Agents for distributed context-aware interaction

Augusto Celentano°, Daniela Fogli*, Piero Mussio*, Fabio Pittarello°

°Dipartimento di Informatica, Università Ca' Foscari di Venezia
*Dipartimento di Elettronica per l'Automazione, Università degli Studi di Brescia

{auce, pitt}@dsi.unive.it, {fogli, mussio}@ing.unibs.it

ABSTRACT

This paper addresses post-wimp HCI system design by introducing the concept of *Interaction Locus* (IL), as a coordination of three-dimensional representation, auditory signs and hypermedia information, which informs the user about the specific nature of the part of the environment that he/she's currently entered. The interaction between a user and the IL is observed and mediated by two agents that, borrowing the terminology from ancient Roman religion, are called the *genius loci* and the *numen* of the user. The *genius* knows the information opportunities of the IL and its interaction possibilities. The *numen* knows the user profile and exploration history across several places, and interact with the *genii* of the ILs in order to give them information about how to help the user in his/her visit. The population of *genii* implements a distributed computation, which becomes ubiquitous when we deal with an IL in the real world. The *numen* maintains the state of the computation for its own user adapting it to the context of the interaction.

1. INTRODUCTION

The issues we address in this paper are related to the so-called *post-wimp human computer interaction*, i.e., interaction that goes beyond the limits of the conventional bi-dimensional interaction based on windows, icons, mouse and pointer [7].

The term post-wimp interaction is usually referred to a number of different metaphors and underlying technologies. Immersive/desktop virtual reality, augmented/mixed reality and ubiquitous computing all belong to this group; they have a number of common features that distinguish them from the other post-wimp paradigms. We refer to them as to *positional interaction paradigms* (PIPs).

The first feature that characterizes PIPs is that humans participate to the interaction with their real or virtual body; at least the human position inside the environment is tracked by some input device and considered as one of the main data for interaction. The subjective presence of the humans and their interaction with objects inside the three-dimensional world (that have a physical/virtual location too) is a key feature that distinguishes them from other paradigms where interaction doesn't take into account the location of the user, or where the interactive entities and objects are placed in an abstract workspace.

In PIPs, user motion (body translation and/or rotation) is an essential part of the interaction inside the environment; typically the user moves along a predefined or free path, and interacts along it in different times and in a number of different physical locations.

If we compare PIPs to the common-sense notion of experience, we can notice a number of similarities; experiences are characterized by the direct involvement of humans in the world; experiences are often the summa of different active phases performed in different places and in different moments.

Human experiences offer usually additional features:

- humans participating to a specific experience start with their own background/knowledge that is enhanced as a result of the experience;
- acquirement of new knowledge is often enhanced by the emotional side of the experience; this situation is connected with the direct involvement of the humans and with the stimulations of all the senses;
- humans use knowledge acquired during the experience as a guide for the rest of the experience or for future experiences.

Our proposal is to enhance interactions allowed by PIPs, trying to emulate the positive features of the human experience; that is the reason why we refer the enhanced interaction activities resulting from our proposal as to *experiential interaction paradigms* (EIPs).

A key component of our proposal is the introduction of a set of *proactive entities* or *agents* [11]; agents act in the virtual/real world keeping track [2]:

- of the initial background of the human involved in the experience;
- of the relevant interaction in order to augment knowledge and to facilitate further interaction.

In this way, agents become aware of the context and able to adapt their behavior to it.

Agents can be embedded in virtual world entities or in the virtual layer that enhances interaction in the augmented/mixed reality entities. When the entities are real or mixed-reality ones the embedded agents constitute an ubiquitous computing context.

2. RELATED WORK

Researchers are becoming aware of the strategic importance of the context in which the interaction process occurs: the use to which the system can be put by the user is not a feature of the system organization but of the domain in which the user operates and belongs to a description of the system in the context of user activity [17].

The influence of context on the interaction process has been explored from different points of views: some researchers study the computational and physical context [6], while others stress the anthropological, cultural, skill aspects [16].

Several research groups tried to cope with these problems by proposing systemic models of the HCI processes. Barnard et al. describe the user-interaction process as a syndetic system, i.e. a system binding up sub-systems of different nature, the human as a psychological sub-system, on the one hand, and the computer as a computing artifact, on the other [4]. The nature of interactive computing as well as the design process of interactive software are being explored and some researchers claim that software in general, and interactive software in particular, is a new medium [1].

Significant results have been obtained adopting a specialized view of context.

For example, the user position is the relevant input information exploited in the *Active Badge* system [18], developed in early Nineties. The system provides a means of locating individuals within a building by determining the location of their Active Badge. The badge transmits periodically a unique infra-red signal; its location (and hence that of its wearer) is determined by a network of sensors that detect these transmissions. The system has been used for a variety of applications, including re-direction of incoming phone calls to re-routing of personal X sessions at any suitable display near to the user position. The latter application allows to eliminate the traditional log-in process, enabling the user to easily make use of computing resources which are to hand [15].

Dey et al. [9] observe that most context-aware applications are built in an ad hoc manner, so determining a lack of generality and requiring for a development of each new application from scratch. The Context Toolkit is the architectural proposal presented in [8][9] which supports the development of context-aware applications by providing the general mechanisms required by context. The toolkit includes three components: context widgets, which mediate between a user and the environment; aggregators, which mediate between applications and elementary widgets, hiding the details related with context-sensing mechanisms; interpreters, which are used to transform low-level context information into higher one. A number of context-aware applications have been developed with the support of this architecture [9].

Adapting the application behaviours to the context is a step towards what Weiser and Brown call "calm" interaction [19], in which users are required to make few active interactions, thanks to the choices made by the devices on the basis of the known context.

Most of the systems described in the literature refer to the single user interaction and do not consider the interactive experience as the result of a set of coordinated interactions performed in different times and places. In our opinion, these interactions allow the development of specialized tools but are inadequate for the cases in which the user must rely on his/her everyday experience in the real world.

In this paper, we make a first step to study these problems by proposing a functional approach to the development of context-aware applications, based on the concept of Interaction Locus enriched by agents in charge of supporting human-computer interaction.

3. THE INTERACTION LOCUS

3.1 Motivation

The *interaction locus* concept was introduced in the context of a research that aimed at finding weaknesses in the current interaction modalities in 3D environments.

Navigation is a primary task for what concerns interaction in real and virtual environments: it is often the main activity performed by the user and in any case, due to the explorative approach that characterizes such interaction, it is a prerequisite for more sophisticated behaviors; in spite of that, current interfaces for virtual and mixed reality don't offer satisfactory solutions to support this task.

Therefore, in order to increase the quality of interaction in three-dimensional worlds a better formalization of the navigation problem was given [13], distinguishing three related primary issues: identification of the scene structure, orientation and navigation. These issues were considered also in the context of real world interaction, in order to extrapolate some useful hints for artificial environments. A number of guidelines for virtual worlds were critically extrapolated from the analysis of the real world and applied to virtual environments; these guidelines included also a strong attention towards *multimodality*, which means giving and receiving different information simultaneously from different senses.

Vision, hearing and touch are all part of the multimodal experience; the efficiency of navigation and the task completion are often related to the simultaneous stimulation of them. This apparent redundancy is often useful in many situations and, in some cases, it can be vital.

In particular, the authors used multimodality to overcome some limits of vision introducing the concept of *interaction locus*, which is an extension of the architectonic concept of *venue*, obtained associating to it multimedia components bearing rich information opportunities; more precisely, the *interaction locus* was defined as *a summa of coordinated three-dimensional representation, auditory signs and hypermedia information*. It is the basic element to give a structure to three-dimensional space. It allows the author of virtual or augmented worlds to inform the user about the specific nature of the part of the environment that he/she's currently entered.

The *interaction locus* helps the user to identify the areas that have a morphologic coherence and, in parallel, that are characterized by homogeneous interaction modalities. It is both a help to recognize a part of the scene that has a visual unity (e.g., a room, a square, a glade in a forest) or that is characterized by homogeneous behaviors (e.g., an industrial plant where some actions are forbidden for security reasons). With the *interaction locus* concept the author is enabled to superimpose to a *raw* scene a virtual entity whose task is to inform the user about the nature of the part of the scene he/she's entered and to present and to mediate the possible interactions inside the area controlled by it.

3.2 An application experience

A number of implementations were built to show how to practically use this concept, in cultural exhibitions promoted by Palazzo Grassi in Venice [12], the worldwide known cultural institution promoting exhibitions ranging from Archaeology to Modern Arts.

As an example we discuss *The Artist's Eyes* environment designed in occasion of the *Picasso 1917-1924* exhibition held in 1998. It is a didactic environment for the Internet user conceived as an invitation to the real exhibition in Venice.

Pavilions are the key elements in *The Artist's Eyes* environment; they are used to build a visible metaphor of the artist's life and work between 1917 and 1924. They are built and regularly placed on an infinite plane (a pavilion for each year) using two orthogonal panels: a photograph about Picasso's life is exposed on one panel; an artist's work, chosen to represent the artist's activity for each year between 1917 and 1924, is exposed on the other panel; hypertextual information about the artist appears on the left side of the environment every time the user approaches to the pavilion.

Besides, users can zoom in and out just clicking over the photograph and the painting; in this case textual information about the chosen image is automatically displayed. Moreover, each pavilion is also a portal towards a three-dimensional animation evoking some key aspects of the artist's composition; the artist's work exposed on the panel becomes the theme of the animation; for example the work chosen to represent the artist's activity for the 1917 pavilion is the curtain painted by Picasso for *Parade*, a ballet written by Jean Cocteau on the basis of Erik Satie's music; when the visitor clicks over the Harlequin hat disposed in the center of the pavilion, the pavilion's panel vanishes and the visitor is introduced into a theatre adorned by the curtain painted by the artist, while the notes composed by Satie for the ballet are played. In this way three-dimensionality and multimodality (to give and to receive multiple simultaneous information) work both as an interface to access information and as a tool to add a further level to the user's experience.

On this environment a set of Interaction Loci has been superimposed, which will be discussed later.

3.3 Definition of Interaction Locus

In the context of this paper an *interaction locus* (IL for brevity) is a connected portion of space characterized by:

- the possibility of perceiving when a user enters and exits the IL; i.e., the IL is well defined and identifiable even if it has no visible borders;
- the presence of an underlying base world, real or virtual. In a virtual environment it is a base geometry describing a realistic, simulated or imaginary world. In a mixed reality environment it is the part of the world to which experiences are attached. In a real world with embedded computing devices it is the part of the world in which the user can interact with the devices. The base world is independent from the IL, and several ILs can be superimposed on it, each corresponding to a different kind of experience the user can live;
- the presence of identifiable interaction devices which support the exchange of information between the user and the world.

The interaction devices receive user input and provide information to the user. They are of three types:

- *interactive objects*, i.e., objects on which the user operates directly, which change their state or their appearance as a consequence of the user interaction. An example of an interactive object in a cultural heritage context is an artwork (e.g. a multimedia installation) revealing information (e.g. by animating itself) when the user approaches or touches it;

- artifacts, i.e., mediators of the interaction between the user and the world, which make evident to the user interaction opportunities which would otherwise be unknown [14]. An example of a mediator is a metaphoric object serving as a button for accessing information when pushed;
- dynamic information objects, which modify their state or appearance, e.g., show or hide themselves, as a consequence of the user interaction on other devices in the environment. An example of information object is a panel on which a text scrolls showing the user information about the place he/she is visiting.

With reference to The *Artist's Eyes* scenario, we can divide the ILs that map this environment in four different classes: the *Intro locus*, the *Pavilion locus*, the *Between locus* and the *Final locus*. The *Intro* and the *Final loci* are the simplest, because they act, respectively, as symbolic places where simple information and interaction activities take place and mark the beginning and the end of the virtual exhibition; the *Between locus* presents to the user some objects that represent the activity and the interest of the artist between the years considered by the exhibition. The *Pavilion locus* is more complex; a variety of interaction can take place inside it, as we described above.

Objects inside interaction loci can be categorized according to the definitions given previously: there are artefacts (the arrows and the harlequin's hat) and interactive objects (the works, the photographs, the ship at the end of the exhibition).

As a more general remark, we note that on a same base world several different interactive experiences can be superimposed, which share the interactive objects which are part of the base world, and use different artifacts and possibly different information objects. Different experiences can be related to different users or to a same user in different times [5]. As a consequence, in the same base world different ILs can be defined, each corresponding to a different interactive experience.

4. AGENTS SUPPORTING CONTEXT-AWARE INTERACTION

The interaction between a user and the environment (i.e., the interactive objects and the artifacts) is *observed* and *mediated* by two agents, one associated to the *locus*, and the other associated to the user. The agent bound to the IL knows the information opportunities of the specific place and the interaction possibilities, therefore is able to assist the user in exploring the place. The agent can perceive user behavior and actions, such as:

- entering and exiting the IL;
- movement, path, focus of attention, point of view;
- interaction with artifacts and interactive objects;
- user dynamics, e.g., speed, time spent in moving and standing in front of information objects, pauses, etc.

The agent is of course limited in perception to what the user is able to manifest of him/herself. Therefore the agent is bound to the presence of a number of "sensors" in the scene, which are the devices which sense the user actions. We will discuss later some issues about sensors and their role.

The way user movement is perceived depends on how the user is represented in the virtual space and varies according to the nature of the environment (virtual, real, mixed, embedded) and the technology used. For example, in a VRML environment the user is represented as a mobile, virtual camera with specific viewing direction and angle; in a real space, the user can be provided with a positioning device (e.g., GPS), or can be tracked by a series of sensors (as in a kind of guided tour). This is a type of interaction which does not raise visible effects but modifies the possibilities that the user has to perceive objects and to interact with them.

Borrowing the terminology from ancient roman religion we call such an agent a *genius loci*, a kind of local divinity who takes care of the place by giving the visitors the opportunity to get most benefit from its exploration.

Also the user has his/her own genius, that we call *numen*, a kind of *guardian angel* who follows the user during navigation by accumulating and managing knowledge about him/her. The *numen* knows the user character (the profile), can accumulate the exploration history across several places, and is able to interact with the *genii* of the different places in order to give them information about how to help the user in his/her visit.

The interaction between a user and a rich and differentiated environment is therefore mediated by two kinds of agents which accumulate, maintain and exchange knowledge about the user (the *numen*), and the interaction place (the *genius loci*).

Figure 1 pictorially describes this scenario. The *numen* and the *genius loci* communicate in a way that will be detailed later. According to the roman iconography, evidenced in studies such as [3], the snake in the figure represents the *genius loci*.

The role of the agents is to provide support to the user for better interacting with the environment. In order to be effective they should therefore exhibit a *proactive* behavior, by anticipating the user needs in terms of actions on interaction objects and artifacts, and level and quality of information provided. In order to get this goal, they organize their tasks in a collaborative way [10].

Moreover, both kinds of agents are context aware. The context of the *numen* is constituted by the user and by the *genius loci*. The *numen* adapts its behavior according to the data obtained by the *genius loci* and to its knowledge of the current state of the user. The context of the *genius loci* is constituted by the user and by the *numen*. Also the *genius loci* determines its own behavior according to the data on the user obtained by the *numen*, to its current perception of the user and to its knowledge of the IL itself.

5. A SYSTEM OF COMMUNICATING AGENTS

A virtual environment is described by a system of ILs and by the population of associated *genii loci* agents. A *numen* following the user in his/her exploration of the environments successively interacts with the *genius loci* agents and determines their actions in support of the user. A *genius loci* observes the user behavior and proactively modifies the effect of user interactions in order to improve it. For example, let us consider a paintings gallery where each painting has a proximity sensor, which enlightens the painting when the user approaches, and has a display aside, where multimedia information can be recalled by selecting it with a set of buttons. A user systematically interested in getting information about the painting technique (which can be accessed by an appropriate button) would benefit form an anticipation of his/her need so that by approaching the painting the display is automatically updated.

An agent could observe this behavior and, by recognizing a recurring pattern in the user actions who, after approaching the painting, selects the "technique" button, could associate the display of the required text to the user movement, rather than to the explicit button selection.

In order to have this behavior, two things must happen:

- the agent must know which information is available at that place, and
- the agent must have some knowledge about the user previous actions in order to be able to infer a modification in the interaction dynamics.

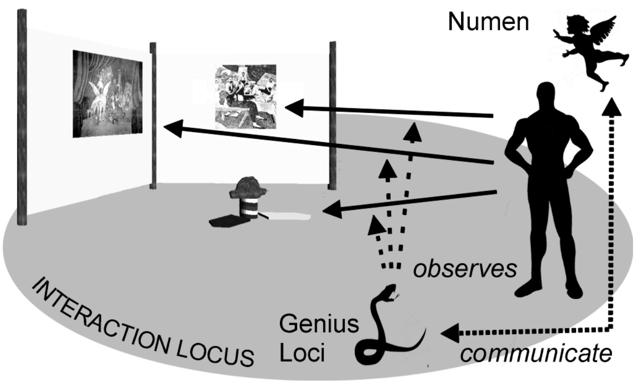

Figure 1. Agents supporting user interaction

Such a knowledge cannot be maintained by only one agent, due to the differences in the various places (the different *interaction loci*), and to the differences in users. A system of agents, like the two kinds proposed, communicating through a suitable protocol, can integrate the knowledge needed at each place and for each user in order to improve the interaction.

The communication protocol between a user *numen* and a *genius loci* is defined by the following steps:

- The protocol is activated when the user enters an IL. As we have discussed earlier, how the user position and movement are tracked is not a major concern, since it can be done in real as well as in virtual or mixed environments.
- The *genius loci* has the knowledge about the *locus* extension and borders, therefore reacts to the user and starts a dialog with the user *numen*, by asking information about the user. The *numen* knows two kinds of information: a user profile, which is a static collection of properties and data about the user, and a user history, which is a set of data collected during exploration, transmitted by the *genii* of the other *loci* visited.
- The *numen* gives the *genius* the information requested, which allows the *genius* to modify, if needed, the properties of interaction in that *locus*.
- The user then interacts in a more effective way, and the *genius* is able to further analyze the user behavior, by possibly inferring new rules about how to improve the user interaction. It is a kind of learning process, driven by the set of rules which are the *genius* knowledge about the interaction opportunities.
- The *genius loci* can make deductions about the user behavior, and change, possibly subject to user's approval, the interaction dynamics. The changes have effect only in that place, since the *genius loci* has no knowledge of interaction opportunities in other *loci* [1].
- On exiting the interaction locus, the *genius loci* returns to the *numen* the result of its deductions. The *numen* decides, according to its own knowledge base (i.e., the user profile and the accumulated information about other *genii* deductions), how to consider the new ones, which are compared and integrated according to the *numen* rules.

[1] We do not enter into details about how the proactiveness of the agents is submitted to the user, who could be asked or not to accept or refuse the agent deductions.

When entering a new interaction locus, the dialog is repeated between the *numen* and the new *genius*. According to this protocol, the *genius loci* maintains knowledge about the place, i.e., it decides how to interpret the user actions according to the place's interaction opportunities, and what kind of deductions to make. The *numen* maintains knowledge about the user, i.e., it decides what deductions made by the *genius* must be considered and how they can be integrated with the static profile and the previous user behavior.

The *genius loci* can maintain also a cumulative knowledge about the behaviors of several users in order to proactively hint the designer of the interactive world about the effectiveness of the design of the locus and of its interaction opportunities. Except for this case, the *genius* does not retain information across the users. Each new user makes the *genius* start from an initial state, which is the state transmitted by the *numen* upon entering the IL.

6. CLASSES OF *LOCI* AND AGENTS

6.1 The agents in *The Artist's Eyes* environment

For a better explanation of the role of the agents we'll examine how in *The Artist's Eyes* environment interaction may be enhanced by their presence.

The case study is an occasion to illustrate also another aspect of the proposal: the organization of ILs and agents in classes. Each class collects the instances of IL homogeneous from the point of view of the interactive experiences possible in them. As a consequence, also the *genii loci* are organized in classes, i.e., they are typed.

When entering an interaction locus a *numen* transmits to the *genius* (or, a *genius* asks the *numen* for) the information which are related to the locus type.

Typed agents can be associated to each IL; besides, an additional agent (the *numen*) has to be associated to the user.

The goal of the agents is to minimize the user interaction effort.

In the case of a 3D environment structured in a manner similar to the Picasso world, a system of coordinated agents allows activities such as:

- *Widget activation forecasting*; the *genii loci* monitor the user interaction inside the *Pavilion IL* and communicate the behaviours to the *numen*; each time the user enters an IL belonging to the class *Pavilion* the *numen* communicates to the local *genius* the previous user history; if the user systematically clicks the Harlequin's hat each time he/she enters a pavilion, the local *genius* will activate automatically the animation without the user's intervention;
- *Zoom forecasting*; if the user, entering a pavillion, systematically clicks the Picasso's photograph to have a zoomed view of it, the local *genius*, informed by the *numen* about the previous user activity in other instances of the same class, automatically activates the zoom function;
- *Motion forecasting*; the *genii loci* related to the *Between locus* typology monitor the presence time of the users inside the area; they communicate the information to the *numen*; if the user's presence inside the environment is lower than a certain time the *numen* will deduct that the user is not interested in this typology of locus and will communicate to the following *genii* of the *Between locus* typology to teleport the user to the next locus.

6.2 Knowledge classes

We need to assume that the knowledge classes associated to different classes of ILs are distinct, even if not necessarily disjoint. It would be desirable to share part of knowledge about different kinds of experience. For example, taking the cultural heritage experiences as an example, it would be desirable to share in different kinds of museums, and with different experience classes, some common behaviors about how to approach and consult the artworks [5]. While it is simple to share behaviors on the basis of the static user profile, to consider in a consistent way the meaning of the user dynamics in different places is much more complicated. However we do not elaborate further on this issue here.

Interaction inside a specific locus can be conditioned by the successful performance of previous interactive phases in other loci; for example, in e-learning contexts, access to loci can be conditioned to the solving of tests in the previous phases. In these situations each *genius loci* has an associated set of conditions to be satisfied in order to allow navigation and/or interaction inside it.

In e-learning contexts, when the user interacts in *locus* A, the *genius* of A observes the successful solving of tests related to this *locus* and communicates the results to the user *numen*; when the user enters subsequently in *locus* B, the *numen* asks to *genius* of B the conditions required for access. The *numen* receives the *genius* requests, compares them with the results of the previous experiences in other *loci* and sends the appropriate answer to the *genius* in order to gain access. Of course, negative answers to the conditions required by the *genius* result in access denials or limited interaction inside locus B.

It should be evident from the previous discussion that proactivity appears in the agents in different ways. The *genii loci* are activated by the user entering the IL, and are initialized with information provided by the *numen*, therefore they are proactive in a very limited sense. They have deductive capabilities, but can propose a different user behavior only when the user is currently acting within their scope.

They are however proactive toward the designer of the interactive world, since they can accumulate knowledge about recurring behavior of different users, therefore suggesting the designer some modifications that could generally improve the user interaction.

The *numen* exhibits most of the proactive capabilities to the user, deciding what information to use, and how, and whether to make the new information available to a new *genius loci* or not. Surprisingly, however, their proactivity does not manifest directly to the user, but is mediated through the *genii* of the ILs that change the interactive objects response to user actions, or suggest users with new interaction opportunities.

7. A GLANCE TO IMPLEMENTATION

Such a functional schema and communication protocol can be implemented by assuming that the interactive environment is made of two kinds of devices: *sensors* and *actuators*. A sensor is an input device that the user can activate directly or as a side effect of movement, which transmits a signal to an actuator, i.e., an object which can change its state and provide some feedback to a user.

According to this view the *genius loci* is a filter agent who mediates between the user actions perceived by the sensors, and the interactive objects. By mediating the objects response, the *genius* can modify the relationships between sensors and actuators, e.g., by introducing by-passes which correspond to the deductions made on the user behavior. Figure 2 depicts the role of the *genius loci* according to this view.

The agents functions are differentiated according to their role. A *genius loci* can:

- recognize a recurring pattern in the user behavior with reference to interaction;
- evaluate the importance of some information, e.g. by tracking the time the user spends in some places or accessing specific information types. The *genius* can also get information by tracking the user movement, direction, focus of attention, if such data are available in the environment implementation;
- capture errors in terms of useless interaction, i.e., interaction that does not modify the system state, does not access information, etc. If such errors are systematic they can be symptom of a poor design;
- change the relationships between sensors and actuators, and in general change the behavior of the place in terms of relationships between the user actions and the system response;
- deduct and suggest changes in the design which can improve the interaction independently from specific users.

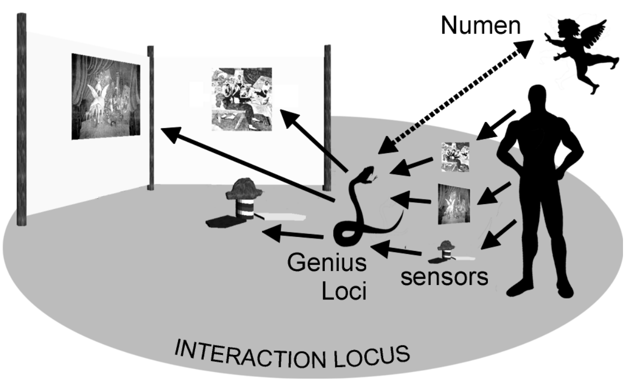

Figure 2. The *genius loci* mediating the user actions

A *numen* can:

- evaluate the importance of the information that user finds in the different places of interaction and the interest for the different types of IL, in relation to the static user profile, by computing recurrences in the user behavior across different ILs, e.g. by measuring the time spent in different types of IL, or the iteration of some experiences, etc.;
- evaluate and integrate the information returned by the *genii loci*;
- suggest the user about interesting or not interesting IL types;
- in some applications, e.g. e-learning applications, the *numen* behaves as a tutor who can force the user to adopt a specific order of visit of different places.

If we model such a system as a set of coordinated processes, the *genii loci* implement a distributed computation system, and the *numen* maintains the state of the computation consistently for each user.

8. CONCLUSION

This paper presents a proposal to enhance Positional Interaction Paradigms by taking into consideration the experience acquired by observing user interaction. To this aim, the concept of Interaction Locus has been considered as the basis for the new experiential interaction paradigms. In these paradigms, the interaction between the user and the entities within an IL is locally personalized to the user preferences and habits. The interaction becomes context-aware because the strategy of navigation within a *locus* emerges from the interaction of the local *genius loci*, aware of *locus* characteristics, with the *numen* of the user, aware of the current state of the user. On the other side, each *genius loci* is an agent embedded in a real or virtual entity of the environment.

9. ACKNOWLEDGEMENT
We acknowledge the support by Italian Ministry of University and Research, MIUR, in the framework of the project *Specification, Design and Development of Visual Interactive Systems, PRIN 2000,* and ex60% funding programme.

10. REFERENCES

[1] P. Armour, The case for a new business model, *Comm. Acm*. 43(8), 19-22, August 2000.

[2] S. Arondi, P. Baroni, D. Fogli, P. Mussio, Supporting Co-evolution of Users and Systems by the Recognition of Interaction Patterns. Accepted for the presentation at *International Conference on Advanced Visual Interfaces (AVI 2002)*, Trento, Italy, May 2002.

[3] J.Th. Bakker, *Living and Working with the Gods. Studies of Evidence for Private Religion and its Material Environment in the City of Ostia*, J.C. Gieben Publisher, Amsterdam 1994.

[4] P. Barnard, J. May, D. Duke, D. Duce, Systems, Interactions, and Macrotheory. *ACM Trans on HCI*, 7(2), 222-262, 2000.

[5] A. Celentano, F. Pittarello, Classes of Experiences: a High Level Approach to Support Content Experts for the Authoring of 3D Environments, in *Structured Design of Virtual Environments and 3D-Components*, Shaker Verlag, 2001.

[6] G. Chen, D. Kotz, A Survey of Context-Aware Mobile Computing, Research. Technical Report TR2000-381, Dartmouth College, Department of Computer Science, 2000.

[7] A. van Dam, Post-WIMP User Interfaces, *Comm. of the ACM*, 40(2), pp. 63-67, 1997.

[8] A. K. Dey, *Providing Architectural Support for Building Context-Aware Applications*, PhD Thesis, College of Computing, Georgia Institute of Technology, December 2000.

[9] A. K. Dey, G. D. Abowd, The Context Toolkit: Aiding the Development of Context-Aware Applications, *Proc. of the Workshop on Software Engineering for Wearable and Pervasive Computing*, Ireland, 2000.

[10] Z. Huang, A. Eliëns, C. Visser, 3D Agent-based Virtual Communities, in *Web3D '02*, Tempe, Arizona, 2002.

[11] N. R. Jennings, An Agent Based Approach for Building Complex Software Systems, *Comm. of the ACM*, 44(4), pp. 35-41, 2001.

[12] http://www.palazzograssi.it

[13] F. Pittarello, Multi Sensory Guided Tours for Cultural Heritage: the Palazzo Grassi Experience, in D. Bearman, F. Garzotto (ed.), *ICHIM '01, International Cultural Heritage Informatics Meeting*, Milan, September 3-7, 2001.

[14] F. Pittarello, M. Pittarello, G. F. Italiano. Architecture and Digital Exhibitions –The Einstein Tower World. *Virtual Environments 1998 - Eurographics Workshop Proceedings*, Springer-Verlag, 1998.

[15] T. Richardson. Teleporting - Mobile X Sessions, in *Proceedings Ninth Annual X Technical Conference*, Boston MA, January 1995.

[16] *The social and Interactional Dimensions of Human-Computer Interfaces*, Thomas P. J. edt., Cambridge University Press, 1995.

[17] F. J. Varela, *Principles of Biological Autonomy*, GSR Amsterdam, North Holland, 1979.

[18] R. Want, A. Hopper, V. Falcao, J. Gibbons. The Active Badge Location System, *ACM Transactions on Information Systems*, Vol. 10, No. 1, January 1992, pp 91-102.

[19] M. Weiser, J. S. Brown. The coming age of calm technology, In *Beyond Calculation: The Next Fifty Years of Computing*. Springer-Verlag, 1997.