
Interactive Multi-Modal Display Spaces for Visual Analysis

Thomas Marrinan

Argonne National Laboratory
Illinois USA
tmarrinan@anl.gov

Arthur Nishimoto

University of Illinois at Chicago
Illinois USA
anishi2@uic.edu

Joseph A. Insley

Argonne National Laboratory
Illinois USA
insley@anl.gov

Silvio Rizzi

Argonne National Laboratory
Illinois USA
srizzi@alcf.anl.gov

Andrew Johnson

University of Illinois at Chicago
Illinois USA
ajohnson@uic.edu

Michael E. Papka

Argonne National Laboratory &
Northern Illinois University
Illinois USA
papka@{anl.gov, nui.edu}

Abstract

Classic visual analysis relies on a single medium for displaying and interacting with data. Large-scale tiled display walls, virtual reality using head-mounted displays or CAVE systems, and collaborative touch screens have all been utilized for data exploration and analysis. We present our initial findings of combining numerous display environments and input modalities to create an interactive multi-modal display space that enables researchers to leverage various pieces of technology that will best suit specific sub-tasks. Our main contributions are 1) the deployment of an input server that interfaces with a wide array of interaction devices to create a single uniform stream of data usable by custom visual applications, and 2) three real-world use cases of leveraging multiple display environments in conjunction with one another to enhance scientific discovery and data dissemination.

Author Keywords

Multiple Display Environments; multi-user interaction; collaboration; input devices; large-scale displays; virtual reality; multi-touch screens; motion capture.

ACM Classification Keywords

H.5.2. User Interfaces: Input devices and strategies, user centered design.

Permission to make digital or hard copies of part or all 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. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

Copyright is held by the owner/author(s).

ISS '16, November 06-09, 2016, Niagara Falls, ON, Canada

ACM 978-1-4503-4248-3/16/11.

<http://dx.doi.org/10.1145/2992154.2996792>

Introduction

Scientific data exploration, analysis, and dissemination have been aided by the use of specialized display technologies and interaction techniques. Large high-resolution displays have been shown to amplify sensemaking abilities and foster collaborative task execution [1]; Virtual Reality (VR) environments, such as head-mounted displays (HMDs), provide additional depth cues and a sense of immersion for three-dimensional data [8]; and GPU enabled clusters allow big data problems to be visualized with ultra-fine detail [7]. Interaction devices also play a role in benefiting exploration, analysis, and dissemination tasks. Devices such as a multi-touch screen, gaming controller, gyromouse, and motion capture system enable users to collaborate and move around a display environment's space freely, without being tied to a desk with a mouse and keyboard. While each of these technologies have provided some level of benefit to scientists, researchers have begun examining the combination of multiple pieces of technology.

Much of the research on Multiple Display Environments (MDEs) surrounds interaction paradigms to support cross-display navigation [10], the development of tools that leverage disparate screens [2,9], and studies to determine MDEs impact on collaborative tasks [5]. In each of these cases, users interface with a single application using multiple display technologies. However, when performing real-world analysis, researchers often require comparing multiple pieces of data from multiple applications to gain meaningful insight. Therefore, we investigated the use of multiple applications networked together with the ability to share certain data, rather than using just one application that runs in a MDE.

Specialized hardware, such as ultra high-resolution displays, HMDs, and collaborative touch screens, are specifically designed to enable analysis for specialized tasks. Ultra high-resolution displays enable scientific visualizations to show extraordinary detail while still being able to provide context to viewers. HMDs block out objects in the physical world and provide additional depth cues to users viewing three-dimensional data, thereby creating an immersive virtual experience. Large touch screen displays, on a tabletop device or vertical flat-screen monitor, enable collaborative interaction since they support multiple simultaneous touch events. Each of these specialized devices, along with personal devices such as laptops and tablets, can be leveraged to support various sub-tasks of an overall data exploration and analysis task. We present our initial findings on three our use cases where data is distributed amongst computing resources, and each display technology in a MDE is running a separate application to fully leverage the affordances of its hardware for exploration, analysis, and dissemination tasks.

Prior to discussing real-world uses of MDEs running networked applications, we present the development and deployment of an input server that interfaces with a wide array of interaction devices to create a single uniform stream of data usable by these custom visual applications. Enabling multiple modes of interaction is important for data exploration and analysis, especially when done collaboratively. Typical computer systems only allow a single user to interact at a time using a mouse and keyboard. Large touch screens are an exception to this general case, where multiple users can simultaneously interact with content on the display. However, enabling collaborative tasks on non-multi-

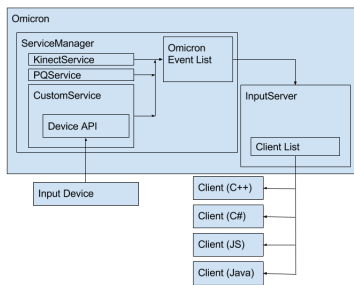


Figure 1. Architecture for the Omicron input abstraction library. The Service Manager receives native events from a variety of input devices and converts them to an Omicron event list. An Input Server application can be launched to stream these events to multiple simultaneous clients (currently implemented in C++, C#, Java, and JavaScript).

touch display environments and not tethering users to a mouse, keyboard, or touchscreen required integrating multiple non-standard input devices, such as gaming controllers and motion tracking systems, into visualization applications. Additionally, by deploying an input service, we allow separate applications to receive input from the same devices simultaneously, thereby enabling multiple applications to receive synchronized updates from a series of input devices.

Multi-Input Service

An input abstraction library, Omicron, was designed to integrate multiple different input modalities. These inputs are abstracted using modular event services, providing clients access to a specific set of input devices. The Omicron library can be used to directly integrate the event stream into a C++ application as a static library or run as a standalone Input Server to stream events using the Omicron connector API to multiple applications or across display clusters. The connector API client interface has been implemented for C++, C#, Java, and JavaScript. The entire library is open-source at <https://github.com/uic-evl/omicron>.

Omicron is comprised of two major elements: the Service Manager and the Input Server. The Service Manager prioritizes the various input service modules that can be attached to the Omicron API. Each service implements the native API of the input device to receive events and is then categorized into a more general Service Type, which identifies similar event types. For example a mouse and touch event share similar base structure such as a 2D position, ID, and an event action (down, move, drag, click). Both these devices would fall under the 'Pointer' Service Type. Similarly a motion tracked object from an optical tracking system and a Kinect skeleton share 3D

positional information and would be categorized as a 'Mocap' Service Type. Other supported Service Types are 'Controllers' (Xbox 360, PS3, Wiimote, gyromouse), 'Keyboard', 'Brainwave' (NeuroSky), and 'Speech' (Microsoft Speech API).

Once native events are received from a device, the Service Manager converts incoming information to Omicron input events, which are time stamped and inserted into an event queue. Omicron's architecture is illustrated in Figure 1.

Applications using Omicron as a static library can access the queue and process events directly. To stream events to other applications or across a computing cluster, an Input Server application is launched and listens for incoming clients using the Omicron connector API. The Input Server manages client connections and Service Type requests, then streams requested events to the clients in either UDP or TCP depending on the priority of an event. Single command events like a button click will be sent on the TCP stream, while continually updating events like a motion tracked 3D position will be sent via UDP.

Multi-Application MDE Use Cases

Networking applications in MDEs and providing an array of collaborative input modalities to users has proven itself as a useful paradigm for a variety of scientific exploration, analysis, and dissemination tasks. This section will cover three unique use cases where multiple separate applications are run and displayed on various computing resources with networked communication. In each case, applications can connect to Omicron's Input Server to support multi-user interaction with input devices best suited for their task (in addition to using native controls such as a mouse and keyboard). Additionally, we are beginning to use



Figure 2. A multidisciplinary team of scientists is shown interactively exploring a 2048^3 voxel micro CT scan of a primate tooth. The left half of the large-scale tiled display is showing a high-resolution image of 3×3 tiles (total image size of 5760×3240 pixels) streamed from ANL's Cooley visualization cluster. A laptop running the interaction client is mirrored on the right half of the large display so the entire group can view the visualization parameters such as the transfer function, material properties, and light positions.

the motion capture data to track users during these tasks so that we can start to understand how researchers utilize the physical space in a MDE.

Visual Exploration for Large Volumetric Data

Interactive exploration of large volumetric data presents a number of challenges that can be addressed using multiple applications in a MDE. Our architecture consists of three main elements: (i) visualization cluster and parallel rendering framework, (ii) tiled display, and (iii) interaction client. We use Cooley, the visualization cluster at Argonne National Laboratory, consisting of 126 nodes (12 CPU cores and 384 GB RAM per node). We run vI3, a framework for large-scale visualization and analysis [3], on Cooley to execute a ray casting volume renderer. We stream images from Cooley to the tiled display using high-bandwidth links. Our tiled display is an 8×4 arrangement of 1080p LCD panels, with a total resolution of 15360×4320 pixels, driven by a Linux workstation containing four GPUs. The workstation runs a streaming client to receive rendered images from Cooley. The final component is the interaction client, which runs on a laptop that communicates with vI3, controlling the position of the camera, the transfer function for volume rendering, and other parameters such as light positions and material properties. Devices such as a keyboard, mouse, or gyromouse are used for interaction.

This setup has been used for interactive exploration of a micro CT scan of a primate tooth. The process delivers a regular grid of 2048^3 voxels. In this experiment, vI3 was run on 16 Cooley GPUs, using a distributed ray casting algorithm with global illumination for soft shadows. vI3 composited images in a 3×3 arrangement, streaming the partial images to

individual monitors on the tiled display. As a result, the total image size was 5760×3240 pixels.

Figure 2 shows Argonne scientists interactively exploring the primate tooth. Among the lessons learned from this experience, it is worth mentioning: (i) soft shadows substantially helped scientists perceive depth in the visualizations; (ii) the large pixel count was crucial to identify regions of interest; (iii) mirroring the laptop user interface on the tiled display was helpful for scientists to understand and tweak visualization parameters, getting immediate feedback on the visual results; (iv) returning to a console to use a keyboard and mouse for interaction was cumbersome, whereas mobile input devices, such as a gyromouse, enable the scientists to interact from anywhere in the room; and (v) scientists quickly learned how to interact with the system, which enabled them to identify regions of interest for further modeling and numerical simulation.

Analysis of 3D Reconstructed Geometric Objects

As described in the previous sub-section, a large-scale tiled display can be coupled with an advanced visualization cluster to stream an interactive visualization session. In addition to using such a configuration for data exploration, it can also be extended to include an HMD, such as the Oculus Rift, to perform further data analysis in VR. For example, we work with a team of neuroscience researchers investigating the mapping of neurons in the brain. An electron microscope is used to capture ultrahigh resolution images of slices from a section of a mouse brain. A supercomputer is then used to segment the data into axons, dendrites, and other structures, creating individual geometric models of each, resulting in an extremely large data set.



Figure 3. On the top we see researchers using a large-scale tiled display to explore models of neurons extracted from electron microscopy images. A subset of these models could be selected for further investigation using a virtual reality head mounted display available across the room. On the bottom, while one investigator explores the data in stereo 3D, others can see a non-stereo version displayed on a large monitor. In addition, each of the researchers in the view on top wears a hat, outfitted with tracking tokens that enables us to uniquely identify and track them. This information will be used as part of a separate effort to investigate how users interact in collaborative spaces such as these.

Leveraging the visualization cluster, which has large amounts of memory and multiple GPUs, several thousand of these complex models could be rendered at interactive rates and displayed on the tiled display. While this in itself could help provide insight, investigating these structures in further detail in stereoscopic 3D could prove to be even more beneficial. For this we use an Oculus Rift. A single workstation is used to drive the VR application, which limits the number of models that can be rendered at a time. Therefore only a subset of structures can be analyzed in further detail at a time. Researchers are able to utilize the tiled display to identify and select a subset of structures to transfer to the Oculus Rift for further investigation. Currently, this process has been done offline, but networking the two applications together and using Omicron would allow researchers to use a controller to select and transfer data from the tiled display to the Oculus in real-time (Figure 3).

Collaborative Meeting for Data Dissemination

After exploration and analysis, researchers must disseminate their findings to collaborators and brainstorm next steps. In this case, we highlight the use of a large ultra high-resolution display, multi-touch screen, and multiple input devices. We have used SAGE2 [4,6], a democratic multi-user windowing environment, for displaying and interacting with content on the large ultra high-resolution display. These meetings are supplemented with a large multi-touch screen used as a digital whiteboard, as well as laptops and tablets for personal use. SAGE2 enables any collaborator to upload new material from their personal device, such as images or PDF documents, as well as reorganize content on the large ultra high-

resolution display. Collaborators can also interact with content on the SAGE2 display by using a variety of input devices enabled through the Omicron Input Server, such as a Wii remote or direct on-screen touch.

The secondary large display, used as a multi-touch digital whiteboard, fosters ad-hoc discussion by enabling any user to generate digital artifacts on-the-fly. In contrast to presenting prepared material, this form of interaction is critical to dissemination and brainstorming sessions, which are less formal and often require adaptation as the group discussion evolves. Since a digital whiteboard application is utilized over a traditional whiteboard, information can easily be recorded and archived. This facilitates data dissemination in two ways: 1) collaborators can easily revisit the manually created artifacts generated during the meeting, and 2) individuals not present at the meeting can view the ad-hoc artifacts from the digital whiteboard in addition to the prepared material shared via SAGE2. Figure 4 illustrates a sample collaborative meeting that uses a MDE for data dissemination.

Conclusion

MDEs present unique opportunities for collaboration and enhancing data exploration, analysis, and dissemination. We have presented our work on enabling scientists to leverage disparate computing and display technologies to facilitate research tasks. Omicron, our multi-input service, enables custom applications to integrate a variety of unique devices for collaborative user interaction. These input devices are leveraged by various applications to allow users to interact with data on display environments ranging from multi-touch screens to ultra high-resolution large-scale displays and immersive HMDs.



Figure 4. Collaborators gather in front of a large-scale display to share research findings and utilize a multi-touch digital whiteboard to brainstorm next steps to advance their work.

We have presented three authentic use cases that highlight how MDEs running separate, but often interconnected, applications enhance scientists' ability to explore, analyze, and disseminate their data. Future work includes analyzing motion capture data to investigate how users utilize the space within an MDE and how to further leverage additional display technologies, including augmented reality systems and mobile devices.

Acknowledgements

This work was supported in part by the DOE (contract DE-AC02-06CH11357), and the NSF (awards OCI-0943559, CNS-0959053, and CNS-0821121).

References

- [1] Andrews, C., Endert, A., and North, C. Space to Think: Large High-resolution Displays for Sensemaking. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10), (2010), ACM, New York, NY, USA, 55-64.
- [2] Geyer, F., Jetter, H., Pfeil, U., and Reiterer, H. Collaborative Sketching with Distributed Displays and Multimodal Interfaces. In ACM International Conference on Interactive Tabletops and Surfaces (ITS '10), (2010), ACM, New York, NY, USA, 259-260.
- [3] Jiang, J., Hereld, M., Insley, J., Papka, M. E., Rizzi, S., Uram, T., and Vishwanath, V. Streaming Ultra High Resolution Images to Large Tiled Display at Nearly Interactive Frame Rate with v13. In IEEE 5th Symposium on Large Data Analysis and Visualization (LDAV), (2015), 133-134.
- [4] Marrinan, T., Aurisano, J., Nishimoto, A., Bharadwaj, K., Mateevitsi, V., Renambot, L., Long, L., Johnson, A., and Leigh, J. SAGE2: A New Approach for Data Intensive Collaboration Using Scalable Resolution Shared Displays. In Proceedings of the IEEE International Conference on Collaborative Computing:

Networking, Applications and Worksharing (CollaborateCom '14), (2014), 177-186.

- [5] Plaue, C. and Stasko, J. Presence & Placement: Exploring the Benefits of Multiple Shared Displays on an Intellectual Sensemaking Task. In Proceedings of the ACM 2009 International Conference on Supporting Group Work (GROUP '09), (2009), ACM, New York, NY, USA, 179-188.
- [6] Renambot, L., Marrinan, T., Aurisano, J., Nishimoto, A., Mateevitsi, V., Bharadwaj, K., Long, L., Johnson, A., Brown, M., and Leigh, J. SAGE2: A Collaboration Portal for Scalable Resolution Displays. In Future Generation Computer Systems, (2016), Vol. 54, 296-305.
- [7] Strengert, M., Magallón, M., Weiskopf, D., Guthe, S., and Ertl, T. Hierarchical Visualization and Compression of Large Volume Datasets using GPU Clusters. In Proceedings of the 5th Eurographics conference on Parallel Graphics and Visualization (EGPGV '04), (2004), Dirk Bartz, Bruno Raffin, and Han-Wei Shen (Eds.). Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 41-48.
- [8] Sutherland, I. E. A Head-mounted Three Dimensional Display. In Proceedings of the December 9-11, 1968, Fall Joint Computer Conference, Part I (AFIPS '68 (Fall, part I)), (1968), ACM, New York, NY, USA, 757-764.
- [9] Van der Laan, M., Kellet, R., Girling, C., Senbel, M., and Su, T. A Collaborative Multi-touch, Multi-display, Urban Futures Tool. In Proceedings of the Symposium on Simulation for Architecture & Urban Design (SimAUD '13), (2013), Society for Computer Simulation International, San Diego, CA, USA, Article 10, 4 pages.
- [10] Waldner, M., Pirchheim, C., Kruijff, E., and Schmalstieg, D. Automatic Configuration of Spatially Consistent Mouse Pointer Navigation in Multi-display Environments. In Proceedings of the 15th International Conference on Intelligent User Interfaces (IUI '10), (2010), ACM, New York, NY, USA, 397-400.