the user with a frame) was judged useful by testers ($M = 3.58$; $SD = 1.16$). Their comments indicated that this information was useful for identifying an artwork and associating it with an icon in the digital map, which is also useful for orientation goals, especially in large museums. However, delays in the localisation support were noticed by some users, which caused some hesitations during the visit. In addition, they judged useful the support given by the enhanced guide for presenting descriptive information about the currently closest artworks ($M = 3.75$; $SD = 0.87$). However, some of them pointed out that additional interaction is needed for deactivating the description of artworks which are not judged interesting.

As for the multi-device support, users appreciated the guide/games offered in the PDA version ($M = 3.67$; $SD = 1.23$) and in the desktop one ($M = 3.82$, $SD = 1.47$). Also the way in which the functionality was split between the PDA and the large screen was rated quite positively ($M = 3.2$; $SD = 1.48$). Regarding the splitting, some of them noted that it might not be easy for the user to follow a description that is displayed partially on a PDA and partially on a large screen because of the consequent division of attention.

When they were asked whether they would have preferred to use a keyboard and mouse for interacting with the large screen instead of the mobile device, six reported preferring the current interaction support, four would have preferred mouse and keyboard, while two replied it was indifferent. The most often given reason for preferring a mouse and keyboard was familiarity with such devices. The ability to select the way to visualise artworks in a section, through icons and/or artwork previews on the large screen, received positive ratings ($M = 3.38$; $SD = 1.5$).

As for the multimodality, testers were also asked to report whether they preferred the vocal part of the user interface or the graphical one. Six users preferred the vocal part, reporting that the vocal support allows the visitors to look at the artworks while listening, which is especially useful in a museum setting where the visual channel is often already overused. One user reported preferring the purely graphical version, three users preferred the vocal + graphic version, two did not report any substantial preference. Regarding the effectiveness of the vocal part, it was rated positively ($M = 3.67$; $SD = 1.22$).

When the users were asked which of the two visits they appreciated more, the majority of them (10 out of 12) answered the enhanced prototype. Among the reasons given by the two people who preferred the basic prototype version, one reported the ability to make a quicker visit, and the other felt that the basic version allowed for more user initiative than the enhanced version.

Regarding the benefits/problems that were reported in the enhanced prototype, on the one hand, some users appreciated the capability of providing more personalised information and also the indication of the localisation within a room, which on the other hand was also reported as the cause of introducing some delays in the responsiveness of the guide.

The user interface for visualising the games associated to the various artworks in the museum was rated relatively high ($M = 3.67$; $SD = 1.37$), though some users reported problems in selecting the icon representing the game associated to a specific artwork. Regarding the content of the games, users judged them quite well ($M = 3.73$; $SD = 1.19$), though they emphasised the importance of providing games that are easy to solve . Games were judged amusing ($M = 3.67$; $SD = 1.15$) and useful in stimulating and improving learning ($M = 3.55$; $SD = 1.37$), one of the preferred games was the quiz.

We also carried out an assessment of the visitors' attitudes to follow the suggestions provided by the system through a log analysis. In our results, 31.5% of the suggestions were actually followed by the users: this can be interpreted as evidence that users tend to visit museums quite autonomously and sometimes do not follow a straightforward strategy in visiting them.

All in all, the application was deemed useful, interesting and with good potential. Users suggested improving the precision of localisation within the museum sections, and the guide's performance in the location support, and simplifying as much as possible the interaction. For example, after the test we changed the modality for activating a game associated with an artwork: initially it was done by selecting the small star icon associated with the artwork icon in the room map but, since it was deemed too small and difficult to select, we introduced a button within the overall artwork description for activating the corresponding game. Some users suggested increasing icon size and picture resolution, to facilitate items selection and also to more easily support the association between virtual and real artworks. Two visitors would have appreciated the dynamic enabling of the available games in order to let the user play only the games related to visited areas, which are the games s/he may be able to solve by benefiting from the museum visit.

8. Conclusions and future work

In this paper, we propose a multi-device, location-aware guide supporting museum visits, which also provides the possibility of enriching the museum visits through individual or collaborative games. Its main contribution is in the ability to exploit multi-device environments, in which users can freely move about with their mobile guide but also exploit large screens connected to stationary PCs when they are nearby. Both the access to museum information and the associated games can benefit from the availability of multiple devices as well as additional services, such as the presentation on the large screen of the locations of other visitors, which are detected through RFID tags. We have also developed an authoring environment, which allows museum curators to easily customize the contents for their museum guides. We have implemented a version of the multi-device guide for two museums (Marble Museum of Carrara and Natural History Museum of Calci). The system is aimed at improving the visitors' experience by extending their interaction with exhibits.

The lessons learned from our experience underscore the fact that museum visitors are increasingly more varied in their interests, background, experience, etc. and all such aspects need to be considered in the design of digital guides, which should be flexible enough to adapt to their needs. In addition, our work also highlighted the importance of understanding how new information technologies might change (possibly improving) the way visitors approach museums. In our study, we have found that multi-device support enabling the users to opportunistically and easily use the devices at hand can be useful for better accessing information related to the museum and its artworks. It can also be useful to improve social interaction and to help users acquire greater knowledge regarding the associated content. The combined use of large screens (whose cost is steadily decreasing) and mobile devices enables solutions that draw on their respective strengths. Users can freely move about and receive information that depends on their interests and location, with the additional possibility to exploit large screens for better accessing more detailed information, additional services and share games and content with other visitors.

Future work will be dedicated to further empirical validation of our guide, aiming to collect both qualitative and quantitative data from a large set of museum visitors, and to identify new ways to enhance personalisation of the guide user interface, in particular in group visits. We will also investigate possible improvements in the localisation capabilities by considering different technologies (such as visual tags). Further multimodal interaction features will be considered, also to obtain a version able to support visually-impaired people through automatic detection of user location and orientation and potential obstacles.

References

Abowd, D.A., Atkeson, C.G., Hong, J., Long, S., Pinkerton, M., 1996. Cyberguide: a mobile context-aware tour guide. Wireless Networks 3 (5), 421–433.

Bell, M., Chalmers, M., Barkhuus, L., Hall, M., Sherwood, S., Tennent, P., Brown, B., Rowland, D., Benford, S., Capra, M., Hampshire, A., 2006. Interweaving mobile games with everyday life. In: CHI 2006 Proceedings Games and Performances, Montréal, Canada.

Bellotti, F., Berta, R., De Gloria, A., Margarone, M., 2006. Guiding Visually Impaired People in the Exhibition. Mobile Guide 06. Turin, Italy.

Bellotti, F., Berta, R., De Gloria, A., Ferretti, E., Margarone, M., 2003. VeGame: exploring art and history in Venice. Computer 36 (9), 48–55.

Brignull, H., Rogers, Y., 2003. Enticing People to Interact with Large Public Displays in Public Spaces, Interact, Zurich, Switzerland (pp. 17–24),

Brown, B., MacColl, I., Chalmers, M., Galani, A., Randell, C., Steed, A., 2003. Lessons from the lighthouse: collaboration in a shared mixed reality system. In: Conference on Humans Factors in Computing Systems (CHI), April 2003. ACM, New York, pp. 577–584.

Cheverst, K., Davies, N., Mitchel, K., Smyth, P., 2000. Providing tailored context-aware information to city visitors. Adaptive Hypermedia and Adaptive Web-Based Systems (AH 2000), vol. 1892. Springer-Verlag, Berlin Heidelberg/New York, pp. 73–85 (LNCS).

Cosley D., Lewenstein, J., Herman, A., Holloway, J., Baxter, J., Nomura, S., Boehner, K., Gay, G., 2008. ArtLinks: fostering social awareness and reflection in museums. In: Proceeding of ACM-CHI, 05–10 April, 2008, Florence, Italy, pp. 403–412.

Ciavarella, C., Paternò, F., 2004. design of a handheld, location-aware guide for indoor environments. Personal and Ubiquitous Computing 8 (2), 82–91.

Dini, R., Paternò, F., Santoro, C., 2007. An Environment to Support Multi-user Interaction and Cooperation for Improving Museum Visits through Games. Mobile HCI 2007. ACM Press. pp. 515–521.

Fraser, M., Stanton, D., Hui Ng, K., Benford, S., O'Malley, C., Bowers, J., Taxen, G., Ferris, K., Hindmarsh, J., 2003. Assembling History: Achieving Coherent Experiences with Diverse Technologies. European Conference on Computer-Supported Cooperative Work (ECSCW). Kluwer. pp. 179–198.

Ghiani, G., Paternò, F., Santoro, C., Spano, L.D., 2008. A location-aware guide based on active RFIDs in multi-device environments. In: Proceedings CADUI 2008. Springer-Verlag, Albacete, Spain, pp. 59–70.

Grinter, R., Aoki, P., Hurst, A., Szymanski, M., Thornton, J., Woodruff, A., 2002. Revisiting the visit: understanding how technology can shape the museum visit. In: Proceedings of the ACM CSCW'02. ACM Press, pp. 146–155.

Hatala, M., Wakkary, R., 2005. Ontology-based user modeling in an augmented audio reality system for museums. User Modeling and User-Adapted Interaction 15 (3), 339–380.

Hindmarsh, J., Heath, C., Vom Lehn, D., Ciolfi, L., Hall, T., Bannon, L., 2001. Social Interactions in Museums and Galleries. Deliverable of the Shape Project (D2.1), Disappearing Computer Project.

Mäntyjärvi, J., Paternò, F., Salvador, Z., Santoro, C., 2006. Scan and Tilt: Towards Natural Interaction for Mobile Museum Guides. Mobile HCI 2006. ACM Press. pp. 191–194.

Leikas, J., Stromberg, H., Ikonen, V., Suomela, R., Heinila, J., 2006. Multi-user Mobile Application and a Public Display: Novel Ways for Social Interaction. PerCom 2006. IEEE Press. pp. 66–70.

Leinhardt, G., Crowley, K., 2002. Objects of learning, objects of talk: changing minds in museums. In: Paris, S. (Ed.), Multiple Perspectives on Object Centered Learning. Lawrence Erlbaum, Hillsdale, NJ, pp. 301–324.

Opperman, R., Specht, M., Jaceniak, I., 1999. Hippie: A Nomadic Information System. Handheld and Ubiquitous Computing. Springer-Verlag, Berlin. pp. 330–333.

Petrelli, D., Not, E., 2005. User-centred design of flexible hypermedia for a mobile guide: reflections on the HyperAudio experience. User Modeling and User-Adapted Interaction 15 (3–4), 303–338.

Petrelli, D., Not, E., Sarini, M., Stock, O., Strapparava, C., Zancanaro, M., 1999. HyperAudio: location-awareness + adaptivity. In: Proceedings CHI '99 Extended Abstracts. ACM Press, pp. 21–22.

Poslad, S., Laamanen, H., Malaka, R., Nick, A., Buckle, P., Zipf, A., 2001. CRUMPET: Creation of User-Friendly Mobile Services Personalised for Tourism. 3G Mobile Communication Technologies, London, UK. pp. 26–29.

Pospischil, G., Umlauft, M., Michlmayr, E., 2003. Designing LoL@, a Mobile Tourist Guide for UMTS. Information Technology and Tourism, vol. 5(3), Cognizant Communication Corporation, pp. 151–164.

Sparacino, F., 2002. The museum wearable: real-time sensor-driven understanding of visitors' interests for personalized visually-augmented museum experiences. In: Proceedings of Museums and the Web, 2002.

Stock, O., Zancanaro, M., Busetta, P., Callaway, C., Krüger, A., Kruppa, M., Kuflik, T., Not, E., Rocchi, C., 2007. Adaptive, intelligent presentation of information for the museum visitor in PEACH. User Modeling and User-Adapted Interaction 17 (3), 257–304.