Bill Tomlinson*
Man Lok Yau
Eric Baumer
Joel Ross
Andrew Correa
Gang Ji
Department of Informatics
5068 Bren Hall
University of California, Irvine
Irvine, CA 92697-3440

Richly Connected Systems and Multi-Device Worlds

Abstract

Many human activities now take place in settings that include several computational devices—such as desktop computers, laptops, and mobile phones—in the same physical space. However, we lack interaction paradigms that support a coherent experience across these collocated technologies and enable them to work effectively as systems. This article presents a conceptual framework for building richly connected systems of collocated devices, and offers two implemented examples of interactive virtual worlds built on this framework. Aspects of this framework include multiple channels of real and apparent connectivity among devices: for example, multiple kinds of data networking, cross-device graphics and sound, and embodied mobile agents that inhabit the multi-device system. In addition, integration of the system with the physical world helps bridge the gap between devices. We evaluate the framework in terms of the types of user experiences afforded and enabled by the implemented systems. We also present a number of lessons learned from this evaluation regarding how to develop richly connected systems using heterogeneous devices, as well as the expectations that users bring to this kind of system. The core contribution of this paper is a novel framework for collocated multi-device systems; by presenting this framework, this paper lays the groundwork for a wide range of potential applications.

1 Introduction

Over the past several decades, computational devices have spread rapidly among many human societies. Because of this extensive growth in usage, devices are frequently located in physical proximity to each other. However, while devices often have the capability to network with each other and with the internet, these collocated devices rarely take full advantage of their physical proximity to each other to help them interact or to facilitate their interactions with people. Rather, they tend to utilize patterns of interaction that were developed before computing devices were so abundant. These single-device interaction paradigms, which work well when humans engage separately with each device, rarely facilitate a coherent experience across several devices. As a result, human users miss out on the significant potential of collocated devices, and may feel overwhelmed rather than supported by the multiplicity of devices present. In order for these devices to integrate smoothly and for people to un-

Presence, *Vol. 18, No. 1, February 2009, 54–71*

---

*Correspondence to wmt@uci.edu.

derstand their collective operation, new interaction paradigms, and new metaphors that guide people's understanding of those paradigms, are needed.

Large collocated groups of people have a wide range of remarkable capabilities—as companies, as orchestras, as armies, as sports teams, as social organizations, and as universities. While computational devices will not form groupings of such complexity and power, they may be able to work together to enable a similar kind of synergy. Just as people can work closely together when physically near each other, collocated devices should be able to cooperate on the task of producing a coherent user experience.

This kind of cooperation requires complex coordination among multiple devices. However, groups of devices are usually interconnected by means of only a single communication channel, such as Ethernet, WiFi, or USB. This single method of connectivity misses a significant opportunity in linking devices into a single system. We propose instead to connect groups of devices via multiple channels—not just through networking, but also through graphics, sound, and other media. Since many of these channels use modalities that are also used by humans (e.g., graphics, sound), this rich connectivity among devices also leads to rich engagement with the people interacting with these systems. This paper proposes a conceptual framework that can help people think about and build multi-device systems. In this framework, the devices are connected to each other using multiple media, some of which may match human sensory capabilities. Our goal in developing this framework is to explore an interaction paradigm suited to collocated, multi-device, interactive systems. This type of interaction should include not only multiple devices, but also the potential for multiple participants as well. By enabling the creation of this form of multi-device HCI, we hope to lay the groundwork for a wide range of other forms of multi-device interactions and thus allow people to benefit from having groups of devices in the same physical space.

To describe this framework for richly connected systems, we begin with an assessment of existing research and related systems. Thereafter we present our framework for richly connected systems, detailing the core components of this framework: collocated data communication, cross-device graphics, multi-device sound, embodied mobile agents that inhabit the multi-device system, and integration of the system with the physical world. Next we introduce two implemented prototypes of interactive exhibits that were built while developing this framework, and describe how these systems apply and demonstrate the framework. We then use these prototypes to evaluate the effectiveness of the framework as a whole, from the perspective of both end-users and developers. Finally, we present a number of lessons learned from these evaluations, and conclude with possible future directions for this framework. Throughout this paper, we seek to demonstrate the viability of collocated multi-device systems in general. These systems are underutilized to date, and we believe they have significant potential for improving the ways in which people interact with the devices that surround them by improving the interaction across and between nearby devices.

2 Related Work

Currently most dominant interaction paradigms in HCI revolve around interacting with a single device at a time (Baecker, Grudin, Buxton, & Greenberg, 1995; Card, Moran, & Newell, 1983; Norman, 1988). Even when people are in groups, the emphasis with mobile devices is often on interacting with a single device, for example as shown in Barkhuus et al. (2005). Although many CSCW systems also incorporate group activity, for example as shown in Grudin (1988), Grudin and Palen (1995), Heath, Luff, Lehn, Hindmarsh, and Cleverly (2002), and Hindmarsh, Heath, Lehn, and Cleverly (2005), the interaction paradigm still primarily revolves around users interacting with a single isolated device, application, or system. In this paper, we explore ways that the heterogeneity and complexity of multiple devices in the real world can be orchestrated to create unified interactive experiences.

This research draws on previous efforts to create an effective user experience for environments containing multiple devices. The Pick-and-Drop system, for example, enables a user to transfer files by picking up files from one computer and dropping them into another computer with a pen (Rekimoto, 1997). It hides the underlying technical details from the users, simplifying the procedure of transferring files and data between multiple devices. In developing the framework described here, we sought to preserve the simplicity and elegance of Pick-and-Drop while enabling interactions with more complex multi-device systems.

Our framework focuses on coordinating multiple collocated devices—devices that are in physical proximity. There are a number of previous systems that involved collocated devices. For example, the PEACH system—Personal Experience with Active Cultural Heritage (Stock & Zancanaro, 2002)—augments museum artifacts near the user with additional information, such as supplemental text and pictures, through the use of a PDA. Similarly, tangible interfaces (e.g., Ishii & Ullmer, 1997) also often involve multiple collocated computational devices. These systems leverage the proximity of collocated devices to improve the user experience.

There are also many different commercial products that involve close coordination of collocated devices. For example, mobile phones and Bluetooth headsets are designed to work closely together, as are video game consoles and wireless game controllers. The HotSync functionality that allows Palm devices to synchronize their data with a desktop computer is also an example of two devices coordinating their actions. However, most of these solutions are specifically tailored for connecting a small, fixed number of devices. The framework described in this paper seeks to coordinate systems that span large groups of heterogeneous devices.

Coordinating devices into richly connected systems requires data networking to connect the component devices. Data communication among collocated devices factors into a number of areas of research that have helped to inform the implementation of this framework. The Meme Tags project enabled the transportation and display of text fragments among small electronic badges (Borovoy et al., 1998), thereby creating a collaborative system using wearable devices. Brumitt et al. created technologies that support intelligent environments in which different devices communicate with each other to provide services for the users (Brumitt, Meyers, Krumm, Kern, & Shafer, 2000). For example, pressure sensors and cameras are used to track users in a room, enabling the working session of that person to be transferred to the computer that is next to him or her.

Graphical connections can play a significant role in creating connections over multiple channels. Previous research in computational cinematography and virtual lighting has greatly informed this project. For example, research in virtual cinematography (e.g., Drucker, He, Cohen, Wong, & Gupta, 2003; He, Cohen, & Salesin, 1996; Tomlinson, Blumberg, & Nain, 2000) has suggested various ways to control the movement of the camera in virtual worlds. In addition, past work in the growing area of lighting design in interactive environments (e.g., El-Nasr, Zupko, & Miron, 2005) has sought to create dynamic lighting design automatically in gaming environments to create a more engaging gaming experience.

Audio is another important aspect to creating richly connected systems and engaging, immersive experiences for users. For most people sound is a pervasive aspect of daily experience, and the unique psycho-acoustic aspect of our hearing helps us to internalize our surroundings through our aural experience of them (Ong, 1982). Creating a well-crafted soundscape for any virtual environment is an important part of engaging and connecting the user with the system. For example, Drettakis et al. use 3D sound to help evaluate urban planning in a virtual environment (Drettakis, Roussou, Reche, & Tsingos, 2007). Other multi-device systems have taken advantage of audio design to create a virtual soundscape, such as Audio Aura (Mynatt, Back, Want, Baer, & Ellis, 1998) and “A New Sense of Place?” (Williams, Jones, Fleuriot, & Wood, 2005). Previous work in using sound in conjunction with mobile devices has also informed our work. Nomadic Radio (Sawhney & Schmandt, 2000), for example, uses an audio interface to deliver notifications to a mobile user.

Our framework also focuses on creating connections between the physical environment and a virtual world. A great deal of research has been done on novel and engaging ways to blur the boundaries between physical and virtual space (e.g., Ishii, Mazalek, & Lee, 2001;

Khoo et al., 2006). Many of these approaches connect the manipulation of a physical object to the manipulation of some digital entity. For example, the phicons in metaDESK allow the user to physically manipulate the location of digitally displayed structures (Ishii & Ullmer, 1997), and the rattling of Live Wire provides a physical and audible clue as to the current state of digital network traffic (Weiser & Brown, 1996). Similarly, this work seeks to emphasize the ways in which physical and virtual worlds overlap, connect, and coexist. However, the goal here is not to enable users to manipulate digital data physically, but rather to allow for a deep set of connections between the physical world and the virtual world to better engage users as they act within a single, hybrid space.

The problem of creating connections between a virtual world and the physical world, and doing so in a fashion that is familiar to the average human being, has been an ongoing subject of research. "Mixed Reality" projects blend between a virtual world and the physical world. Human Pacman (Cheok et al., 2003) and Age Invaders (Khoo et al., 2006), for example, are played in both the physical and the virtual worlds. Similarly, MR MOUT (Mixed Reality for Military Operations in Urban Terrain) focuses on the construction of algorithms that allowed for color transferring and shadowing between physical and virtual entities (Hughes, Reinhard, Konttinen, & Pattanaik, 2004). "Virtual Light" uses a virtual flashlight to emit an image that represents the same image a user would experience with a physical light source (Naemura, Nitta, Mimura, & Harashima, 2002). It focuses on the correlation between virtual and physical entities, and uses a virtual image to emulate a physical object under lighting. Our work, like these systems, seeks to create human interactions across both physical and virtual space.

Our framework also uses embodied agents to help form connections between devices, as well as between the user and the system. For the purpose of this research, we use Maes' definition of agents: "computational systems that inhabit some complex, dynamic environment, sense and act autonomously in this environment, and by doing so realize a set of goals or tasks for which they are designed" (Maes, 1995). A great deal of work has been conducted on the creation of believable agents (e.g., Bates, Loyall, & Reilly, 1994; Blumberg & Galyean, 1995; Gratch et al., 2002; Perlin & Goldberg, 1996); this research incorporates work in embodied mobile agents (Tomlinson, Yau, & Baumer, 2006), which builds on these previous efforts.

There have been several systems that involved multi-device interactions with agents. The AgentSalon project is a system with desktop computers and mobile devices to facilitate face-to-face communication (Sumi & Mase, 2001). A large display is shared with multiple users, and each user has a mobile device, such as a PDA, which holds an animated agent. The agents can be transferred to the large display, where they engage in automated conversation. These automated conversations are intended to facilitate conversations between the users. Agent Chameleons is a system with agents that can transfer between robots and virtual environments (O'Hare & Duffy, 2002). The project explores the embodiment of agents in robotic platforms as well as in virtual environments. Each platform has different associated behaviors and capabilities; the agent understands the platforms that it can migrate to and evaluates which one is more appropriate for the current situation. The agent will then migrate to that platform and continue to function.

In applying this framework, we choose to use virtual worlds as a core interaction paradigm. Virtual worlds provide a number of benefits over other styles of interaction, such as windowing systems or text-based interfaces. They promote social engagement and activity (Barab, Thomas, Dodge, Carteaux, & Tuzun, 2005). They allow for an increased sense of embodiment and presence (Biocca, 1997). Finally, virtual worlds are able to augment and enhance understanding and education (Roussos et al., 1999). While virtual worlds are only one possible interaction paradigm for richly connected systems, they provide a useful starting point for exploration of this area.

All of these projects represent the type of interaction that would be enabled by the spread of collocated multi-device systems. However, the research described here is different from these prior works in several ways: it presents a broader framework for the development of

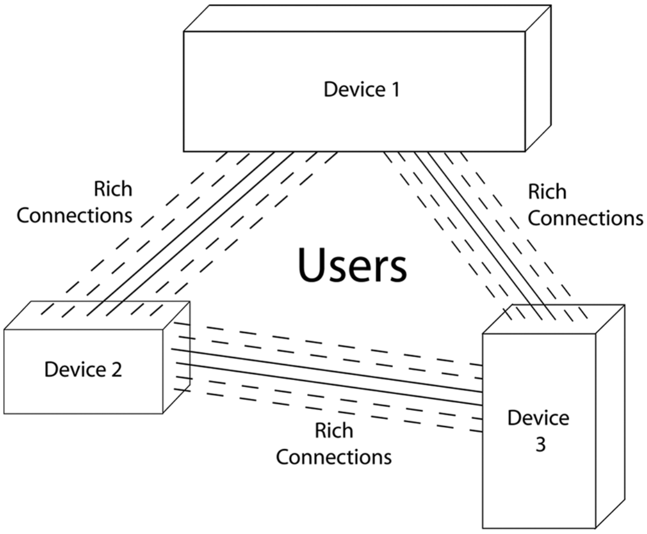

Figure 1. *The organization of a richly connected system. Each of several devices is connected to the others by actual channels such as data networking (solid lines) and apparent channels such as cross-device graphics (dashed lines).*

this type of system, it has a significant focus on the heterogeneity of the component devices, and it emphasizes using multiple modalities to form connections between devices, thus leading to a richer engagement with the people who use these systems.

3 Richly Connected Systems

The core idea presented by this article is that systems of devices may be connected via multiple channels, some of which may be perceivable by the people who interact with the system. This rich connectivity among devices, as well as among devices and people, leads to new forms of interactivity and new capabilities not possible with less coordinated systems.

Figure 1 depicts the main organization of a proposed framework for developing these richly connected systems. In this framework, heterogeneous devices are connected and coordinated through a wide variety of channels simultaneously: not only direct data network connections, but also apparent channels of connectivity such as graphics, sound, and other media. Furthermore, these channels are coordinated among themselves—media connections correspond to one another as if providing a single multimodal link between devices. The user, who is aware of and may interact with many of the communication channels, can perceive this unification. Thus for the user, devices are strongly coupled despite their differing capabilities and forms.

This framework differs from other multimodal interfaces in that the rich connections occur between the devices within the system, as well as between the user and the system. Although the user is aware of the multiple channels and thus may have multimodal interactions with the system, the coordination between these inter-device connections can allow for a more unified interaction experience. By creating linked communication channels for the user to interact with, we can enable novel forms of interaction with multiple heterogeneous devices.

In this article we explore a number of potential communication channels for richly connected systems. First we discuss the data connections, which allow for the passing of information as well as the coordination of other connection media. We then examine some user-perceivable modalities, specifically graphical and sound connections. Finally we describe the use of computational agents and integration with the physical world to create a unified interpretation for interaction with a system of richly connected devices.

3.1 Multi-Device Data Networking

The most typical way for groups of devices to be connected is in a digital network, with data packets exchanged through various wired or wireless technologies. Richly connected systems often involve this kind of data networking, and use it as a foundation for the other forms of inter-device connectivity. These networks allow for the devices to coordinate other channels of communication, such as by transferring system data or synchronizing user output.

To perform this data sharing, systems may use common protocols such as Ethernet, WiFi, Bluetooth, and/or Infrared (IrDA) to ensure rapid and reliable information transfer within the system. Richly connected systems may also use less traditional methods for coordinating data among devices. A system may transfer data using audible frequencies, or by using vision systems to determine the state of nearby devices. Each of these methods of data communication has a different set of characteristics that may make it well suited for a particular system, as well as provide coordination with other modalities. For example, IrDA requires line of sight, which links the physical modality to the data transfer, while vision systems might link data transfer to graphical displays. In either case, device coordination must rely on some method for communicating data to support other modes of communication between devices.

3.2 Multi-Device Graphics

Computer graphics are human-perceivable media that can be used as a communication channel in richly connected systems. For example, by having a single graphical entity (such as an icon or a virtual character) seem to move from one device to another during a data transfer, we increase the apparent coordination between the devices. Creating this illusion requires carefully coordinated timing, as well as graphical animations that support the idea of virtual characters who are actively mobile, rather than just passively transferred.

When two devices have different graphical capabilities, it may not be possible for both devices to represent a graphical entity in the same way; when the entity is transferred between such devices, it undergoes a stylistic transformation that may cause a visual discontinuity. One way to improve graphics and animation across heterogeneous devices involves separating the moment of cross-device transfer from the moment of stylistic transformation. The technique involves implementing an explicit stylistic change on the device with superior graphical capabilities, shortly before or after the cross-device transfer. Making the stylistic transformation explicit, and separating it temporally from the transfer, helps to preserve the identity of the graphical entity as it moves between devices, strengthening the apparent connection between the devices.

The use of virtual lighting may also need to be considered in creating the illusion of a unified graphical space across the multi-device system. For example, a problem that occurs in virtual worlds displayed on multiple devices is that graphical shadows cast by virtual entities may not be coordinated across the different devices. This problem is compounded in a richly connected system, where the graphical modality may be linked with the physical modality—in this situation, the shadows should change according to the physical motion of the device to coordinate the communication channels.

3.3 Multi-Device Sound

Multi-device systems may also be connected through the audio channel. Sound may be used in conjunction with visual displays to create video communication, or as its own user-perceivable modality. When dealing with multi-device systems, synchronizing the audio across different devices is often the key concern for coordination. This concern becomes even more pressing when the audio channel is linked with the graphical channel: although not as important an issue for background sound effects, the synchronization of graphical and audio events can be crucial for maintaining a coordinated connection. Coordinating audio with physical movement can also provide some unique challenges for sound designers, as the locations of specific speakers may not be known and fixed.

There may also be problems when creating audio for systems of heterogeneous devices. Just as devices may have different graphics cards or resolution displays, devices may also have different sound cards with distinct hardware capabilities, as well as disparate speaker setups ranging from monaural or stereo to 7.1 surround sound. Adapting audio output across devices with different capabilities can be done using different sets of sound samples, or may involve more complex audio manipulation.

3.4 Multi-Device Agents

A significant part of creating a coherent interactive system is making that system comprehensible to users. To provide an interpretation for the linked connections in a richly connected system, our framework suggests the use of autonomous and semi-autonomous computational agents that inhabit the world. Embodied mobile agents (EMAs) (Tomlinson, Yau, et al., 2006) should be able to operate on any of the heterogeneous devices in the multi-device system, and to transfer seamlessly between devices along the multimodal connections.

Graphical EMAs can take a range of forms, including animated characters (personified agents with the appearance of sociality), animated creatures (agents modeled after nonhuman species), and animated objects (based on inanimate objects in the real world). These agents provide a mechanism for coordinating interactions on the various devices. Seeing the same agent performing similar tasks across different devices can make a system's operation more transparent. Mobile agents, whether embodied in an animated form or not, can be an effective means of transferring data among devices, providing an interpretation for the interlinked connections in a richly connected system.

One of the challenges with EMAs in heterogeneous multi-device environments is enabling the agents to transfer believably across devices. There are a number of factors that contribute to this believability, including the agent's appearance on both sides of the transfer, the timing of the movement between devices, graphical effects that can support animations, and careful integration with other modalities, such as sound. When moving across devices that have significantly different graphical styles, these agents may be enhanced via heterogeneous animation techniques (Tomlinson, Yau, & Gray, 2005) that help smooth over differences (e.g., screen size, resolution) that might otherwise compromise the continuity of the animated transfer between devices.

Taken together, the multiple connections among the devices in these systems enable the creation of richly connected systems. While the potential connectivity among the devices in this kind of system is not limited to the channels described here, these media provide examples of the many different ways in which collocated devices may be connected.

3.5 Physical World Integration

Since the physical context of the multi-device system is an important factor in the way the system is used, significant integration between the real environment and the digital media can contribute to the viability of the system as a whole. The incorporation of various sensing technologies—GPS, light sensors, cameras, microphones, accelerometers, and so on—can help the physical reality surrounding the devices have an impact on the data within them, linking tangible or physical modality with other modes of communication. For example, graphics and sounds can react and interact based on inputs from these sensing components. Strengthening the connection between the elements of the multi-device system and the physical location in which that system is deployed can help the devices interact more effectively with each other, and help facilitate the human-computer interaction created by the system.

3.6 Examples of Richly Connected Systems

Richly connected systems can enable a wide variety of novel interactions by using multiple channels to connect devices. For example, one could envision a mobile photo sharing application using cell phones and a central stationary display, perhaps in the user's home, that provides rich connections using network, graphical connections, and physical world integration. Mobile users would be able to send photos to the central display using SMS, a type of networking connection. The central display would be oriented like a table, so that when the shared photo arrives, its graphical location on the central display would reflect the physical location from which the user sent the photo. That is, if the mobile device were on the east side of town, a photo sent from that device would appear on the east side of the central display, with the distance from the center corresponding to the mobile device's distance from the display. This richly connected system, incorporating networking, graphical connections, and physical world integration, would enable users to gain a new sense of their environment through the position of shared photos on the central display.

Some existing work can also been seen as examples of richly connected systems. For example, Human Pacman (Cheok et al., 2003) can be seen as a somewhat richly connected system, closely correlating networking, graphical connections, and physical integration. The graphical elements on the player's mixed reality display are based on the player's physical position and orientation. Interaction with various Bluetooth-enabled physical objects and virtual cookies provide further coupling between these three connections, enabling a novel and engaging experience through the richness of the connections.

The next section describes the use of rich connections to create another type of experience: collocated virtual worlds. Such worlds not only rely on many rich connections between devices, but they demonstrate how richly connected systems can be used to create engaging and novel forms of interaction with multi-device systems.

4 Collocated Virtual Worlds

Two prototype systems have been implemented using this framework. These systems are virtual worlds that exist among multiple collocated, richly connected devices. Virtual worlds implemented on a richly connected system enable a kind of interactive experience that resembles those found in augmented reality and multiplayer online games, but with greater potential for physical collaboration and coordination due to the collocated nature of the experience. The goal is to enable multiple participants to become immersed together in a multi-device virtual world. In particular, the virtual worlds described here were designed as educational environments to help children engage with complex bodies of content, taking advantage of physical embodiment and social factors to promote learning.

4.1 The Island Metaphor

One way of framing the heterogeneous interaction inherent in collocated virtual worlds is through the island metaphor (Tomlinson, Baumer, & Yau, 2006). In this metaphor, stationary devices are treated as islands of virtual space, separated from one another by seas of physical space. Mobile devices act as virtual rafts, allowing virtual entities to move between different islands. The island metaphor is highly appropriate for collocated systems for a number of reasons. First, by comparing the land/water distinction to the virtual-space/real-space distinction, the metaphor offers an intuitive model for how real space and virtual space relate to each other. Users can take advantage of previous notions of the process of traveling, such as having a specific method for getting into and out of a raft. People understand that they cannot simply climb into a raft when it is in the middle of the ocean—it must be physically close to do so, and thus the physical proximity requirements of the collocated virtual world make sense.

The island metaphor also offers a highly social and natural interaction paradigm. Multi-device systems lend themselves to social interaction in two main ways. First, there is a structural parallel: just as the multiple devices are working together to form the virtual world, multiple people may work together to experience and interact with that world. Second, and more practically, a different participant may manipulate each of the devices, thereby enabling multiple participants to engage with the system simultaneously. The island metaphor leads to interactions that involve multiple people: even if everyone carries his or her own virtual raft, they must interact at the same virtual islands, thereby requiring them to be in physical proximity to each other to transfer characters between islands. Through this process, systems based on the island metaphor may encourage a more social style of interaction with the virtual world, and possibly promote social interactions among users.

This is not to say that the island metaphor is the only interaction metaphor that could be used for richly connected systems. Other interaction metaphors could bias the experience toward work-related topics (e.g., an office space metaphor), toward purely social interaction (e.g., a nightclub metaphor), or various other domains (e.g., sports, parenting, politics). The island metaphor lends its own particular focus to the experience built using it. Careful consideration of the core metaphor being used can help shape the interactions that it encourages.

We now describe the two projects created following the presented framework for richly connected systems and based on the island metaphor: the Virtual Raft project and the EcoRaft project. Although the two projects have different applications, the technical details are very similar.

4.2 The Virtual Raft Project

The Virtual Raft project (Tomlinson, Yau, O'Connell, K. Williams, & Yamaoka, 2005) is a multi-device game that teaches color theory. It consists of three desktop computers and three tablet PCs. The desktop computers display virtual "islands" and the tablet PCs display virtual "rafts." Each island contains a central bonfire and several virtual characters that hold torches of the same color as the central bonfire. The bonfires are colored red, green, and blue—three additive primary colors. When a tablet PC is brought close to a desktop computer, a virtual character "jumps" from the virtual island to the virtual raft, moving from the desktop to the tablet PC. As the tablet PC is moved close to a different tablet or desktop, the character jumps from the current raft to this new raft or island. When a user transfers a virtual character to a new island, the torch color of the virtual character will mix with the central bonfire, changing its color. For example, a character with a red torch arriving at an island with a blue flame will create a violet fire. Mixing all three colors together creates a white flame. The goal of the game is to create white fires on all three islands, which requires each island to have at least one character from each of the other islands. Users interacting with the system can discover additive color mixing, such as how combining blue and red creates violet while combining red and green creates yellow.

4.3 The EcoRaft Project

The EcoRaft project (Tomlinson, Yau, Baumer et al., 2006) is built using the same platform. It is a multi-device museum exhibit that helps children learn about restoration ecology, and was developed in collaboration with an ecology professor and her students who study restoration of Costa Rican ecosystems. Like the Virtual Raft project, the EcoRaft project consists of three desktop computers and three tablet PCs. Each desktop computer contains a virtual ecosystem, modeled after real ecosystems in Costa Rica. The ecosystems are made up of Coral trees, Heliconia flowers, and different types of hummingbirds. One of the desktop computers represents a national park, that always thrives with all of the species. The other two desktop computers represent more delicate ecosystems. These computers are each connected to a silver button, which when pressed removes all of the plants and animals from that island. Pressing the button represents the ability to over farm an island and devastate the ecosystem. Users can help restore the ecosystems by transferring plants and animal species from the National Park to the devastated islands. These species are carried using the tablet PCs (see Figure 2), which represent virtual collection boxes and can be used to physically carry the virtual species between the virtual islands. Each box can only carry a single species, so users must work together to repopulate an island. This activity teaches users that the destruction of ecosystems is very easy, and that restoration, while difficult, is still possible.

Figure 2. *Several children interact with a multi-device virtual world in the EcoRaft exhibit.*

5 Applying the Framework

In this section, we discuss in detail the ways in which the Virtual Raft and EcoRaft projects apply our framework for richly connected systems. We focus on how the multiple, coordinated connections between devices allow for novel interactions with this collocated virtual world as we explore the different components of the framework.

5.1 Multi-Device Data Networking

The Virtual Raft and EcoRaft projects use a typical digital network to transfer data information between devices. In both projects, the desktop computers are equipped with Ethernet connections and IrDA adapters, while mobile computers have built-in WiFi and IrDA capabilities. The projects use IrDA to detect when the tablet PCs are close to the desktops. IrDA is used because it operates over relatively short distances (compared to Bluetooth or wireless Ethernet) and can detect another device within approximately a 30° angle, which allows the desktop computers to determine the proximity and orientation of the tablet PCs. This allows the data connection to be linked to the physical interaction with the devices. However, as users frequently move around the space in both projects, the line of sight requirement and the slow transfer speed of IrDA make it difficult to transfer data (specifically, character information) over that connection. In coordinating the data communication with the physical and graphical modalities, we found it was hard for users to hold up the tablet PC and maintain a line of sight connection while the virtual characters graphically transfer between computers. As such, we chose to only use IrDA to detect the proximity between the devices and to fetch the unique network name of the nearby device. The system then uses a WiFi connection to send the actual data for the virtual character. WiFi is faster and does not require line of sight, so the character can still transfer even if the user breaks the IrDA connection by moving the tablet PC. Thus these projects actually coordinate two different channels of data communication, along with graphical and physical channels.

5.2 Multi-Device Graphics

These projects use the data connections to determine when an animated character should be transferred. This transfer requires coordinating the exchange of the character data with the graphical display of the character moving between the devices. To do this, we have the sending computer transfer the character information well before it starts the animation of the character moving off the device. Once the receiving device gets the character data, it creates a graphical clone of the character based on that information—a virtual character that looks and acts exactly the same as the original. However, the new character is created off-screen at first, waiting for a specified time period before it begins its own transfer animation and appears on the receiving computer's display. This time delay ensures that the character has already disappeared from the first device, creating a fluid transfer and the illusion that the same virtual character moved from one device to the other. In this way we can synchronize and coordinate the transfer across both modalities (data network and graphical) to provide the illusion of a single entity moving between devices and creating a novel interaction in which users can carry virtual characters between virtual islands.

We have also experimented with the integration of virtual lighting models that spread and persist across multi-device systems, thereby increasing the apparent graphical connectedness. When virtual characters are transferred between devices, shadow orientation information could be copied along with the character data. Thus transferring a character could propagate the shadow direction to the receiving computer. Through this process, the shadows on multiple virtual spaces could be calibrated to match each other, thereby unifying the graphics across devices more fully.

5.3 Multi-Device Sound

A connection over the audio channel is also coordinated with the data network and graphical connections when a character transfers between devices. For example, in the Virtual Raft project, when a character jumps from an island to a raft, a sound clip of a voice saying “Whee!” is played on the island, followed by a splashing water sound played on the raft device when the character lands. This aural continuity between the devices helps to support the visual continuity achieved by properly timing the animations that occur during transfer—by linking the sound connection with the other communication channels, we increase the believability of the transfer, as well as cue the user where he or she should look to follow the action. Combining these modalities into a richly connected system allows us to create a more compelling user interaction.

We also integrate the sound communication channel with the physical world, linking the physical modality with the aural. In the EcoRaft project, for example, instead of simulating a hummingbird flying by emitting the sound from permanent speakers set around the space, the hummingbird’s wing beats emit from the tablet while it is being physically carried, causing the physical action to be coordinated with the audio output. In addition, we coordinate the physical position of the islands with sound media in creating a unified experience for the user: the ambient sound emitted by each stationary computer serves as in implicit indication of that island’s ecological status. A completely deforested island is almost totally silent, with only the sound of a faint, haunting wind blowing in the background. As restoration progresses, the background audio is filled with the sounds of rustling leaves, birds chirping in the distance, and other rainforest noises. Each island has a characteristic set of background sounds, such that the full audio aesthetic of the space can only be appreciated when all the islands are fully restored. Furthermore, each individual species of hummingbird has a slightly different call, so that as different species of hummingbirds are brought to different islands, participants are immersed in a fuller, more complex soundscape. By coordinating the audio channel with the physical world, we create an interaction in which the user can simply stand in the middle of the installation and very quickly get a rough impression of the state of the entire system: both the ecological state of the three islands, and the activities occurring in the space.

5.4 Multi-Device Agents

Embodied mobile agents inhabit the virtual worlds of both implement systems. In the Virtual Raft project, humanoid agents move among the islands and rafts. In the EcoRaft project, hummingbirds and plant seeds move between the island and the virtual boxes. These agents help to create a unifying interpretation for the multiple connections between devices—the data network, graphical, audio, and physical channels come together as an animated character moving from one device to another. Thus these agents enhance the comprehensibility of the projects.

Agents can also help foster a sense of social presence in participants (Lee & Nass, 2003), and can be seen as a channel for communication: a channel of embodied social information. The believability and engagement that agents can bring to even a single-device system can help reinforce the sense of connectedness among groups of collocated devices. In both the Virtual Raft and EcoRaft projects, the agents moving graphically and audibly among the screens help participants grasp the organization of the collocated virtual world, allowing users to interpret the richly connected multi-device system.

5.5 Physical World Integration

An important step in integrating a richly connected system with the physical world involves the physical placement of devices. In both the Virtual Raft and the EcoRaft installations, the three large displays are situated roughly in a circle facing one another, so as to give the impression of three distinct virtual spaces that are not directly connected to one another. Also, these large displays are placed at the edge of the interaction space, so that participants cannot pass behind them. This placement helps to give the impression that, rather than the virtual world being contained within the display, the display is a window to a virtual space on the other side of the screen. Such placement both supports the island metaphor and helps to integrate the system with the physical world.

Another example of this integration involves the use of webcams in these projects. Other researchers have used closed circuit cameras in interactive contexts, for example to situate museum-goers within an interactive art piece (Heath et al., 2002; Hindmarsh et al., 2005). In the projects presented here, webcams are not used to place images of the participants within the installation, but rather to serve as the virtual characters' "eyes" out into the physical world. In the Virtual Raft, when a user approaches an island, the webcam mounted atop the large display detects the motion. In response, the characters stand up and approach the front of the screen, giving the appearance that they are walking toward the users who are walking toward them. If there is no motion for a period of time, the characters turn around and return to the central fire. Similarly, in EcoRaft, if the webcam detects motion, hummingbirds will fly up to the front of the screen and hover for a moment directly in front of users.

Finally, in both projects the mobile devices are also integrated with the real world. In both projects, the tablet PCs are equipped with accelerometers, which are able to detect orientation about two axes. When the tablet is tipped front-to-back or side-to-side, the virtual entities contained therein graphically react to the device's physical orientation. The humanoid characters in the Virtual Raft project try to balance on the raft—if the participant tilts the raft too much, the character falls in the water and his or her torch is extinguished. In the EcoRaft project, seeds react to tilting by rolling around inside the virtual box the tablet represents, and hummingbirds try to fly to the highest point in their virtual cage. In this way, virtual entities can react to the physical orientation of the device on which they are located, helping to blur the boundary between physical and virtual spaces.

By having multiple devices respond to the same real world phenomena, another rich connection is formed within the system. The devices are linked by their common response to an external stimulus, a coordination that can be observed by the user. By blurring the boundaries between physical space and virtual space, the fact that devices are physically distinct can also be blurred—rather than having multiple heterogeneous devices that act independently, the user instead is able to interact with a single, richly connected system.

6 Evaluation

The projects implemented using this framework have been shown to a wide variety of audiences, with over 3,000 participants at a number of different conference venues including ACM SIGGRAPH (Tomlinson, Yau, Gray, et al., 2005), ACM CHI (Tomlinson, Yau, O'Connell, et al., 2005), and CSCL (Tomlinson, 2005). In addition, these projects have been demonstrated to several hundred participants in a university lab space, as well as through temporary installations at the Discovery Science Center (DSC) in Santa Ana, CA. During these exhibitions, the project teams observed and took notes about users' interactions with and reactions to the installations. A number of different specific evaluations were also performed, including a series of open-ended, semi-structured interviews with participants at SIGGRAPH and DSC, which were aimed at evaluating the educational efficacy of the EcoRaft system (Tomlinson, Baumer, Yau, Carpenter, & Black, 2008).

Before presenting the results of these evaluations, it is important first to describe their goal. Similar to the way in which tangible interfaces (Ishii & Ullmer, 1997) enable new modes of human-computer interaction, the main purpose of using richly connected systems to build multi-device worlds is to enable new types of interactions with computational devices that were not possible in previous systems. Thus, the evaluation presented here does not take the form of proving or disproving a quantifiable hypothesis about whether or not these systems increase a sense of presence, engagement, immersion, or any other specific aspect of systems for virtual worlds. Indeed, doing any sort of with/without comparison would not make sense or even be possible here, since the example systems described above would not even be comparable without being richly connected. Rather, the evaluation focuses on understanding and explicating the subjective experiences users have with these systems. To that end, the authors employed qualitative methods (cf. Friedman et al., 2007) from anthropology and social sciences (J. Lofland & L. Lofland, 1995), including note taking, semi-structured interviews, transcript analysis, and video analysis. These data were evaluated using an open-ended, iterative coding process; researchers review the data, making note of various themes. These themes then inform further iterative analysis, until all data is thematically coded, enabling an evaluation of this framework that spans multiple deployments of multiple systems. The evaluation here presents users' experiences, as captured through these qualitative methods, with richly connected multi-device systems built using the framework presented.

6.1 User Experiences

During the various deployments, users on the whole enjoyed interacting with these richly connected multi-device worlds. When asked about preferences between the multi-device interactions and single-device ones, most preferred the multi-device interactions, responding, for example, "Yes, this is much better," "I actually like the fact that you can move things around," or "it destroys the normalcy of what the society thinks is like normal computer interfacing."

Participants also appreciated the physical aspects of the project: "I like it . . . 'cause you are walking and moving . . . You feel like you are carrying the hummingbird . . . instead of just clicking and dragging." Some participants commented that they "liked the physicality of it, . . . the fact that you walked around." Such responses speak to the strength and importance of physical world integration in richly connected systems; allowing participants to share the same space with virtual entities leads to a sense of connection not possible in other systems. Another participant offered that the way the virtual entities reacted to the tablet's orientation made them into "quasi-physical objects," which was enabled by the close coupling between the graphical and physical connections in the system. These comments indicate not only that these systems successfully met the framework's requirement for physical world integration, they also indicate that physical integration, in conjunction with other means of connection, enables participants to engage with the system in ways not possible using only one or two types of connections.

The authors have also conducted an evaluation of EcoRaft that focuses specifically on the system's efficacy in facilitating children's learning about restoration ecology (Tomlinson, Baumer, Yau, Carpenter, & Black, 2008). That evaluation consisted primarily of semi-structured, open-ended interviews with children and adults who had interacted with the EcoRaft exhibit, supplemented by observations of participants interacting with the system. The transcripts of those interviews and observation notes were analyzed using the iterative coding and qualitative methods described above. While that evaluation focuses mostly on EcoRaft's educational aspects, some of the findings are relevant to EcoRaft as a richly connected multi-device system—in particular, how the rich connections between devices help to create an environment conducive to collaborative, discovery based learning. A portion of those findings is summarized here.

One of the most prominent educational findings was that each of the 14 students interviewed mentioned the importance of collaboration. In discovery based learning, students must collaborate with one another, suggesting courses of action and rejecting or confirming one another's hypotheses. During observations, students were noted telling one another about the ecosystem's dynamics, for example, why a seed would not grow or why a hummingbird flew away, and suggesting alternate paths to accomplishing the restoration task. This collaborative form of inquiry was supported partly by the rich connections supplied by the system. For example, the integration of the system with the physical world meant that, in order for one student to demonstrate to another how to properly transfer an organism, both students often held the tablet at the same time, necessitating close cooperation. Similarly, the rich connections—through graphics, sound, networking, and physical integration—brought students together as they watched and discussed the graphical and aural movement of an organism between two physically proximate devices, a movement enabled partially by the coupling between physical and data network connections. Interactions such as these demonstrate yet another strength of richly connected systems: the ability to encourage social interaction among users.

The system design also supports different roles for participants to play. Richly connected systems can provide a wealth of interaction modalities: EcoRaft allows participants to carry three different species of organisms on three different tablets, to directly affect those organisms through the physical orientation of the tablet, to protect the islands by preventing others from pressing the silver buttons, to affect hummingbird behavior through the webcams, to perceive their impacts on the organisms and on the ecosystem as a whole through visual and aural feedback, and to interact in a variety of other ways. Furthermore, because of its physical configuration, EcoRaft also provides participants indirect ways of interacting with the installation, through observing, questioning, or making suggestions to those interacting with it. Using a richly connected system in the learning context for which the multi-device world was designed enabled the support of a broad range of learning styles.

These evaluations also led to a set of unanticipated findings about unintended results of certain physical aspects of the system. For example, one adult noted that, due to the weight and heft of the tables, the objects contained therein acquired a certain preciousness. "When you hold something with two hands" as one must do with the tablets, "you know that it's very important." On the other hand, "mov[ing] it with a mouse drag and drop, you kind of lose that interactivity that makes something sacred." While the incorporation of this theme was an unintended effect of the physical integration afforded by the tablets, the importance and sacredness of the objects being carried on the tablets helps participants assume the role of restoration ecologists that is central to EcoRaft. These findings also point to the central importance of physical world integration as a key type of connection that helps enrich the whole system's connectedness.

7 Lessons Learned

The evaluation section above demonstrates the ways in which richly connected systems can enable novel types of interactions. However, it may be helpful to describe lessons learned during the process of creating these systems—the practical aspects of how we went about building these multi-device worlds using richly connected systems. This section serves as a summary of several major conclusions gained from the designers' and researchers' experiences with the systems, and is presented in the interest of facilitating the development of other multi-device worlds and richly connected systems.

One of the characteristics of most current groups of collocated devices is some degree of heterogeneity. While the range of devices available today is certainly different from the range of devices that will be available in several years, it is very likely that absolute standardization across devices will not occur, and heterogeneity will continue to factor prominently in multi-device computing. The differences among devices can be both a benefit and a challenge for developers. For example, having different types of devices enables each device to be used in ways for which it is particularly well-suited: mobile devices can move around the space, high resolution monitors can display crisp imagery, and devices with accelerometers can harness their physical capabilities. Having devices serve different roles can enable a wide variety of different combinations of devices and functionality.

However, implementing software for systems of heterogeneous devices is frequently challenging. Developing software for multiple platforms has proved to be a bit arduous, as current widely available programming tools lack functionality to aid development in which code is written on one computer and immediately run on another. Even though the Virtual Raft and EcoRaft projects were written in Eclipse, one of the most popular integrated development environments for Java programming, we had no automated way to write code for two or more separate interacting programs on one of our desktop PCs, compile the programs, transfer the appropriate programs to various tablets and desktops, and run all the programs on their respective devices. While Eclipse does provide some remote debugging tools, including the ability to connect to a remotely running Java virtual machine, it lacks the sort of all-inclusive solution required. In addition, many of the features coders have become accustomed to (such as the ability to perform an automatic stack trace in Eclipse) were not available to our team when writing code for a remote machine, as remote debugging tools were not available at that time. To help patch this hole in tool functionality, we are developing a plugin to Eclipse that will facilitate the development of multi-device systems by addressing the issues mentioned above and automating the debugging cycle for multi-platform systems. Finding or building an effective set of tools for multi-device development would be an excellent first step for other teams seeking to develop richly connected systems.

An interesting realization that we had while watching people interact with the multi-device exhibits was that children appear to expect rich connectivity. For example, during one interaction, a group of about 10 children were using the system, but there were only three tablet PCs. One seven-year-old child, who was not carrying a tablet, walked up to one of the virtual islands and held out his hand toward it, palm up. He had grasped the gist of the core interaction—that the characters on the monitor would jump out of the screen—but had decided that the tablet might not be necessary; perhaps any flat surface would do. This willingness to participate in richly connected systems may not be something that people need to learn; we have vast experience in it from our multimodal interpersonal interactions. Rather, those of us who have been using sparsely connected devices for the past several decades may need to unlearn our expectations about the impoverished state of current multi-device HCI.

Nevertheless, many of the participants who will engage with these kinds of systems over the next few years will have these limited expectations for inter-device interactivity. Therefore, maintaining an awareness of the expectations that users will bring is a critical part of designing novel systems. Graphics, sound, and physical interaction are the modalities that people expect from most consumer-grade computing devices and for which those devices are designed, so these modalities are useful for near-term multi-device systems as well.

Over the longer term, the concept of a richly connected system suggests that designers should seek to connect the devices through as many different channels as possible. Humans should share at least some of these channels—for example: graphics, sound, and other capabilities within the human sensory repertoire. While a given interaction may not require every possible communication channel across devices, it may be hard to appreciate the benefit of connecting a channel until after the connection has been made and designers have an opportunity to experiment with it in concert with other media.

8 Future Work

This research is continuing to progress in a number of areas, including the creation of a series of exhibits based on different regional ecosystems, and the development of a new framework for interaction with multi-device systems.

8.1 Network of Exhibits

The EcoRaft exhibit described above focuses on a Costa Rican rain forest ecosystem. The future stages of this project will develop a network of six interactive museum exhibits based on common themes in restoration ecology. Each of the six exhibits will address an ecological issue that is relevant to the geographical region in which it is displayed. For example, an exhibit in Florida might feature the snakehead fish and one or more native species of fish, while an exhibit in Minnesota might feature wolves and rabbits. Each participating museum will be able to run the regional content that was developed for the other museums as well, thereby encouraging repeat visitation and helping visitors learn about ecological principles that stretch across different ecosystems. To deploy these exhibits, we are recreating the EcoRaft project in a more extensible form, creating a kind of API that other developers can use to create their own exhibits.

8.2 Human-Mediated Networking

In both the Virtual Raft and EcoRaft projects, users carry information (embodied in the form of graphical agents) between computers. They are thus helping to transfer this information around the system, participating in a form of networking. We are currently exploring the opportunities afforded by this type of multi-device interaction, where human actions help to network computers, which we call *Human-Mediated Networking* (HMN). Just as computer mediated communication deals with the ways in which computers facilitate interactions between people, HMN addresses ways in which people help coordinate interactions among devices. Having humans carry embodied agents between computers, as in the described projects, is an example of such networking. By taking advantage of human actions and behaviors, we can allow for new types of systems, which may be more environmentally sustainable as networking resources are replaced by networking labor.

9 Conclusion

This paper has presented a framework for the design and implementation of richly connected systems on networks of collocated heterogeneous devices. This framework coordinates data networking, graphics, sound, physical world integration, and embodied mobile agents into a unified device-to-device connection, allowing for novel forms of interaction. The framework was used in the production of two interactive projects: the Virtual Raft project and the EcoRaft project. An evaluation of these projects was also offered, describing the ways in which this framework allowed for the creation of an engaging experience for users. This paper also described a number of lessons learned in the development of this framework. While richly connected systems are not a solution to the entire broad problem of enabling people and devices to work together more effectively, this framework does provide a possible means for creating a coherent interaction paradigm across multiple collocated devices.

As computational devices become more common across human societies, the potential usefulness of groups of these devices that are physically proximate to each other increases significantly. The growing frequency with which people find themselves in the presence of several different devices necessitates more effective ways for people to engage with those devices as systems, rather than in isolation. Just as people can achieve greater functionality when they work together, so too can devices become more useful and effective when they are enabled to operate smoothly together. Richly connected systems are just one example of this potential future application domain.

Acknowledgments

This research was funded the Nicholas Foundation, the Emulex Foundation and NSF grant IIS-0644415. It was supported by the California Institute for Telecommunications and Information Technology (Calit2) and the Donald Bren School of Information and Computer Sciences. The authors thank the other members of the Social Code Group, who contributed greatly to the implementations and ideas described here.

References

Baecker, R. M., Grudin, J., Buxton, W. A. S., & Greenberg, S. (Eds.). (1995). *Readings in human-computer interaction.* San Mateo, CA: Morgan Kaufmann.

Barab, S., Thomas, M., Dodge, T., Carteaux, R., & Tuzun, H. (2005). Making learning fun: Quest Atlantis, a game without guns. *Educational Technology Research and Development, 53*(1), 86–107.

Barkhuus, L., Chalmers, M., Tennent, P., Hall, M., Bell, M., Sherwood, S., et al. (2005). Picking pockets on the lawn: The development of tactics and strategies in a mobile game. *Ubiquitous Computing: Ubicomp 2005*, 358–374.

Bates, J., Loyall, A. B., & Reilly, W. S. (1994). An architecture for action, emotion, and social behavior. *Selected Papers from the 4th European Workshop on Modeling Autonomous Agents in a Multi-Agent World, Artificial Social Systems,* 55–68.

Biocca, F. (1997). The cyborg's dilemma: Embodiment in virtual environments. *Humanizing the Information Age, Proceedings of the Second International Conference on Cognitive Technology,* 12–26.

Blumberg, B. M., & Galyean, T. A. (1995). Multi-level direction of autonomous creatures for real-time virtual environments. *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques,* 47–54.

Borovoy, R., Martin, F., Vemuri, S., Resnick, M., Silverman, B., & Hancock, C. (1998). Meme tags and community mirrors: Moving from conferences to collaboration. *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work,* 159–168.

Brumitt, B., Meyers, B., Krumm, J., Kern, A., & Shafer, S. A. (2000). EasyLiving: Technologies for intelligent environments. *Proceedings of the 2nd International Symposium on Handheld and Ubiquitous Computing,* 12–29.

Card, S., Moran, T., & Newell, A. (1983). *The psychology of human-computer interaction,* p. 488. Mahwah, NJ: Lawrence Erlbaum.

Cheok, A. D., Fong, S. W., Goh, K. H., Yang, X., Liu, W., & Farzbiz, F. (2003). Human Pacman: A sensing-based mobile entertainment system with ubiquitous computing and tangible interaction. *Proceedings of the 2nd Workshop on Network and System Support for Games,* 106–117.

Drettakis, G., Roussou, M., Reche, A., & Tsingos, N. (2007). Design and evaluation of a real-world virtual environment for architecture and urban planning. *Presence: Teleoperators and Virtual Environments, 16*(3), 318–332.

Drucker, S. M., He, L., Cohen, M., Wong, C., & Gupta, A. (2003). Spectator games: A new entertainment modality for networked multiplayer games. Retrieved October 5, 2008. Available at http://research.microsoft.com/~sdrucker/papers/spectator.pdf.

El-Nasr, M. S., Zupko, J., & Miron, K. (2005). Intelligent lighting for a better gaming experience. *CHI '05 Extended Abstracts on Human Factors in Computing Systems,* 1140–1141.

Friedman, D., Leeb, R., Guger, C., Steed, A., Pfurtscheller, G., & Slater, M. (2007). Navigating virtual reality by thought: What is it like? *Presence: Teleoperators and Virtual Environments, 16*(1), 100–110.

Gratch, J., Rickel, J., Andre, E., Cassell, J., Petajan, E., & Badler, N. (2002). Creating interactive virtual humans: Some assembly required. *Intelligent Systems, IEEE, 17*(4), 54–63.

Grudin, J. (1988). Why CSCW applications fail: Problems in the design and evaluation of organizational interfaces. *Proceedings of the 1988 ACM Conference on Computer-Supported Cooperative Work,* 85–93.

Grudin, J., & Palen, L. (1995). Why groupware succeeds: Discretion or mandate? *Proceedings of the Fourth European Conference on Computer-Supported Cooperative Work,* 263–278.

He, L., Cohen, M. F., & Salesin, D. H. (1996). The virtual cinematographer. *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques,* 217–224.

Heath, C., Luff, P., Lehn, D. V., Hindmarsh, J., & Cleverly, J. (2002). Crafting participation: Designing ecologies, configuring experience. *Visual Communication,* (1), 9–33.

Hindmarsh, J., Heath, C., Lehn, D. V., & Cleverly, J. (2005). Creating assemblies in public environments: Social interaction, interactive exhibits and CSCW. *Computer Supported Cooperative Work (CSCW), 14*(1), 1–41.

Hughes, C. E., Reinhard, E., Konttinen, J., & Pattanaik, S. (2004). Achieving interactive time realistic illumination in mixed reality. *Army Science Conference.*

Ishii, H., Mazalek, A., & Lee, J. (2001). Bottles as a minimal interface to access digital information. *CHI '01 Extended Abstracts on Human Factors in Computing Systems,* 187–188.

Ishii, H., & Ullmer, B. (1997). Tangible bits: Towards seamless interfaces between people, bits and atoms. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems,* 234–241.

Khoo, E. T., Lee, S. P., Cheok, A. D., Kodagoda, S., Zhou, Y., & Toh, G. S. (2006). Age invaders: Social and physical inter-generational family entertainment. *CHI '06 Extended Abstracts on Human Factors in Computing Systems,* 243–246.

Lee, K. M., & Nass, C. (2003). Designing social presence of social actors in human computer interaction. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems,* 289–296.

Lofland, J., & Lofland, L. (1995). *Analyzing social settings: A guide to qualitative observation and analysis.* Belmont, CA: Wadsworth.

Maes, P. (1995). Artificial life meets entertainment: Lifelike autonomous agents. *Communications of the ACM, 38*(11), 108–114.

Mynatt, E. D., Back, M., Want, R., Baer, M., & Ellis, J. B. (1998). Designing audio aura. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems,* 566–573.

Naemura, T., Nitta, T., Mimura, A., & Harashima, H. (2002). Virtual shadows in mixed reality environment using flashlight-like devices. *Transactions of the Virtual Reality Society of Japan, 7.*

Norman, D. (1988). *The psychology of everyday things.* New York: Basic Books.

O'Hare, G. M. P., & Duffy, B. R. (2002). Agent chameleons: Migration and mutation within and between real and virtual spaces. *The Society for the Study of Artificial Intelligence and the Simulation of Behaviour (AISB),* 3–5.

Ong, W. (1982). *Orality and literacy* (2nd ed., 2002). New York: Routledge.

Perlin, K., & Goldberg, A. (1996). Improv: A system for scripting interactive actors in virtual worlds. *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques,* 205–216.

Rekimoto, J. (1997). Pick-and-drop: A direct manipulation technique for multiple computer environments. *Proceedings of the 10th Annual ACM Symposium on User Interface Software and Technology,* 31–39.

Roussos, M., Johnson, A., Moher, T., Leigh, J., Vasilakis, C., & Barnes, C. (1999). Learning and building together in an immersive virtual world. *Presence: Teleoperators and Virtual Environments, 8*(3), 247–263.

Sawhney, N., & Schmandt, C. (2000). Nomadic radio: Speech and audio interaction for contextual messaging in nomadic environments. *ACM Transaction on Computer-Human Interaction, 7*(3), 353–383.

Stock, O., & Zancanaro, M. (2002). Intelligent interactive information presentation for cultural tourism. In *Proceedings of the International CLASS Workshop on Natural Intelligent and Effective Interaction in Multimodal Dialogue Systems,* 28–29.

Sumi, Y., & Mase, K. (2001). AgentSalon: Facilitating face-to-face knowledge exchange through conversations among personal agents. *Proceedings of the Fifth International Conference on Autonomous Agents,* 393–400.

Tomlinson, B. (2005). A heterogeneous animated platform for educational participatory simulations. *Proceedings of the 2005 Conference on Computer Support for Collaborative Learning; Learning 2005: The Next 10 Years,* 677–681.

Tomlinson, B., Baumer, E., & Yau, M. L. (2006). The island metaphor. *ACM SIGGRAPH 2006 Research Posters,* 84.

Tomlinson, B., Baumer, E., Yau, M. L., Carpenter, F. L., & Black, R. (2008). A participatory simulation for informal education in restoration ecology. *E-Learning, 5*(3), 238–255.

Tomlinson, B., Blumberg, B., & Nain, D. (2000). Expressive autonomous cinematography for interactive virtual environments. *Proceedings of the Fourth International Conference on Autonomous Agents,* 317–324.

Tomlinson, B., Yau, M. L., & Baumer, E. (2006). Embodied mobile agents. *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems,* 969–976.

Tomlinson, B., Yau, M. L., Baumer, E., Goetz, S., Carpenter, L., Pratt, R., et al. (2006). The EcoRaft project: A multi-device interactive graphical exhibit for learning about restoration ecology. *CHI'06 Extended Abstracts on Human Factors in Computing Systems,* 1445–1450.

Tomlinson, B., Yau, M. L., & Gray, J. (2005). Heterogeneous character animation: How to make an interactive character jump between stationary and mobile graphical computing platforms. *ACM SIGGRAPH 2005 Sketches,* 18.

Tomlinson, B., Yau, M. L., Gray, J., Baumer, E., O'Connell, J., Williams, K., et al. (2005). The virtual raft project: A network of mobile and stationary computer systems inhabited by communities of interactive animated agents. *ACM SIGGRAPH 2005 Emerging Technologies,* 30.

Tomlinson, B., Yau, M. L., O'Connell, J., Williams, K., & Yamaoka, S. (2005). The virtual raft project: A mobile interface for interacting with communities of autonomous characters. *CHI'05 Extended Abstracts on Human Factors in Computing Systems,* 1150–1151.

Weiser, M., & Brown, J. S. (1996). Designing calm technology. *PowerGrid Journal, 1*(1), 75–85.

Williams, M., Jones, O., Fleuriot, C., & Wood, L. (2005). Children and emerging wireless technologies: Investigating the potential for spatial practice. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems,* 819–828.