

Towards A Usable API for Constructing Interactive Multi-Surface Systems

Chris Burns, Teddy Seyed, Theodore D. Hellmann, Jennifer Ferreira, Frank Maurer, Mario Costa Sousa
University of Calgary, Department of Computer Science
2500 University Drive NW, Calgary, Alberta, Canada, T2N 1N4
{chris.burns, teddy.seyed, tdhellma, jen.ferreira, frank.maurer, smcosta}@ucalgary.ca

ABSTRACT

Research into multi-surface systems goes back for more than thirty years, yet these systems have not been taken up in real-world settings. We believe the reason for the lack of adoption is that constructing multi-surface systems is costly and requires specialist knowledge of tasks related to device discovery, cross-platform interoperability, networking, and spatial tracking. These tasks represent a significant distraction from implementing features that actually matter to end users. While some APIs exist for supporting the set-up of multi-surface systems, they are directed at specialist developers. We propose to develop a highly learnable API for constructing multi-surface systems, which is targeted at non-specialists.

Keywords

Multi-surface system, API Usability, API Design

INTRODUCTION

Multi-surface systems (MSS) integrate multiple, heterogeneous computing devices into a single application solution. They rely on Natural User Interface (NUI) approaches (as opposed to multi-display environments that primarily focus on WIMP-based interfaces on multiple displays). With over thirty years of research in MSS, real world and deployed systems are rare. We hypothesize that this is due to the cost and complexity of developing such applications for use in an interactive space. Currently, a developer creating such an application would have to be knowledgeable about device discovery, cross-platform interoperability, networking, spatial tracking, and other tasks. While these ancillary tasks are necessary for the application as a whole to work naturally, these tasks represent a significant distraction from implementing features that matter to end users.

To reduce the amount of this specialist knowledge that developers require to build applications for use in interactive spaces, developers need advanced and usable tool support. We believe this tool support should take the

form of a highly usable and reusable application programming interface (API). This API should allow developers to focus on development of their applications rather than on ancillary tasks like interpreting video input, interpreting depth sensor data or advanced 3d graphics processing. The API should also include functionality for handling difficult, ambiguous cases for interaction in multi-surface interactive spaces – such as automatically determining the devices that different users intend to interact with. These situations will become increasingly common as interactive spaces increasingly incorporate multiple devices in one room, e.g., smartphones and tablets. It is important that this API be highly usable – especially in terms of learnability – so that it can be easily adopted and used by typical development teams.

In this paper, we discuss existing applications for interactive spaces, challenges to the development of such systems, and the characteristics of an API that would be better able to support their development. A brief description of our work towards the creation of such an API is also included at the end of this paper.

INTERACTIVE SPACES

Several definitions of multi-surface interactive spaces have been proposed. In 2011, Gjerlufsen et al. used the following to describe interactive spaces: “Multi-surface environments are ubiquitous computing environments where interaction spans multiple input and output devices and can be performed by several users simultaneously” [6]. However, in 2006, Shen et al. made use of a much more specific definition: “By using the term multi-surface, instead of multi-display, we emphasize the nature of many of today’s interactive walls, tables, Tablet PCs, desktop displays, laptops and PDAs that often can be interacted upon in addition to be merely the visual display” [10]. We follow the latter definition because it allows us to focus on systems in which devices participate in both interaction and display of data. Few real-world systems actually fit under this description, i.e., systems that are deployed and in active use outside of research labs.

The research literature has many examples of prototypes of multi-surface systems created and used within research laboratories. The earliest example is i-Land, a system developed by Streitz et al. in 1999 [11]. This early interactive system included a tabletop and wall display. Many of the interactions with this system focused on

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PPD’12, May 25, 2012, Capri, Naples, Italy.

Copyright 2012 ACM 1-58113-000-0/00/0010...\$10.00.



Figure 1: “Pouring” data from an iPad to a tabletop.

transfer of data between devices without tracking those devices. Since then, interactive spaces have been extended to include tracking of people and devices in interactive spaces using motion capture systems like those produced by VICON¹. The WILD room, for example, employs such a tracking system to support interactions based on the position of individuals and items in the interactive space [1]. Code Space is a system developed at Microsoft Research that uses depth-sensing cameras to support interactions between a digital tabletop and a large-format wall display [3]. This system employs touch gestures as well as gestures performed in the air to support collaborative meetings. Touch gestures are performed on the device, using the multi-touch capabilities of the device. Gestures performed in the air are physical gestures, such as waving or pointing. Both types of interactions are accomplished through integration of smartphones into the interactive space.

Real-world products are harder to find. One real-world example is the PBCave². This system is commercially available, can be used for information visualization and is composed of a digital tabletop and wall display. It does not, however, support integration with now-ubiquitous smartphones and tablets to provide spatial tracking.

The lack of examples of interactive spaces in use outside the laboratory is troubling given the long history of research in this field. In order to address this issue, we should first ask: what makes development of applications for use in interactive spaces difficult? We explore these development challenges in the next section.

DEVELOPMENT CHALLENGES

We believe that there are two causes of the lack of real-world interactive spaces: (1) the cost of hardware commonly used in these systems, and (2) the complexity of developing applications on top of a multi-surface environment. We believe that both of these difficulties can, at least partially, be overcome with an API which supports the use of widely available hardware (cost reduction) and which developers can use to effectively incorporate their

¹<http://www.vicon.com/>

²http://www.pbworld.com/capabilities_projects/visualization/cave.aspx

applications into an interactive space (complexity reduction). The API needs to provide a certain set of features which we'll discuss in the following.

Spatial Tracking

Spatial tracking allows a system to track the position of people, devices, and other items in the interactive space. Couple with a layout of the space, a variety of proxemics interactions can be supported. In the past it was necessary to use expensive technology to accomplish tracking. With the introduction of the Microsoft Kinect, it is now possible to get motion-tracking using more widely-available technology. Using the Kinect together with orientation-aware devices – such as Apple's iPhone and iPad – it is possible to track both the position and orientation of these devices. However, compiling low-level information about the position and information of a device into meaningful high-level information is a complicated task. These tasks require a significant amount of mathematical knowledge and also involve cross-platform communication between components of the system.

Heterogeneous Device Integration

Digital tabletops are large, touch-enabled surfaces that are good for collaborative work. Examples of these devices include the Evolve One³, the SMART Table⁴, and the Microsoft Surface⁵ which all run Microsoft Windows. On the other hand, most smartphones and tablets run either iOS or Android operating systems. In order to build a multi-surface system that supports common existing devices, developers are required to overcome the significant challenges that exist in trying to communicate information between these platforms. What is needed is a software architecture that supports device discovery and message passing between heterogeneous devices.

API DESIGN

In this section, we propose the two main features of an API for the development of applications that run in multi-surface interactive spaces. We believe supporting spatial awareness and communication between devices would simplify the implementation of these systems. Such multi-surface interactive spaces would be useful in the domains of data visualization and analysis in the context of decision support. Users would be professionals within the same discipline or closely related to the data and task in hand.

Spatial Awareness

The API must be aware of people and devices in the room. Spatial awareness is the ability to determine position and orientation over time – including mobile devices such as tablets and smartphones and fixed devices such as digital tabletops and wall-sized displays. The layout of the room

³ http://www.evolve.com/en/hardware/multi-touch_table.php

⁴ <http://smarttech.com/table>

⁵ <http://www.microsoft.com/surface>



Figure 2: A “Flick” gesture transferring data from an iPad to a wall sized display

will be specified using a custom web based tool that allows end-users to define the locations of items within the room. The location of mobile devices can be tracked using a Kinect while orientation data can be captured from the gyroscopes already built into many mobile devices. This data can be transmitted back to the system and integrated with the fixed position data to support spatial interactions.

Communication

The API must also allow all devices in the system to communicate without requiring developers to consider the platform-specific details of each device. This should be accomplished using HTTP as the transportation layer, and implementing a REST-ful interface for each device. To simplify communication tasks further, client libraries should be implemented for each common platform, such as Android, iOS and Windows. Developers could simply subscribe to communication events using the language common to their platform, for example, Java for Android or Objective-C for iOS. To preserve consistency across all platforms, the API will set minimum hardware requirements for supported devices.

SUPPORTED INTERACTIONS

Several types of interactions should be supported by the API to help designers and developers in the creation of multi-surface applications. Each interaction should be treated as a *first-class* event in the API. Developers will be able to assign certain functionality to be triggered on each device when it receives a specific event. For example, when a flick gesture is sent from one device to another, the system will alert the target device with a notification of this event. This design will allow designers and developers to add these interactions into their application with very little time and effort. This section further describes the different types of gestures that should be supported by the API.

Proxemic Interactions

Proxemic interactions have been explored in detail in the field of HCI. Centrally, Marquardt et al. presented Proximity Toolkit, an API for supporting proxemic interactions [8]. Developers can define functionality to be triggered based on the spatial information. This toolkit, however, relies on the use of VICON cameras. However, it

is now possible use the the Kinect for motion capture instead at a much lower cost.

Physical Gestures

Physical gestures performed in the air are triggered by users moving their arms, hands, or fingers. For example, in some applications, users are able to control applications by pointing at icons as a means of choosing a selection or waving their hands to go back up one level of a directory hierarchy. The implementation of these gestures in an application is currently a very low-level, complicated process. In order to promote the use of these interactions in multi-surface applications, the API should provide high-level methods to allow these gestures to be used by designers and developers who do not have the specialist knowledge that would normally be required to implement these gestures.

Device Gestures

A device gesture is a gesture performed by a user physically interacting with or moving a mobile device. These gesture types can be subdivided into *control* and *information-passing* gestures. Control gestures are distinct from gestures which pass information. A basic set of gestures will be developed by the API designers while leaving flexibility for new gestures to be added by application developers.

1.1.1 Between-Device Control Gestures

Gestures can be used to allow one device to control another. Touch interactions with a smartphone could be used to trigger events on another device in the interactive space. In the literature there is an example of using a multi-touch tabletop to control the interface of a mobile phone, such as work by Olsen et al who presents a paper on “spilling” control from one a mobile phone to a multi-touch tabletop [9].

1.1.2 Information –Passing Gestures

Certain gestures can be used to support information transfer between applications running on different devices e.g. flicking summoning or pouring. Gestures performed with devices seem especially appropriate for this kind of task. Some work in the literature exists describing these gestures. Bhandari and Lim describe a system in which an entire smartphone is moved in a specific way to trigger interactions [2]. For example, they proposed the rotate gesture, which is triggered by rotating the device from the horizontal to the vertical position. A throw and pull gesture was shown by Döring et al. [5] for communicating data between a tabletop and a mobile device, and this gesture was also used by Daschelt and Buchholz for communicating data to large public displays [4]. Finally, the chucking gesture – a one handed gesture using a mobile phone – was proposed by Hassan et al. as another simple data-transfer gesture [7].

Existing APIs

Gjerlufsen et al. [1] developed an API for developing multi-surface systems. This API provides much of the functionality required for developing these systems. It

includes support for heterogeneous devices (devices on different platforms), a communication layer, and access to position information for devices in the interactive space. While these are desirable characteristics, the API was developed using a data-oriented programming model rather than the more common object oriented approach. This approach is less common and thus, may make it more difficult for typical developers to use this API. Both these issues make it difficult for non-expert developers to use.

Another existing API is the Proximity Toolkit developed by Marquadt et al. [8]. This API allows developers to work with proxemic interactions, but it does not provide a communication layer or support for other interaction types.

STATE OF IMPLEMENTATION

We are currently developing a prototype multi-surface system. This prototype uses an Evolve Tabletop, a SMART board and 2 iPads. The system also uses the Microsoft Kinect to gather positional tracking and iPad's internal gyroscope to provide orientation tracking. We are experimenting with several physical gestures for this system, such as flick, summon and pour, see Figure 1 and Figure 2. In the flick gesture, the users swipe across the iPad while pointing in the direction of the target device. The summoning gesture allows users to pull content from a target, by rapidly pulling back the iPad. Finally, by placing the iPad over the target tabletop and rotating it on its side, the pouring gesture allows users to transfer content to the tabletop. We intend to extend this prototype into a complete API over the next several months.

FUTURE WORK

Our focus is on easing the implementation work involved in setting up multi-surface interactive spaces. Therefore, we consider an API usable if it exhibits high *learnability*. All the tasks that are common to implementing spatial awareness and communication between devices should be supported by the API. The learnability of the API can be evaluated with user studies – having users carry out small but specific tasks with the API in a similar way to the study by Stylos and Myers [12]. To make sure that the API is widely adopted we will work to ensure that all major use cases, required by application developers, are supported by the API.

CONCLUSION

This paper has suggested that a lack of real world adoption of interactive multi-surface systems is due to the lack of a learnable API. The API would need to run on common platforms and to solve ancillary tasks common to setting up any multi-surface system. These tasks include spatial tracking and communication. We propose that a highly learnable API targeted at non-specialist developers will reduce costs and make interactive multi-surface systems more widely available.

ACKNOWLEDGMENTS

We would like to thank our colleagues for their useful discussions and advice. We also thank the anonymous

reviewers for their careful and valuable comments and suggestions. This research was partially funded by the NSERC/AITF/Foundation CMG Industrial Research Chair in Scalable Reservoir Visualization, and by the NSERC SurfNet Research Network.

REFERENCES

1. Beaudouin-Lafon, M. Lessons Learned from the WILD Room, a Multisurface Interactive Environment. *Proc. IHM '11*, 105-112.
2. Bhandari, S and Lim, Y. Exploring Gestural Mode of Interaction with Mobile Phones. *Proc. CHI EA '08*, 2979-2984.
3. Bragdon, A., DeLine, R., Hinkley, K., and Morris, M.R. Code Space: Touch + Air Gesture Hybrid Interactions. *Proc. ITS '11*, 212-221.
4. Dashelt, R. and Buchholz, R. Natural throw and Tilt Interaction between Mobile Phones and Distant Displays. *Proc. CHI EA '09*, 3253-3258.
5. Döring, T., Shirazi, A.S., and Schmidt, A. Exploring Gesture-Based Interaction Techniques in Multi-Display Environments with Mobile Phones in a Multi-Touch Table. *Proc. AVI '10*, 419.
6. Gjerlufsen, T., Klokmose, C. N., Eagan, J., Pillias, C., and Beaudouin-Lafon, M. Shared Substance: Developing Flexible Multi-Surface Applications. *Proc. CHI '11*, 3383-3392.
7. Hassan, N., Rahman, M.M., Irani, P., and Graham, P. A Chucking: A One-Handed Document Sharing Technique. *Proc. INTERACT '09*, 264-278.
8. Marquadt, N., Diaz-Marino, R., Boring, S., and Greenberg, S. The Proximity Toolkit: Prototyping Proxemic Interactions in Ubiquitous Computing Ecologies. *Proc. UIST '11*, 315-326.
9. Olsen Jr., D. R., Clement, J., and Pace, A. Spilling: Expanding Hand-Held Interactions to Touch Table Displays. *Proc. TABLETOP '07*, 163-170.
10. Shen, C., Esenther, A., Forlines, C., and Ryall, K. Three Modes of Multi-Surface Interaction and Visualization. *CHI 2006, TR2006-025*.
11. Streitz, N.A., Geißelr, J., Homber, T. et al. i-Land: An Interactive Landscape for Creativity and Innovation. *Proc. CHI '99*, 120-127.
12. Stylos, J. and Myers, B. A. The Implications of Method Placement on API Learnability. *Proc. SIGSOFT '08*, 105-112.