Multiple Virtual Rafts: A Multi-User Paradigm for Interacting with Communities of Autonomous Characters

Bill Tomlinson[1], Jesse Gray[2] and Man Lok Yau[1]

[1]Arts Computation Engineering program
University of California, Irvine
Irvine, CA 92697-3425
wmt@uci.edu, mlyau@uci.edu

[2]Robotic Life Group
MIT Media Lab
Cambridge, MA 02139
jg@media.mit.edu

ABSTRACT

This paper describes a multi-user version of the "Virtual Raft Project" being exhibited in the Interactivity Program at CHI 2005. The Virtual Raft Project is an interactive installation in which communities of autonomous animated characters inhabit desktop "virtual islands." A human participant may transport the characters between the islands via a mobile device-based "virtual raft." This paper describes an implementation of a multi-user version of this project, in which several virtual rafts may be used simultaneously to carry characters among the islands. The multi-user experience improves on the single-user original in four ways: increased throughput, increased collaboration among the participants, increased enjoyment for the participants, and the introduction of a new mode of interaction (characters jumping directly from one raft to another). This paper also provides a preliminary evaluation of the entire system through observations from a deployment of the Virtual Raft Project to approximately two hundred people.

Author Keywords

Tangible interfaces, intuitive interfaces, mobile devices, autonomous characters, interactive animation

ACM Classification Keywords

H5.2 Information interfaces and presentation (e.g., HCI): User Interfaces; I.3.7 Three Dimensional Graphics and Realism: Animation

INTRODUCTION

The Virtual Raft Project is an interactive installation that is being exhibited in the Interactivity program at CHI 2005 [7]. This installation features several computer screen-based "virtual islands" – virtual environments populated by small communities of autonomous animated characters. Each participant interacts with the system by means of a mobile device-based "virtual raft," which enables him or her to carry characters from one virtual island to another. When the participant brings the raft close enough to an island, a character can jump from the island to the raft (see Fig. 1). The human may then transport the character on the raft to another island, where it may disembark. The character on the tablet PC also responds to the motion of that device in real space; when the participant tilts the tablet, the character needs to try to keep its balance. In addition, a simple machine vision system on each island allowed the characters to respond to the movements of participants in front of that screen. The ability of the characters to move from screen to screen while taking into account the physical locations of these screens, as well as their ability to respond to the movements of the screen itself, all help endow the character with a greater perceived physical presence. This island/raft paradigm allows the virtual character to break free from the confines of a single computer screen.

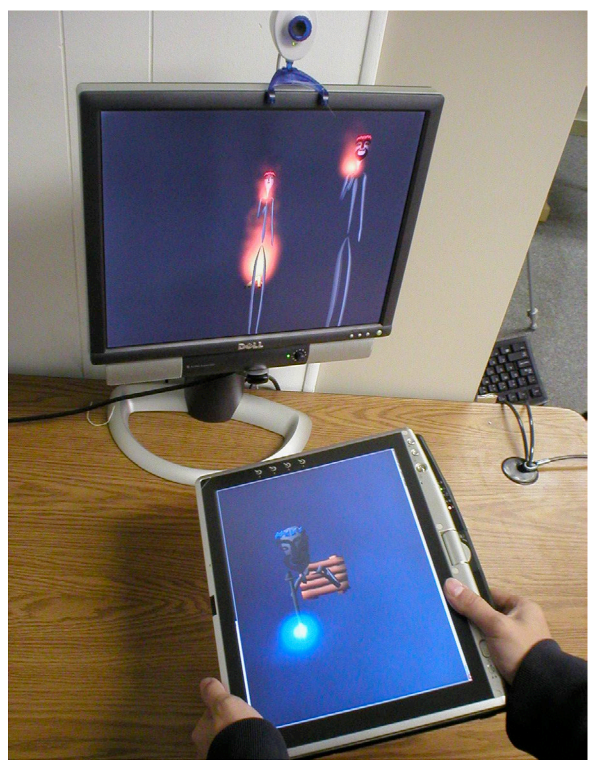

Fig. 1: The virtual raft paradigm lets characters to jump seamlessly between desktop screens and mobile devices.

CHI 2005, April 2–7, 2005, Portland, Oregon, USA.
ACM 1-59593-002-7/05/0004.

This paper presents a multi-user version of the Virtual Raft Project in which multiple rafts may carry characters at the same time. This implementation had four major impacts on the installation as a whole. First, it increased the throughput of the installation significantly, allowing as many as twenty people to be engaged simultaneously. Second, it increased the collaboration among the people interacting with the system, enabling them to work together to achieve common goals. Third, it appeared to increase the enjoyment that most people felt while interacting with the system. Finally, it made possible a new mode of interaction, in which characters could be allowed to jump from raft to raft, rather than between islands and rafts. These impacts were observed during a deployment to two hundred people in November 2004. These impacts suggest that the multi-user version of the system is a substantial advance over the single-user interaction, and a significant contribution to the field of human-computer interaction.

RELATED WORK

Many innovative research projects have been conducted involving multi-user interaction with mobile computing devices. Klopfer and Woodruff [5] experimented to see how teachers can use wearable and handheld computers to improve their teaching and to explore how the technology can help us understand the collaborative learning process. Flintham et al. [4] created the avant-garde games “Can You See Me Now?” and “Bystander” that allowed online participants to collaborate with mobile participants on city streets via sophisticated GPS systems. Borovoy et al. [1] implemented a system to facilitate interaction among the participants at a conference. Staffan et al. [3] implemented a game that uses the physical world as a game board.

Numerous researchers have explored the field of autonomous animated characters (e.g., [2, 6]). Computer games such as EA’s The Sims and toys such as Bandai’s Tamagotchi have also offered innovative work in the area of real time characters.

None of these systems have allowed the seamless transfer of animated characters among heterogeneous computational platforms, or the use of accelerometers to enable a tactile interaction among participants and virtual characters.

MULTI-USER INTERACTION

The virtual raft paradigm is being used both as a research platform for exploring interfaces and character believability, and also as an educational tool to encourage users to engage with specific content domains in an interactive setting. To this end, the system features a simple additive color theory game. Each of three islands begins with a different color fire – red, green and blue – and three characters that carry torches of that same color. Whenever a character goes to a new island, it brings its torch and contributes its torch color to the fire. The goal is to get all three colors on all three islands, producing three white fires. Tilting the raft too far to one side will cause a character to stumble on the raft and dip its torch in the water, quenching its fire and preventing it from delivering any color until it can re-light its torch at another fire. This system provides a narrative framework for the interactions that groups of people have with the system, and allows them to have a specific goal in mind while interacting with the virtual characters, islands and rafts.

IMPLEMENTATION

This system runs on standard desktop computers and tablet PCs. The desktop computers are the islands on which the characters gather, and the tablet PCs each display one raft on which a character can travel between islands. The desktop computers are equipped with a web-cam each to detect the proximity of humans to the screen and an IrDA dongle to detect rafts that are within range. The tablet PCs have IrDA built in, as well as 2-axis accelerometers that can translate the motion of the device into motion of the virtual raft, increasing the user's connection with the character.

The initial version of this project used a web-cam to detect the presence of the virtual raft, which was wrapped in bright yellow material to aid detection. Once the raft was detected by an island's web-cam, the island would transfer a character to or from the raft over wireless ethernet. This worked well in the single user version since the goal was to detect only the one raft; there was no need to distinguish between different rafts. With multiple rafts, however, it becomes important to identify which raft is which. To solve this problem, we switch to IrDA for proximity detection. Wireless ethernet is still used for the actual character transfer, although in the future the system may be modified to use IrDA for both detection and transfer.

A compliant IrDA device is required to implement the Link Management Protocol, which includes a Device Discovery service. This service provides the names of all the IrDA devices in range. Once a computer receives the name of a device in IrDA range, it can connect to it via TCP/IP over wireless Ethernet and negotiate character jumping.

The maximum range of IrDA operation is approximately one meter, which is ideal for our purposes. One meter is an acceptable distance for the characters to jump across, so we do not need to attempt any additional range estimation: if a device is discovered by the IrDA Device Discovery service then it is within jumping range. IrDA also requires that the devices be pointing to within 30° of one another. The virtual raft on the tablet PC was modeled to have a prow, encouraging users to hold the tablet PC in the correct orientation. The IrDA dongles for the desktop computers were attached unobtrusively beneath each monitor.

The switch to IrDA also has increased the robustness of the system; IrDA is a well established technology that does not need to be re-calibrated in new lighting conditions, as a vision based system might. IrDA is widely available on portable electronics, so the switch to IrDA may ease porting the raft software to other portable devices.

In the initial system, using the cameras mounted on the island computers for raft detection necessitated an asymmetric design for the communications protocol between the raft and the islands - only the islands could detect the raft. However, now that all devices use IrDA in the same way, the system has been modified so that the communications protocol is symmetric. Rafts and islands all speak the same language, so they can ask a device if there is space for another character to jump on (and answer that query if asked by those devices), as well as send characters to other devices (and receive them from those devices). Without any extra work, then, other interactions automatically function as a user would expect: for example a character can jump from their raft straight to an unoccupied raft (see Fig. 2).

Fig. 2: Interactors allow a character to jump from one raft to another.

DISCUSSION

The multi-raft version of the virtual raft project was exhibited in November 2004 to a group of two hundred people at the opening ceremony of a new building on a university campus. Although it was not possible to collect statistical data about the interactions in this venue, several impacts of the multi-user implementation became evident through observation.

First, the multi-user implementation enabled a significant increase in the throughput of the installation. Groups of twenty people at a time visited the installation space, and nearly all of these individuals appeared to be engaged with the experience. Visitors appeared to take on several different roles in the interaction. There were *primary interactors*, who carried the rafts around the space and moved characters from island to island. Each primary interactor was flanked by several *advisors*, who actively watched the interaction and offered advice to the primary interactor. There were several *connectors*, people not affiliated with a certain raft who kept tabs on the overall progress of the group and would offer suggestions about which islands needed which color characters. There were *islanders*, who remained near a given island and interacted with the characters on that island through the vision system. Finally, there were *spectators*, who enjoyed viewing the installation as a whole and watching the movements of other participants. People sometimes switched roles; for example, advisors were often interested in being primary interactors, but were waiting their turn with the interface.

A second impact enabled by the multi-user version of the system was an increase in the collaboration among participants. While interacting with the system, participants talked to each other to get the task done, they helped each other learn the interface and they worked together to get the right characters on their rafts. Having three rafts allowed a parallel exploration of the interface functionality. If one person was having trouble with the interface (which was surprisingly uncommon, but did occur on several occasions), a person who had seen someone else with a different raft would often offer advise about how to do it better. In fact, in several groups this feedback loop led to a certain kind of superstitious behavior that the raft needed to be tilted forward in order to help the characters jump. If there was any delay in detection, interactors would often tilt the tablet up slightly so that the character on the tablet could better "see" the island it was jumping to, indicating that they were buying in to the lifelike nature of the characters rather than trying to look for the underlying mechanisms. This behavior would then spread among other interactors in the group. This spread was made possible because there were three rafts operating at the same time.

A third impact of the multi-user implementation was an increase in the enjoyment that people appeared to feel while interacting. There was a substantial amount of social interaction among the groups of participants, many of whom did not know each other before the tour. They talked together, laughed and appeared to derive pleasure from helping each other and working together to achieve a common goal. The virtual raft experience appears to get better when several people are interacting with it at once.

Finally, the multi-user implementation enabled an additional mode of interaction that wasn't possible in the single raft version. This interaction involved allowing characters to jump from raft to raft. While this interaction was not a necessary part of accomplishing the group's common goal, more than a few participants were interested to see if characters could jump between rafts, and were very satisfied to see that this interaction was, in fact, possible.

The virtual raft interaction appears to be intuitive for most participants. Even with no demonstration of how the proximity detection or jumping functioned, people held the rafts up to the LCD screens of the islands and succeeded in getting characters to jump on the raft on their first try. Most of the participants in this deployment interacted with the system in nine groups of about twenty at a time, as well as a number of smaller groups. Each large group was asked to complete the simple color-blending game that required at least six specific character transfers (so that one character of each color made it to each of the three islands). All nine groups succeeded in this task within approximately five minutes, thereby demonstrating that the raft paradigm was usable by novice participants in a very short period of time. After accomplishing the goal, the group seemed satisfied with the interaction, several times breaking into spontaneous applause at their success.

FUTURE WORK

The future work for this project involves fixing problems that became apparent during the public deployment, and also moving forward with several research topics.

The greatest difficulty that people had that prevented successful operation of the system was holding the raft in the wrong orientation. This occurred infrequently, but it indicates that the affordance provided by the interface could be improved to better indicate directionality. The prow that was added to the graphical raft on the tablet solved this problem to a large extent, but other indications of the directionality of the raft (e.g., hand-grips on the sides) could offer further improvements.

Extensions to the project will include making a system in which the multi-user element is crucial to the group's success, for example where two or more people have to take action simultaneously for the goal to be achieved. Also, the project has research under way to increase the social intelligence of the autonomous characters in the system, so that they are able to learn about their experiences on the rafts and tell each other about them. Finally, the production team has a collaboration under way with a regional science center to produce an ecological version of the project, in which several animated species will inhabit the islands, and visitors to the center will serve as vectors for the dispersal and migration of these species.

CONCLUSION

This paper has presented a novel multi-user interactive project in which participants help use mobile device-based virtual rafts to help autonomous characters move among computer screen-based virtual islands. This multi-user implementation improves upon a previous single-user version by increasing the throughput of the installation, enabling greater cooperation among participants, augmenting the enjoyment that people feel when interacting with the system, and introducing a novel mode of interaction in which characters may jump from one mobile device to another. These four advances became evident during a deployment of the project to approximately two hundred participants.

The Virtual Raft installation is based on a novel metaphor between the real space/virtual space distinction and the land/water distinction. In this "Island metaphor," graphical characters inhabit virtual space, which is like land to them, but cannot travel through real space, which is represented as water, without human help. The engagement that people appear to have felt with this metaphor points to an interesting additional result of this project. As computer programs become more autonomous, the desktop metaphor becomes a bit less appropriate. In the real world, files, folders and rubber bands reside on desktops; autonomous entities generally do not. Thinking about a computer screen as an island of virtual space could help people conceive of more effective ways to interact with the systems running there. While the island metaphor isn't likely to supplant the traditional computer desktop, it might serve as a useful way to think about future computational systems and the way people will interact with them.

ACKNOWLEDGMENTS

The authors would like to thank Eric Baumer, Sara Goetz, Jessica O'Connell, Ksatria Williams, So Yamaoka, Thomas Alspaugh and Ben Pillet for their contributions to the project. This project was supported by the ACE (Arts Computation Engineering) program, the Donald Bren School of Information & Computer Sciences and the Claire Trevor School of the Arts at UCI, by the UC Game Culture & Technology Lab, and by the California Institute of Telecommunications & Information Technology.

REFERENCES

1. R. Borovoy, F. Martin, S. Vemuri, M. Resnick, B. Silverman, and C. Hancock, "Meme tags and community mirrors: moving from conferences to collaboration," in *1998 ACM conference on Computer supported cooperative work*, 1998.
2. R. Burke, D. Isla, M. Downie, Y. Ivanov, and B. Blumberg, "CreatureSmarts: The Art and Architecture of a Virtual Brain," in *Game Developers Conference*, San Jose, CA, USA, 2001.
3. J. Falk, P. Ljungstrand, S. Bjork, and R. Hansson, "Pirates: proximity-triggered interaction in a multi-player game," in *CHI '01 extended abstracts on Human factors in computing systems*: ACM Press, 2001, pp. 119--120.
4. M. Flintham, S. Benford, R. Anastasi, T. Hemmings, A. Crabtree, C. Greenhalgh, N. Tandavanitj, M. Adams, and J. Row-Farr, "Where on-line meets on the streets: experiences with mobile mixed reality games," in *Conference on Human factors in computing systems (CHI 03)*: ACM Press, 2003, pp. 569--576.
5. E. Klopfer and E. Woodruff, "Platforms for Participatory Simulations - Exploring Systems and Generating Discourse with Wearable and Handheld Computers," in *Conference on Computer Supported Collaborative Learning*, Oslo, 2003.
6. K. Perlin and A. Goldberg, "Improv: A System for Scripting Interactive Actors in Virtual Worlds," in *ACM SIGGRAPH 96*, New Orleans, LA, 1996.
7. B. Tomlinson, M. L. Yau, J. O'Connell, K. Williams, and S. Yamaoka, "The Virtual Raft Project: A Mobile Interface for Interacting with Communities of Autonomous Characters," in *CHI 05 Conference Abstracts and Applications*, Portland, OR, 2005.