Embodied Mobile Agents

Bill Tomlinson
University of California, Irvine
Computer Science Bldg, Room 430A
Irvine, CA 92697
1-949-824-9333
wmt@uci.edu

Man Lok Yau
University of California, Irvine
127B Computer Science Trailer
Irvine, CA 92697
1-949-824-9333
mlyau@uci.edu

Eric Baumer
University of California, Irvine
127B Computer Science Trailer
Irvine, CA 92697
1-949-824-9333
ebaumer@uci.edu

ABSTRACT

The move in many societies toward individuals having multiple networked computational devices – workstations, notebooks computers, mobile phones, PDAs – radically changes the ways in which people engage those devices. However, we lack interaction paradigms that enable a coherent experience across these technologies. One possible approach to this problem involves the use of embodied mobile agents (EMAs), that is, graphically animated, autonomous or semi-autonomous software systems that can migrate seamlessly from one computational device to another. This paper describes an interactive museum exhibit that was implemented with EMAs, discusses the opportunities and challenges presented by this new form of agent, and consider other potential applications for EMAs. While not a universal solution to challenges of interacting with heterogeneous networks of devices, embodied mobile agents can help to provide a coherent user experience across multiple computational devices.

Categories and Subject Descriptors

H.5.1 [**Information Interfaces and Presentation**]: Multimedia information systems – *animation; artificial, augmented, and virtual realities.* I.3.6 [**Computer Graphics**]: Methodology and Techniques – *interaction techniques.*

General Terms

Design, Human Factors.

Keywords

Animated agents, virtual embodiment, mobile agents.

1. INTRODUCTION

Over the past several decades, computational devices have spread rapidly among many human societies. These devices are now sufficiently common that some people have more than one of them, for example owning a workstation, a notebook computer, a PDA, and a mobile phone. These devices are often networked to each other and to the Internet, together providing the platform for an individual's personal information space. However, the varied interaction paradigms that we use when engaging these devices do not always facilitate a coherent experience across multiple devices. In order for these devices to integrate smoothly together and for people to understand their cooperation, new interaction paradigms are needed.

This paper proposes the use of embodied mobile agents[1] (EMAs) – that is, graphically animated, autonomous or semi-autonomous software systems that can migrate seamlessly from one computational device to another – to help enable the development of more intuitive multi-device interaction paradigms. These EMAs may either be animated characters (personified agents designed to elicit social responses from people), animated creatures (biomorphic agents designed to offer some of the capabilities of animals), or animated objects (virtual agents designed to mimic real-world mobile physical objects). EMAs can help to streamline the process of working with several devices simultaneously, to increase the understanding that people have for how information is flowing among their devices, and to help different people collaborate using the capabilities of their multiple computational devices.

Applications of EMAs could reach across a wide variety of domains. In education, EMAs could be used to build multi-device, agent-based participatory simulations that enable people to interact with and learn about a variety of content domains from ecology to economics to the spread of disease. In an industrial setting, EMAs could be used to visualize the flows of information among human collaborators and their computational devices. In entertainment, EMAs could make possible entire new genres of games, with agents jumping quickly between devices carried by players in real space. Other application domains could include social computing systems, healthcare systems, and simulation & training.

To create these kinds of applications, there are a number of problems that need to be solved. How does an agent's embodied form adapt to the capabilities of each device? How can an agent be perceived by people as the same entity as it moves from device to device? How can an animated agent be socially and emotionally engaging across devices? Is it possible for the affordances of the heterogeneous networked devices to facilitate the presence of agents? By addressing these and other related problems, this research project has demonstrated that EMAs can be a robust and deployable technology.

AAMAS'06, May 8-12, 2006, Hakodate, Hokkaido, Japan.

---

[1] The term "embodied mobile agents" was chosen over "mobile embodied agents" because of research on the latter term deals primarily with the navigation sense of the word "mobile". Here, the focus is on creating embodied forms of mobile agents in the sense of [12].

Embodied Mobile Agents have the potential to change the way in which people interact with our growing networks of stationary and mobile computational devices. From interactive educational systems to new forms of entertainment to multi-device industrial applications, EMAs could have a significant positive impact on society.

2. RELATED WORK

The EMAs in the two projects discussed in this paper have three main characteristics: they are animated (featuring expressive device-specific graphics), they are interactive (by means of engaging interfaces and autonomous behavior), and they are able to operate seamlessly across multiple devices (such as desktop computers, tablet PCs, PDAs and cell phones).

This work is most similar to two previous projects. The project relates closely to the work of Rekimoto [23], whose Pick-and-Drop system allowed people to drag digital objects from one computer's screen to another. The proposed research builds on this idea, but places the autonomy in the software itself, rather than having the transfer happen only when specifically controlled by a user. The work also draws inspiration from McIntyre *et al.* [18], who developed a system in which an embodied mobile agent could transfer among different workstations, interacting with people via webcams at each workstation. The proposed research is different from this work in two ways: first, the proposed work deals with heterogeneous networks – engaging the challenges that come from operating on different computational platforms – and second, it focuses on the moment of transfer as an important part of the process – the animation, sound, etc. that enables people to see where the agent has come from and where it is going.

Several other research projects have also combined physical elements and virtual agents. Sumi and Mase's AgentSalon [27] is an interactive installation that uses animated agents on large displays and handheld computers to facilitate face-to-face communication. PEACH [26] features agents on handheld computers that provide information to museum visitors. Agent Chameleons [20] described agents that can migrate between robotic and virtual platforms. The project described here is different from AgentSalon, PEACH, and Agent Chameleons as this project places greater emphasis on the animated transfer between devices and on seamless cross-device 3D character animation. Agent Chameleons also does not have a tangible interface.

The development of EMAs is based on a wide range of previous work in autonomous agents and multi-agent systems. Embodied agents (those with a virtual or physical body through which they sense and act on their environment) have also been a topic of much recent research, having been studied in fields including robotics [3] and human-computer interaction [7].

Mobile agents have also been an active area of research and development in recent years, e.g., [22]. Gray *et al.* [12] defined mobile agents as "programs that can move through a network under their own control, migrating from host to host and interacting with other agents and resources on each." (p. 8) Note that the term "mobile agent" is sometimes used to refer to an agent that is able to navigate around its world. While the mobile agents described in this paper are also able to navigate around their worlds, the term "mobile" is used here to denote the interdevice capabilities of the agents, rather than their world navigation capabilities.

Many researchers have explored different aspects of mobile agent technologies, including mobile agents for networked virtual environments [9], mobile agents for network management [28] and a virtual machine for mobile agents [19]. However, non-animated mobile agents are not designed to be engaging for the users in an interactive setting, as the users will only see the result, not the process in between. This engagement is especially important for applications in the domains of entertainment and education. To enhance this engagement, the mobile agents described in this document are embodied as graphically animated entities.

To enable EMAs to animate expressively, this research builds on previous work from both academic and industrial contexts. Academic researchers have looked closely at the topics of expressiveness and believability in autonomous agents, focusing most frequently on interactive virtual humans (e.g., [11]). The computer/video game industry has also expended great effort to make engaging embodied agents. Many character-based games have been exceedingly successful financially, selling millions of copies and generating billions of dollars in sales. Some commercial products have explored the use of virtual characters in handheld devices. Tamagotchi is a handheld device inhabited by a virtual character that owners can play with and feed. Animal Crossing is a computer game that allows the transfer of virtual characters from one virtual environment to another using a memory stick.

Special effects technicians have experimented with breaking the plane of the screen as a way of increasing the believability and impact of the characters and events portrayed. For example, a show at Walt Disney World features several Muppets walking back and forth between different monitors. This project also builds on work in distributed display environments, which investigates the ways that multiple monitors may enhance the user experience, e.g. [13]. Other researchers have explored ways to use distributed display environments to represent virtual environments, e.g. [10]. The project described here seeks to have the same seamlessness in transfers between computing machines as exists in these transfers between displays.

The animation of the EMAs moves beyond this previous work in several ways. This research focuses on the moment of transfer from one device to another as a critical component in the success of EMAs. No previous project has sought to produce expressive real-time animation across the boundary between devices. The interdevice animation in these projects involves the close coordination of autonomous behavior, interactive animation, real-world sensing, interdevice communication, interface design, physical set construction, and graphical and sound effects. This research has explored the details of each of these aspects as it applies to interdevice animation, and the persistence of an agent's identity across devices. This emphasis helps to connect the real and virtual worlds by enabling the EMAs to appear to move through real space. People are accustomed to seeing animated entities on single screens, but the movement of an EMA from one screen to another helps to reduce the inherent separation that usually occurs between the physical world and a virtual environment. Finally, it engages the physical sensing capabilities of the multiple computational devices on which the EMAs operate. The Virtual Raft Project [30] used accelerometers, webcams and IrDA to enhance the believability of EMA animation. EMAs sense both their virtual environment and aspects of the real world through a blackboard architecture through which entities in the virtual environment exchange information.

The research described here also builds on previous work in autonomous behavior and interface design. In the area of autonomous behavior, Perlin and his colleagues have done pioneering work in creating synthetic actors, e.g., [21]. Cassell and her colleagues built virtual humanoids with the ability to express themselves like real people, in particular in the area of embodied conversational agents, e.g., [5]. Blumberg and his colleagues have contributed in several areas including agent learning [1].

In order to enable EMAs and people to interact with each other, the EMA system features input and output interfaces that connect the two. These interfaces are informed by previous work in engaging and tangible interface design [15]. In the areas of pervasive computing and mobile entertainment, researchers have explored the use of tangible hardware and social computing to enhance user experience in a reality game space [6]. Borovoy and his colleagues developed Meme Tags that allows the transfer of text fragments among mobile devices [2]. Research has been done to investigate ways to enhance the interaction paradigm between a user and a cellular phone [25]. Hand and head gestures have also been suggested as an interface to a mobile interactive paradigm [17].

3. EMA-BASED EXHIBITIONS

Embodied mobile agents have been implemented in two major projects in the past year – the Virtual Raft Project and the EcoRaft Project.

3.1 Virtual Raft Project

The first, the Virtual Raft Project, is an interactive exhibit involving three standard computer workstations and three tablet PCs. The workstations represent "virtual islands" populated by groups of animated humanoid EMAs and the tablet PCs represent "virtual rafts" that participants can carry between the islands (see Figure 1). When a raft is brought up to one of the islands, an EMA can jump onto it. When the participant then carries that raft to a different island, the EMA can jump off the raft onto this new island. A video demonstrating the Virtual Raft Project can be found at:

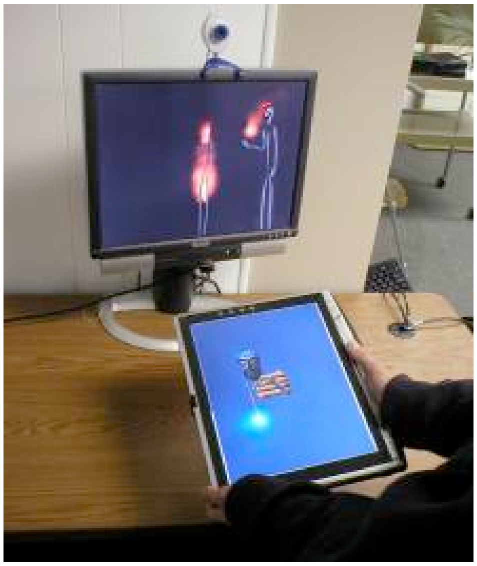

Figure 1: Embodied Mobile Agents may animate seamlessly between a desktop computer and a tablet PC.

http://www.ics.uci.edu/~wmt/movies/VirtualRaftProject.mov

This project enables people to engage physically with EMAs in three main ways. First, the act of moving the tablet PCs between the islands gives people a physical connection to the space and enables them to control the movements of EMAs among the islands. Second, webcams above each of the virtual islands, running a simple background subtraction algorithm, enable the EMAs to react to the presence of people standing in front of that island and respond to their motion. Third, accelerometers in each tablet PC let the EMAs react to the physical motion of the tablets as people carry them, swaying back and forth as the tablets are carried between islands. These three components of the implementation contribute to a sense that the EMAs inhabited the same physical space as people.

The Virtual Raft Project began several key research directions that relate to the proposed work. First, it offered a novel and intuitive interaction paradigm for a person to engage with autonomous animated characters. Second, the project began the exploration of EMAs as viable elements in a multi-user experience. Third, it explored the relevance of EMAs as part of an affective interaction, having a social and emotional impact on people. Fourth, it developed the idea of heterogeneous character animation – that is, animation that enables a real-time animated character to animate seamlessly among two or more heterogeneous devices.

Through four public exhibitions (the ACM CHI 05 conference, the CSCL 05 conference, the Games Learning and Society 05 conference, and the opening of a new building on the campus where it was developed) and numerous smaller scale demonstrations, approximately 800 people have interacted with the Virtual Raft Project to date.

3.2 EcoRaft Project

A second EMA-based project, the EcoRaft Project, was an extension of the Virtual Raft Project developed in collaboration with a team of restoration ecologists and the staff of Discovery Science Center, a science museum in Santa Ana, CA. Instead of featuring humanoid EMAs, this exhibit contained EMAs in the form of animated animal and plant species and was designed to enable 8- to 12-year-old children learn about restoration ecology. The installation consists of three stationary computers that represent virtual islands, and three mobile devices that represent rafts or boxes. Each virtual island represents a different ecosystem. The ecosystems can be populated with hummingbirds, coral trees, and heliconia flowers. Participants can use the mobile devices to transport species from one island to another by bringing a mobile device near one of the stationary computers. One of the virtual islands is the national forest, which has a fully populated ecosystem and can act as the reserve, while the other two virtual islands can be deforested by the press of a button. Participants can repopulate a deforested island by bringing different species in the right order by means of the mobile devices.

In addition to developing further the other themes mentioned above, the EcoRaft Project began to explore the viability of EMAs as tools for education. The Virtual Raft metaphor provides a set of organizing principles for an interactive

Figure 2: A group of children engaged with the Virtual Raft Project.

educational simulation. It distinguishes the operational domain of virtual creatures from that of people, thereby providing a way of explaining how people are meant to interact with the system. It provides inspiration for subject matter that may be effectively conveyed through this kind of simulation. And it provides guidance for the visual, acoustic, tactile, and other elements of the interactive experience. While the platform is not ideal for all topics of education, it does provide a useful framework for exploring topics that deal with groups of people (e.g., social or cultural education), groups of biological organisms (ecology, evolution, etc.), movements between those groups (invasive species, cultural migration, etc.), and potentially other topics (economics, math, languages) as well [29].

The EcoRaft Project was exhibited in the Emerging Technologies program at SIGGRAPH 2005, and was shown to approximately 2000 people over five days in summer 2005. In addition, the project was exhibited publicly at Discovery Science Center on two occasions, and was shown to a group of children in the research lab where it was developed (see Figure 2), with a total of approximately 200 participants.

4. IMPLEMENTATION

The hardware for both the Virtual Raft Project and the EcoRaft Project includes three stationary computers and three tablet PCs. All the computers have 3D graphics cards, infrared (IrDA) sensors, and Ethernet connectivity. Each stationary computer is connected to a webcam and each tablet PC has built-in accelerometers. Both the Virtual Raft Project and EcoRaft Project are written primarily in Java with a module written in C++ to interact with the IrDA sensors. The graphics are displayed using JOGL, the Java version of OpenGL graphics library. A number of factors combined to create an engaging moment of transition for the EMAs – animation, graphical effects, sound, autonomous behavior, networking, interface elements, and interaction context. Following are examples from both of these projects to demonstrate various ways to design engaging EMAs.

Animation is an important element of EMAs, as it represents the embodiment of the agent. It is also one of the distinct features and advantages of EMAs comparing to ordinary mobile agents. One of the challenges is that the animation must look similar on both sides of the transfer, so the users will believe that the character that appears on the source and the destination devices are actually the same character. To contribute to this believability and engagement, the EMAs on different sides of the transfer should also be consistent. The EMA before and after transfer should appear the same and the EMA should appear from the expected direction. In the Virtual Raft Project, the jump animation of a humanoid in the source computer and the landing animation in the destination computer are technically two separate animations and are played in two different machines, but the timing is precise so as to simulate a fluid and seamless transfer between the two devices.

Graphical effects can also augment the animations to make the experience more engaging and life-like. For example, the Virtual Raft includes water ripples near a moving raft, and the hummingbirds in the EcoRaft have iridescent feathers on their breast. However, the level of details of the graphical effects needs to be balanced to maintain the frame rate. For example, since the graphical capabilities of the tablet PCs did not match those of the desktop computers, the iridescence for the hummingbirds was disabled on the tablets to maintain an acceptable frame rate.

In addition, sound plays a significant role in enhancing the believability and engagement of EMAs. Sound serves as a transition indicator to convey a more convincing sense of continuity than animation alone is able to as agents move between platforms. For instance, in the EcoRaft Project when a seed rolls from one device to another, it emits a sound like a slide whistle. This slide whistle sound played an ascending tone from the device of origin followed by a descending tone on the destination device. Similarly, hummingbirds emit a distinctive set of sound effects when flying between devices. Using complementary pairs of sounds emitted from the two devices in sequence supports the illusion that a single agent leaves one platform and then appears on the other. They also provide confirmation for the user that a transfer is occurring. Furthermore, each type of agent has its own signature sound. For example, the each species of hummingbird sings a different tone, with the tones of the four species combining to form a chord. The signature sounds attached to each entity texture the soundscape much as the instrumental signatures of various characters do in Prokofiev's "Peter and the Wolf." Using sound in these various ways strengthens the connection between the embodied agent and its surroundings by giving an agent multiple avenues through which to contribute to its acoustic environment.

People often feel more engaged when they see directly visual feedback [14]. The EMAs in our implementation incorporate certain behaviors in response to human activities. Each stationary computer, which represents a virtual island, has a webcam mounted on top of the display. It allows the EMAs in the virtual world to sense the activities of users in the real world. When the webcam detects human motion, the humanoids in the Virtual Raft get up and approach the virtual camera, and the hummingbirds in the EcoRaft will have a higher probability of flying near the virtual camera. This response from the EMAs gives the impression that they have an awareness of the physical world. Also, the tablet PCs have built-in accelerometers, which can detect the tablet's motion in two axes. Our implementation allows the EMAs to respond to this data, thus reacting to the physical orientation of the tablet. In the Virtual Raft, humanoid EMAs try to balance on the raft; in the EcoRaft, hummingbirds fly to highest point of the tablet, and seeds will roll inside the box. The added interactivity helps increase the engagement for the users interacting with the exhibit. Furthermore, these examples show ways in which EMAs are able to respond to their physical environment, blurring the line between the physical and virtual worlds.

In the EcoRaft installation, many people felt that they needed guidance to learn the ecological principles in the installation. Since this installation was designed as a science museum exhibit and it is not feasible to have a museum staff member available to demonstrate to the visitors at all times, our research team added a text overlay on the screen showing the status of the ecosystem. An accompanying voiceover also plays when the text overlay changes. This feedback loop allows visitors to try the installation and learn how to interact without needing someone to explain the system to them. It also provides an additional channel for educational content.

The transfer of EMAs from one computer to another computer also involves networking. With other mobile agents, the transfer of an agent involves packaging up the agent's behavior and memory, sending the agent data to the destination computer, and then reassembling the agent on the destination computer when it receives the data. Only one agent exists during the entirety of the transfer. The sequence of transferring an EMA is similar to the transfer of a mobile agent, involving packaging up data, sending to the destination, and reassembling the agent. In the case of an EMA, the transfer must occur during the animation so that the agent appears to transfer seamlessly; otherwise the time that it takes to transfer the agent would interrupt the animation and break the fluidity. The user will see the beginning of a transfer animation on the source device while the agent's data is being sent across. When the destination computer receives the agent data, it starts reassembling the agent and starts playing end of the transfer animation. At this point, two animated agents representing the same character exist in two different computers at the same time, but only one is visible to the user. When the whole animation sequence is complete, the agent in the source computer will be removed from the virtual world and the agent in the destination computer will continue to execute. Timing is very important in this process. If the timing is off, users may see both agents at the same time, or they may need to wait for a long time before the agent appears on the destination computer, thus decreasing the believability.

The tablet PCs have built-in infrared (IrDA) sensors and the stationary computers are connected to external IrDA dongles. The devices use the IrDA sensors to detect proximity and orientation of the other devices and TCP/IP to exchange animated characters. In this implementation, IrDA is used to detect proximity, as it is directional and the range is small. IrDA sensors require line of sight to detect each other and the detection angle is about thirty degrees, which make them good candidates to detect proximity. The sensing range of other technology, such as Bluetooth or GPS, is too broad and not precise enough for this purpose. This ensures that islands will only receive signals from tablet PCs that are in close proximity, in front of the island, and oriented correctly. When the infrared device detects that another device is in range, the two computers establish a network connection (TCP/IP) through Wireless Ethernet (802.11) and use this network connection for data transfer. TCP/IP was chosen for the communication protocol because it is faster than IrDA, especially for transferring all the attributes of the EMAs. If we had used IrDA to transfer the attributes of the EMA, it might have introduced a visible pause during the data transfer and interrupted the flow of the animation. While this method is somewhat more complex, it serves to minimize transfer delays, since wireless Ethernet allows for much faster transfer than IrDA.

5. EVALUATION

In all, approximately 3000 people have interacted with the two EMA implementations. At each exhibition of the projects, the research team observed people interacting with the system in various capacities, ranging from listening to an explanation of the system, to actually using the tablets, to watching people use the tablets, to explaining and demonstrating the system to others. During these exhibitions, groups of participants were successful in achieving basic tasks with EMAs such as distributing them evenly among specified devices with minimal introduction and in short periods of time.

In order to evaluate EMAs' capability to engage participants, the research team analyzed participants' interactions with one of the EMA-based installations described above, the EcoRaft project. A combination of qualitative techniques were employed, including non-interactive observation, participant observation, and semi-structured open-ended interviews [16]. Also, some of Druin's techniques for working with children [8] were employed. Non-interactive observation and participant observation were employed with nearly all of the 3000 participants. Interviews were done with 40 participants – 35 of them interviewed both before and after the interaction, and 5 interviewed only afterward. These interviews were conducted by one interviewer and one notetaker. The exact methodology differed slightly at the three deployment sites for these projects due to the exigencies of the various environments.

The first deployment site was in the Emerging Technologies program at the 2005 ACM SIGGRAPH Conference in Los Angeles, California. The researchers engaged in non-interactive observation; in addition, at SIGGRAPH the researchers conducted semi-structured interviews of one to two participants at a time. During set blocks of time, random participants were approached and asked if they would like to participate in an evaluation of the installation. The process began with a preliminary interview, which was done before participants had a chance to interact with the system. Participants were asked about their initial impressions regarding the EcoRaft exhibit, as well as their demographic information and technological background. A research member then gave a demonstration of the exhibit to the subject while the interviewer and notetaker observed the participants' interaction with the system. After finishing the interaction, the participants returned for a follow-up interview in which they were asked about their general impressions, their understanding of the ecological concept behind the game, the ease or difficulty in comprehending and using the system, their perspective on the realism of the exhibit, and whether or not the interface was natural and intuitive.

The second deployment site was at the Discovery Science Center, a hands-on science museum in Santa Ana, CA. The participants consisted of children, ages 7 to 13, visiting with family members or with day camps and summer schools. In addition to making non-interactive observations, the researchers asked some of the children to describe their experiences in an interview and give their impressions about the system. Interviews were conducted with children randomly selected from those who interacted with the system for at least three to five minutes. Those children who agreed to participate in the interview were brought, with their parent or guardian accompanying, to a quiet room separate from the exhibit space for the interview. Interview questions were similar to those asked of other participants, as well as questions about whether they worked with other children and what they had learned.

The third setting was in a research lab on the campus of UC Irvine. The subject group consisted of local home-schooled children ages 5 through 9, all of whom were familiar with the Virtual Raft Project. The participants were split into focus groups of 4 to 5 children accompanied by parents or guardians. Before interacting with the system, participants were interviewed to determine their familiarity with rainforest ecology and restoration, as well as general familiarity with technology and what they remembered from the previous interaction. The focus groups were then combined and allowed to interact with the system together as one large group. After the interaction, follow-up interviews were conducted with the same small focus groups. The questions in these interviews were similar to the questions asked of the children at Discovery Science Center.

Researchers' observations and participants' responses indicated that participants of varying ages and levels of technical experience were able to quickly learn how to interact with the system. When asked about the interaction with the EMAs, participants indicated that "it seemed fairly intuitive" and that the interaction "was really easy, it's very accessible, that's definitely a very big plus." One of the most engaging aspects was what happened while the object being transferred was located on the participant's tablet. As described in the implementation section, accelerometers in the tablet PCs allowed virtual objects contained therein to move around as if affected by gravity, transforming these virtual objects into "quasi-physical objects," as one participant put it. According to another, "the ability to actually roll the seed around makes it an actual real object." This "quasi-physicality" made the EMAs seem so much a part of the physical world that many participants, especially the children, tried to interact with them by touching the screens of the tablet or the stationary computers. Also important was the seamless transfer between devices. "I am really impressed with how well images transfer from the tablet right to the screen," said one participant. "Everything seems to jump out." The coordination of animation, sound, networking, and other aspects of the implementation created a physicality that served to increase the engagement level for the large majority of participants.

However, the evaluations demonstrate that not every aspect of the EMAs was as successful at engaging participants. In the EcoRaft, very few participants noticed the effect of the webcams. When asked about them, one participant knew that they had "something to do with the recognition of something [in front of the screen]… it senses that basically there's something trying to interact." However, this individual thought that it was related to sensing the presence of the tablet, not the motion of participants, and had no idea that it affected the behavior of the hummingbirds. With the Virtual Raft, though, far more people noticed the change is the humanoids' behavior when a participant approached; many thought the humanoids were coming over to say hello when participants approached the screen. It could be that people are more willing to attribute that sort of intelligent behavior to humanoids rather than hummingbirds, but it is more likely due to the fact that the change in the humanoids' behavior is very sudden and obvious in comparison with that of the hummingbirds. The important point here is that people's effects on the behavior of EMAs should be obvious, such as that of the accelerometers, and less subtle, such as that of the webcam in EcoRaft.

The evaluations suggest that the EcoRaft is an effective tool in teaching about the process of restoration ecology. One of the core concepts behind the installation is that one cannot simply plant a bunch of seeds and expect them to grow, but rather specific species have to be brought to deforested areas in a specific order [4]. This concept is key to the process of restoration ecology. Many participants, when asked about the concept behind the installation, responded with comments such as "[the process] needs to be [done] in a specific order, that you need to put the tree first before you take the flower over there," or "you can't just put a bird there and expect it to stay without having a plant and the plant needs to be fed by the trees."

These evaluations demonstrate that EMAs are an effective approach to designing engaging autonomous agents. They point to some of the aspects of the design that act to increase that engagement, chiefly animation across heterogeneous devices and ways for the agents to react to the state of the surrounding physical world, as well as some other aspects that do not help as much to increase engagement. Furthermore, these evaluations show that EMAs can be and have been deployed as an effective educational tool, harnessing the increased levels of engagement they provide to teach concepts from restoration ecology.

6. DISCUSSION

This research anticipates the continued expansion of mobile devices into many aspects of life around the world. Many of these devices will have a variety of input and networking technologies such as high-bandwidth connections, global positioning, machine vision, and voice recognition. This is already the case in many places in the world. Software agents will operate on many of these machines, performing complex tasks both autonomously and at the explicit request of users. These agents will operate readily across devices, taking action fluidly on heterogeneous computational systems. The research described in this document examines the process by which these agents may be embodied on heterogeneous networks of devices. As the various aspects of embodied mobile agents begin to be addressed, this area may begin to develop as a platform for educational technology, new forms of computer gaming, and other work-related and social multi-device interactions.

EMAs could have applications in a number of different areas. As has been described above, EMAs could be applied to the area of ecology education by creating a participatory simulation in which several computers serve as virtual habitat niches populated by several EMA species, and mobile devices let children move EMAs from niche to niche. The EcoRaft Project seeks to serve this purpose. In addition, EMAs could be useful in teaching a wide range of complex topics, from economics (the flow of trade) to physiology (the transport of nutrients in the body) to chemistry (the reactions of chemicals as they mix) to global health issues (the transmission of disease). The project seeks to design EMA technologies that are sufficiently general to be applicable in these and other domains, as well as to provide a specific example of the application of EMAs to education through the EcoRaft Project.

An industrial application could involve a system to support a group of people collaborating on a project, where the flows of information are embodied as animated information objects moving from device to device. For a person on the East Coast, EMA-based information objects coming in from the West Coast could appear on the west side of his or her PDA's screen; similarly EMAs coming from upstairs could drop from the top of a desktop monitor. Size, speed, appearance and behavioral characteristics could provide immediate feedback about the content of the information object, and interactions among EMAs could convey the relationships among them. This kind of contextual information could help each person to understand both the geographical and the conceptual dynamics of the project.

An entertainment application could involve a game in which several humanoid EMAs are jumping rapidly from device to device, and the goal is to catch certain EMAs, while passing off others to one's competitors. This game would be a kind of mixed reality "hot potato" game, incorporating aspects of both virtual world and real world engagement.

These potential applications point to the relevance of EMAs to several domains that currently exist or are not far into the future. In the longer term, EMA research could also help to create entirely new forms of human-machine interactions, where autonomous systems play a more significant role in computation. Currently, many people are skeptical of computational systems with too much autonomy, and there is a focus on interacting "through" computers rather than interacting "with" computers [24]. As computational systems grow more powerful, it may be relevant for paradigms to develop through which people can utilize this potential. EMAs may be of some use in designing systems where people interact effectively with multiple computing devices and autonomous computational entities.

7. FUTURE WORK

With regard to the EcoRaft Project, the research team is working on a number of topics that will enable the permanent installation of the exhibit at Discovery Science Center. Specifically, the tablet PCs need to be mounted on rolling platforms that will prevent them from being dropped by participants and will also store longer-lasting batteries. As an alternative, the research is also being done on implementing a version for smaller, more durable mobile devices such as PDA's, as discussed below. Also, the flat panel displays need to be covered with Plexiglas to protect them from interactions where participants assume that they must be touch screens, and poke them with their fingers or bump them with the tablet PCs. There are numerous robustness challenges beyond the software itself in an installation designed for 250,000 annual visitors.

In addition to designing the first museum exhibit, the team is beginning the development of a network of regionally specific ecological exhibits for science centers and museums around the world. Each exhibit will feature a collection of animal and plant species that are tailored to a specific regional ecosystem, to help visitors learn about the ecological issues of their region. For example, a Minnesota exhibit might focus on the reintroduction of wolves, while other exhibits might address the spread of invasive species such as snakehead fish or American bamboo. These exhibits will be built on the same basic Virtual Raft platform.

The authors are also currently working on the integration of Palm devices into the system, and anticipate a number of opportunities and challenges in this regard. The differences in input, output and networking capabilities will require additional effort to incorporate these devices. For example, the smaller size of the device will entail a variety of modifications to the graphical representations of the EMAs. In addition, the limited memory and processing capacity will also entail the implementation of new subsystems to enable EMAs to scale effectively in both memory footprint and processing needs.

With regard to developing EMAs for broader application, security becomes a significant concern. An autonomous agent capable of moving between devices could easily be seen as a kind of computer virus. There are at least two main potential security problems. First, an agent from one's own device could migrate to someone else's device, carrying with it potentially sensitive information. Second, an agent from an unknown or undesirable source could migrate to one's own device, carrying with it a harmful payload or at least being a new kind of spam. Therefore, having a substantial mechanism for preventing unwanted mobility of agents will be an important part of this project.

In addition to these topics, there are significant areas of future work in the integration of EMAs with existing software systems. For example, how can existing mobile agents be given an animated embodiment so that people can perceive their actions more clearly, and how could existing embodied agents be enhanced by allowing them to cross smoothly between devices? These areas of future work could be relevant to a great number of current agent-based systems.

8. CONCLUSION

This paper has proposed that embodied mobile agents (EMAs) can help to provide a coherent user experience across networks of multiple computational devices. Two interactive projects implemented with EMAs have been demonstrated at several major conferences (ACM CHI, ACM SIGGRAPH, CSCL), a public science center, and a number of other venues, to a total audience of approximately 3000 participants. Evaluations of the EMA system based on these participants' experiences were presented, demonstrating some of the advantages and challenges of EMA-based systems. While EMAs are not a solution to all multi-device interactions that people might wish to have, they could be an effective tool in the effort to create seamless interactions across the heterogeneous networks of devices that people increasingly use for learning, work, and play.

9. ACKNOWLEDGMENTS

The authors would like to acknowledge the support of the Donald Bren School of Information and Computer Science, The California Institute for Telecommunications and Information Technology (Calit2), the Nicholas Foundation, and the Emulex Corporation for their continuing support of this project. The authors also thank several anonymous reviewers for their contributions to the paper.

10. REFERENCES

[1] Blumberg, B., Downie, M., Ivanov, Y., Berlin, M., Johnson, M.P. and Tomlinson, B., Integrated learning for interactive synthetic characters. in Proceedings of the 29th annual conference on Computer graphics and interactive techniques, (San Antonio, Texas, USA, 2002), ACM Press, 417--426.

[2] Borovoy, R., Martin, F., Vemuri, S., Resnick, M., Silverman, B. and Hancock, C., Meme tags and community mirrors: moving from conferences to collaboration. in Proceedings of the 1998 ACM conference on Computer supported cooperative work, (1998), ACM Press, 159--168.

[3] Brooks, R.A., Human Level Cognition in Embodied Robots. in Proceedings of the International Joint Conference on Neural Networks (IJCNN '93), (Nagoya, Japan, 1993), 1079–1084.

[4] Carpenter, F.L., Nichols, J.D., Pratt, R.T. and Young, K.C. Methods of facilitating reforestation of tropical degraded land with the native timber tree, Terminalia amazonia. Forest Ecology and Management, 202. 2004. 281-291.

[5] Cassell, J., Bickmore, T., Billinghurst, M., Campbell, L., Chang, K., Vilhjálmsson, H. and Yan, H., Embodiment in Conversational Interfaces: Rea. in Proceedings of the CHI'99 Conference on Human Factors in Computing Systems, (Pittsburgh, PA, USA, 1999), 520-527.

[6] Cheok, A.D., Yang, X., Ying, Z.Z., Billinghurst, M. and Kato, H. Touch-Space: Mixed Reality Game Space Based on Ubiquitous, Tangible, and Social Computing. Personal Ubiquitous Comput., 6 (5-6). 2002. 430--442.

[7] Dourish, P. Seeking a Foundation for Context-Aware Computing. Human-Computer Interaction, 16 (2, 3 & 4). 2001. 229-241.

[8] Druin, A. Cooperative inquiry: developing new technologies for children with children. Conference on Human Factors in Computing Systems. 1999. 592-599.

[9] Fabre, Y. A framework for mobile-agents embodied in X3D networked virtual environment Proceeding of the eighth international conference on 3D Web technology, ACM Press, Saint Malo, France, 2003.

[10] Ferreira, A.G., Cerqueira, R.F.G., Celes, W. and Gattass, M. Multiple Display Viewing Architecture for Virtual Environments over Heterogeneous Networks ACM SIGGRAPH, ACM Press, 1999, 1-10.

[11] Gratch, J., Rickel, J., André, E., Badler, N., Cassell, J. and Petajan, E. Creating Interactive Virtual Humans: Some Assembly Required IEEE Intelligent Systems, 2002, 54-63.

[12] Gray, R., Kotz, D., Nog, S., Rus, D. and Cybenko, G., Mobile Agents: The Next Generation in Distributed Computing. in Proceedings of the Second Aizu International Symposium on Parallel Algorithms/Architectures Synthesis, (Fukushima, Japan, 1997), IEEE Computer Society Press, 8-24.

[13] Hutchings, D.R., Czerwinski, M., Meyers, B. and Stasko, J. Exploring the Use and Affordances of Multiple Display Environments. Workshop on Ubiquitous Display Environments at UbiComp. 2004.

[14] Hutchins, E.L., Hollan, J.D. and Norman, D.A. Direct Manipulation Interfaces. Human-Computer Interaction, 1 (4). 1985. 311-338.

[15] Ishii, H. and Ullmer, B. Tangible bits: towards seamless interfaces between people, bits and atoms. in Proceedings of the SIGCHI conference on Human factors in computing systems, ACM Press, 1997, 234--241.

[16] Lofland, J. and Lofland, L. Analyzing Social Settings: a guide to qualitative observation and analysis. Wadsworth Publishing Company, Belmont, CA, 1984.

[17] Lumsden, J. and Brewster, S. A paradigm shift: alternative interaction techniques for use with mobile & wearable devices. in Proceedings of the 2003 conference of the Centre for Advanced Studies on Collaborative research, IBM Press, Toronto, Ontario, Canada, 2003, 197-210.

[18] McIntyre, A., Steels, L. and Kaplan, F., Net-mobile embodied agents. in Proceedings of Sony Research Forum, (1999).

[19] Meunier, J., A Virtual Machine for a Functional Mobile Agent Architecture Supporting Distributed Medical Information. in 11th IEEE Computer Based Medical Systems conference, (Stamford, CT, 1999).

[20] O'Hare, G.M.P. and Duffy, B.R., Agent Chameleons: Migration and Mutation within and between Real and Virtual Spaces. in The Society for the Study of Artificial Intelligence and The Simulation of Behavior (AISB 02), (London, England, 2002).

[21] Perlin, K. and Goldberg, A., Improv: A System for Scripting Interactive Actors in Virtual Worlds. in Proceedings of ACM SIGGRAPH 96, (New Orleans, LA, USA, 1996), 205-216.

[22] Peysakhov, M., Artz, D., Sultanik, E. and Regli, W. Network Awareness for Mobile Agents on Ad Hoc Networks. Autonomous Agents & Multi Agent Systems. 2004.

[23] Rekimoto, J. Pick-and-drop: a direct manipulation technique for multiple computer environments. in UIST '97: Proceedings of the 10th annual ACM symposium on User interface software and technology, ACM Press, 1997, 31--39.

[24] Roush, W. Social Machines. Technology Review, August. 2005.

[25] Sacher, H. and Loudon, G. Uncovering the new wireless interaction paradigm. interactions, 9 (1). 2002. 17-23.

[26] Stock, O. and Zancanaro, M. Intelligent Interactive Information Presentation for Cultural Tourism. International Workshop on Natural, Intelligent and Effective Interaction in Multimodal Dialogue Systems. 2002.

[27] Sumi, Y. and Mase, K., AgentSalon: Facilitating face-to-face knowledge exchange through conversations among personal agents. in International Conference on Autonomous Agents, (Canada, 2001), ACM Press, 393-400.

[28] To, H.H., Krishnaswamy, S. and Srinivasan, B. Mobile agents for network management: when and when not! Proceedings of the 2005 ACM symposium on Applied computing, ACM Press, Santa Fe, New Mexico, 2005.

[29] Tomlinson, B., A Heterogeneous Animated Platform for Educational Participatory Simulations. in Computer Supported Collaborative Learning (CSCL), (Taipei, Taiwan, 2005).

[30] Tomlinson, B., Yau, M.L., O'Connell, J., Williams, K. and Yamaoka, S., The Virtual Raft Project: A Mobile Interface for Interacting with Communities of Autonomous Characters. in Conference Abstracts and Applications, ACM Conference On Human Factors In Computing Systems (CHI 2005), (Portland, OR, 2005).