

CupHolder: A Multi-Person Interactive High-Resolution Workstation

Mark Hereld*

Michael E. Papka*

Justin Binns*

Rick L. Stevens*

Argonne National Laboratory, University of Chicago, and the Computation Institute

ABSTRACT

Demand for high resolution visualization, large pixel real estate collaborative workspaces, and interactive computer interfaces continue to drive researchers to develop new physical portals connecting them to their computational tools, to their data, and to their colleagues around the world. In this paper we describe CupHolder, a high performance workstation designed to support interactive collaboration and research activities. It is configured to enable display of high resolution imagery while enabling a comfortable interactive environment for one to several co-located researchers. Moreover it is driven by a high performance commodity cluster that provides substantial local rendering muscle as well as a high performance interface to Grid-based computational tasks. CupHolder is comprised of commodity components. It is primarily novel because it represents an integration of these components into a new form factor that we believe is a useful precursor and testbed for future integrated workspaces.

CR Categories: B.4.2 Input/Output Devices – Image display, H.5.2 User Interfaces – Input devices and strategies, H.5.2 User Interfaces – Interactive styles, H.5.3 Group and Organization Interfaces—Computer-supported cooperative work, I.4.9 Image Processing and Computer Vision – Applications, C.0 General – Hardware/software interfaces

Keywords: tiled displays, human computer interface

1 INTRODUCTION

Commodity display and computation resources play an increasing role in scientific research at all levels. Many approaches to aggregating these resources have been applied to the task of creating high performance facilities in the form of large cluster computing systems and large format tiled displays.

Creating large format high-resolution display systems has a relatively long history which includes much work on projection-based tiled displays [9]. For a summary of issues and pointers to relevant research see Hereld et al. (2000) [3]. However, a path to higher pixel density systems leveraging now ubiquitous LCD flat panel display technology has become increasingly appealing [7]. Displays of this kind with as many as 100 Mpixels (LambdaVision) have been developed within the OptiPuter project [8]. A table format display prototype with a clear plastic protective sheet has also been developed by the same project, the LambdaTable. These often enormous displays are driving development of new technologies for content delivery over very high bandwidth networks [6]. The table format displays invite




Figure 1: Our StationOne concept drawing from 1998. It was a combination of functional display regions driven by groupings of computers selected for best support of the various functions comprising the workstation: graphics, control, and interface. In addition the frame was conceived as a flexible substrate for customization by the user – placement of user devices, surfaces and hooks for personal effects.

interactive styles that include gestural input, as well as the inclusion of physical objects to bridge the gap between the real and virtual [5]. Many interesting variations on these themes have been explored. The Magic Table combines projector, camera, machine-vision and physical tokens on a table surface to create a multi-person interactive whiteboard [2]. And the Escritoire is a multi-resolution projection-based personal writing table [1]. None have folded the high performance visualization display capabilities into the personal interactive workspace.

We have long been experimenting with display system and workstation form factors in search of designs that will naturally support the high performance demands of today's scientific applications (Fig. 1). In addition to high resolution, vast numbers of pixels provide a canvas for support of interactive collaborative environments enabling the side-by-side presentation of complex simulation results and video streams of distant collaborators. A progression of devices and technologies map the arc of our work to date: ActiveMural, μ Mural, and the AccessGrid. With the development of Grid infrastructure in mind, we have become keenly interested in developing a prototype interface to potentially huge computational and data resources. In particular, we seek an instance of such an interface with a form factor that is suitable for an office.

In this paper we describe a new configuration of commodity components which targets high performance display and computing while supporting collaborative work styles and enabling interactive input modalities based on computer vision techniques. In the following section we will describe the design of the CupHolder. In section 3 we will discuss applications that

*e-mail: {hereld, papka, binns, stevens}@mcs.anl.gov

have debuted on it. And in section 4 we will add observations and comments about the future of this design point.

2 CUPHOLDER DESIGN

In this section we will discuss the various aspects of the CupHolder design (Fig. 2), what motivated our choices, and how these were ultimately implemented.

Visualization Field. The main surface for display of demanding visualizations and multimedia content is provided by the 3 x 3 array of Dell UltraSharp 2001FP flat panel displays (20.1 inch, 1600 x 1200 pixels). The aggregate 4800 x 3600 pixel surface is bright and compelling. Each display panel is driven by a PC (3.6 GHz P4, 1MB cache, 2 GB ram, 750 GB disk, on-board copper gig-e network) with an nVidia GeForce 6600GT 128MB graphics card in the PCI-E x16 slot. We selected this box for its form-factor – the Shuttle XP is about 8 inches wide, 8 inches tall, and 12 inches deep – so that we could house them within the envelope of the CupHolder frame.

Console and Table. To this field of pixels we added two multi-purpose functional units. The console field is provided by 3 of the same flat panel displays at a 45 degree angle adjoining the bottom of the visualization field. The table field is provided by another 3 LCD panels lying horizontally and adjoining to the bottom of the console field (Fig.2). This substantial dollop of pixels is driven by a single Dell XPS 600 workstation (3.2 GHz Pentium D) with two dual-headed 256MB nVidia 7800 GTX adaptors in the two 16x PCI-E slots of the box and a single nVidia GeForce FX 5200 Ultra in a PCI slot. It has 4GB of ram and 1 TB of disk. These two ranks of displays are covered by a cold formed transparent polycarbonate cowling that protects the six displays from spilled liquids and damage from objects placed on it (keyboards, mice, coffee cups).

Human Factors. In designing the CupHolder we wanted a

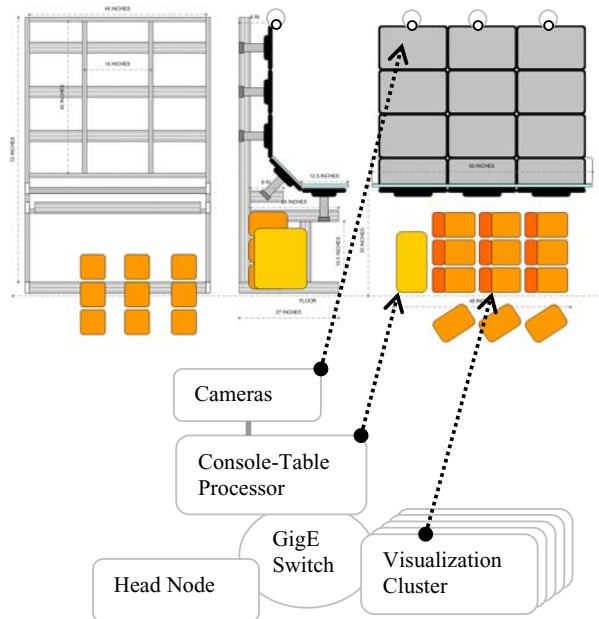


Figure 2: Diagram of the CupHolder. The 3 x 3 array of LCD flat panel displays are driven by a cluster of PCs. The console bank of 3 displays and the table bank of 3 displays are driven as a single desktop by a 6-headed PC. Three USB cameras provide overlapping coverage of the console, table, and surrounds from above.

system that would fit in an office, serve one to several people, and encourage a comfortable work style. To this end, we decided on that it should be used principally while standing, mainly to support several people. The integrated horizontal surface provides a convenient resting place for personal effects, a keyboard, and future tangible user interface devices.

Interactive Interface Testbed. The volume above the table and the space in front of the visualization field are targets for gestural input (Fig. 3). Cameras (two or three) are stationed above the CupHolder and casually positioned to take in the target volume. Overlap in the field-of-view of these provides for modest positional resolution in the vertical direction.




Figure 3: The CupHolder. The angled console and horizontal table surfaces both provide substantial pixel real estate for application control interfaces that may involve gesture recognition.

3 APPLICATION EXPERIMENTS

We have engaged the CupHolder in a number of tests to see how it performed and how it felt for a number of use cases. All of these application experiments were carried out informally at SuperComputing 2005, on its maiden voyage.

Monolithic Visualization. In this use case a single application drives the cluster, using Chromium [4] or homegrown solutions, to produce a visualization that covers the 3 x 3 array of LCDs as if they were a single panel. In this environment the user is usually interested in the large number of pixels to provide a canvas on which a large number of data points can be mapped. The mullions provide a windowed view into the visualization that users feel comfortable with, much like looking out a normal window. For many situations care needs to be taken that the mullions appear as obstacles in the view path and not simply as discontinuities to enhance the sense of realism for the user.

Distributed Simulation Monitor. Here, the application drives the panels separately, deriving rendered results from processes running on Grid computing resources around the country (Fig. 4). The central tile shows overall progress of the simulation with a composite view of the entire visualization. This mode of operation can be combined with the approach described above merging multiple panels to form a larger high-resolution view. This mode of use provides an effective way to monitor a large number of different events, be they the results of a simulation, the results of a simulation from multiple views, or using different visualization

methods for the same data. It can be used to provide a method of presenting diagnostics of how the simulation is performing or other information about the environment in which the simulation is being performed.




Figure 4: The CupHolder as a demonstration vehicle at SuperComputing. This application shows CupHolder in its capacity as high performance visualization front end to a large distributed simulation run on Grid computers at several locations around the country and one in the United Kingdom.

Cartographic Navigation. Here we begin to explore the look and feel of interacting with flat 2-D representations using hand gestures to select viewing modes and to navigate by pan and zoom (Fig. 3). For these experiments we developed simple techniques for interpreting camera images and for passing events to the pre-existing application program.

The image processing prototyping was done in Matlab. A calibration step required the user to place his hand into a zone within the camera field of view where the mean and standard deviation of hue and saturation were computed and used for future identification of skin tone in the field. The left and right images were then processed to provide 2D position of the left and right hand in camera coordinates. The left hand image was tested for proximity to virtual button positions. The right hand image was tested for significant changes of position and interpreted as “mouse” input to be used for pan and zoom operations – hand orientation was interpreted as mouse up or mouse down state.

A python interface to the Matlab code was provided using PyMat. Detection of state changes from the button interface on the left hand and position changes from the mouse interface on the right hand were passed to the application program via its xmlrpc interface.

4 DISCUSSION

Although we have only begun to apply the prototype CupHolder to the various tasks for which it was designed, it is clear that the combination of visualization real estate, integrated (spill-proof) display table and control console, cameras, and its standing configuration provide a compelling interface.

The large display surface available for parallel rendering is very appealing (bright, crisp, dense in pixels) and useful for the range of applications tested on it. Pan and zoom on huge images filling the display was compelling. Applications which provided distinct components fit naturally into the segments of the tiled display.

The console and table sectors of the workstation provide a compelling environment for program development as well. The large available real estate in the table and console display fields is useful for coding, monitoring, and googling while testing applications on the visualization field. Of all the display systems we have designed, this one elicits that greatest “*I want one!*” reaction from people who see it.

With the limited number of “flying hours” that we have accumulated on this system we can make a few tentative observations about its ergonomics.

- Standing in front of the display and discussing visualizations in groups of two to five is quite effective and comfortable.
- Facing the display and working with keyboard and mouse on the table surface is comfortable as well, but depends on the height of the table which is fixed in the current design.
- Facing the table and controlling applications with hand gestures can be taxing on the back. This probably derives mostly from the slow frame rate of the current input processing which in turn requires the user to move slowly with arms cantilevered out over the control surface.
- Working at the console while seated in a “standing chair” is likely to provide added comfort for protracted sessions.

In this paper we have described a novel integration of commodity display, PC, and camera components into a functional unit that combines high performance visualization and interactive capabilities in a multi-purpose workstation. We described several use cases. This prototype device will allow us to experiment with and develop interactive input methods, distance collaboration, and remote visualization.

As a prototype, CupHolder will allow us to evaluate this breed of high performance workstation and continue to design and develop improved systems. Some of the areas that will need attention are summarized in this list of notes.

- The single most damning aspect of this approach to high resolution large format displays is the regrettable width of the border zone on each LCD panel. Some of this frame width could no doubt be eliminated at the packaging level. But there are likely more fundamental issues that would require creative re-engineering of the panels to allow for minimal abutting losses on all four of the sides of each panel.
- Another interesting possibility is to marry this display system with a pixel-based network transport, such as Sage [6]. Such an approach would reduce the expense of the local computing fabric, and would accrue benefit from the straightforward homogeneous model for content display – pixel shipping, that is.
- Because of its concave shape the CupHolder form factor provides a natural volume for capture and interpretation of gestural input, an area that needs further development in this context.
- The flat surface of the table field, backlit with a dense array of pixels provides an interesting arena for mixed mode user interface combining aspects of tangible UI wherein objects placed on the surface and provide means for interacting with displayed information on the table.
- We intend to integrate sound more tightly with the design. This enhancement will support distance collaboration applications.
- The present physical design would benefit from the addition of more unprogrammed space in the form of hooks, racks, and surfaces for personal items and customization.

5 ACKNOWLEDGMENTS

The authors acknowledge help from the other members of the Futures Laboratory, particularly Ti Leggett, in preparing for and executing the CupHolder deployment to SuperComputing 2005. This work was supported by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Contract W-31-109-ENG-38 with additional support provided by the Pantheon Testbed project.

REFERENCES

- [1] Mark Ashdown and Peter Robinson, "The Escritoire: A personal projected display", Proceedings of the 11th International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision (WSCG 2003), Pilsen, Czech Republic, 3-7 February 2003, pages 33-40.
- [2] Bérard, F., "The Magic Table: Computer-Vision Based Augmentation of a Whiteboard for Creative Meetings," in IEEE workshop on Projector-Camera Systems (PROCAMS 2003), Nice, France.
- [3] Mark Hereld, Ivan R. Judson, Rick L. Stevens, "Tutorial: Introduction to Building Projection-based Tiled Display Systems", *IEEE Computer Graphics and Applications* 20(4): 22-28 (2000).
- [4] G. Humphreys, M. Houston, Y.-R. Ng, R. Frank, S. Ahern, P. Kirchner, and J. T. Klosowski, "Chromium: A Stream Processing Framework for Interactive Graphics on Clusters," in *Proceedings of ACM SIGGRAPH 2002*, San Antonio, pp 693-702.
- [5] Ratti, C., Wang, Y., Ishii, H., Piper, B., Frenchman, D., "Tangible User Interfaces (TUIs): A Novel Paradigm for GIS", in *Transactions in GIS*, 2004, 8(4): pp. 407-421.
- [6] Renambot, L., Rao, A., Singh, R., Jeong, B., Krishnaprasad, Naveen, Vishwanath, V., Chandrasekhar, V., Schwarz, N., Spale, A., Zhang, C., Goldman, G., Leigh, J., Johnson, A., "SAGE: the Scalable Adaptive Graphics Environment", in *Proceedings of WACE 2004*, Nice, France.
- [7] Timothy A. Sandstrom, Chris Henze, and Creon Levit, "The hyperwall", in *IEEE Proceedings of the International Conference on Coordinated and Multiple Views in Exploratory Visualization (CMV'03)*, ed. Jonathan C. Roberts, p.124.
- [8] Larry L. Smarr, Andrew A. Chien, Tom DeFanti, Jason Leigh, Philip M. Papadopoulos, "The OptIPuter," *Communications of the ACM*, Volume 46, Issue 11, November 2003, pp. 58-67.
- [9] R. Surati. "Scalable Self-Calibrating Display Technology for Seamless Large-Scale Displays", PhD thesis, Massachusetts Institute of Technology, 1999.