

Report Documentation Page

*Form Approved
OMB No. 0704-0188*

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE OCT 2012	2. REPORT TYPE	3. DATES COVERED 00-00-2012 to 00-00-2012			
4. TITLE AND SUBTITLE Poster: Building a Large Tiled-Display Cluster		5a. CONTRACT NUMBER			
		5b. GRANT NUMBER			
		5c. PROGRAM ELEMENT NUMBER			
6. AUTHOR(S)		5d. PROJECT NUMBER			
		5e. TASK NUMBER			
		5f. WORK UNIT NUMBER			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory, 4555 Overlook Ave., SW, Washington, DC, 20375		8. PERFORMING ORGANIZATION REPORT NUMBER			
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)			
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)			
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES IEEE VisWeek Workshop on Large Data Analysis and Visualization, Seattle, Washington, 14-19 Oct 2012.					
14. ABSTRACT This poster describes the design and construction of a world-class tiled-display cluster in our laboratory space. It features a 245.76 megapixel display surface (60 WQXGA monitors) with a 270 field of view. The display is powered by a 18-node rack-mount machine cluster located in an adjacent room. We will summarize the system design and provide advice for anyone interested in building a similar cluster.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 2	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Poster: Building a Large Tiled-Display Cluster

Jonathan W. Decker*

Mark A. Livingston†

ABSTRACT

This poster describes the design and construction of a world-class tiled-display cluster in our laboratory space. It features a 245.76 megapixel display surface (60 WQXGA monitors) with a 270° field of view. The display is powered by a 18-node rack-mount machine cluster located in an adjacent room. We will summarize the system design and provide advice for anyone interested in building a similar cluster.

Index Terms: I.3.2 [COMPUTER GRAPHICS]: Graphics Systems—Distributed/network graphics;

1 MOTIVATION AND DESIGN

Regardless of the resources are available to a single workstation, cluster computing remains a simple means to theoretically multiply computational throughput. In the same way, individual displays or projectors can be combined to expand the visual real estate of system. A tiled-display cluster is a combination of networked machines and multiple monitors. Tiled-display clusters allow researchers to employ these resources on a single large-scale visualization or series of visualizations. Applications include immersive virtual environments, presentation demonstrations, group collaborative views, and command-center views.

We endeavored to build the most advanced tiled-display cluster possible within a limited laboratory space. We examined existing systems at other institutions and this informed our design [5] [3]. One motivating goal was to replace a older CAVE (Cave Automatic Virtual Environment) system in a smaller space using consumer hardware. However, we were also moving away from research into virtual environments and focusing heavily on large-scale VisAnalytics. Instead of one large view, we may want to have a number of separate views of different, equally important data sets.

For these reasons, decided on high-resolution desktop monitors over larger format, thin bezel displays so that we could get reasonably close to the displays and still see new detail. We chose 30" WQXGA desktop monitors as the tiles for our display wall (WQXGA resolution is 2560x1600).

Figure 1 shows a visual summary of our design process. We choose the monitors first, developed a 3D sketch of how the monitors would fill the space, commissioned a designer to design a stand-alone support structure for the monitor wall, and placed a machine cluster opposite the monitor space to run the wall.

The end result was an 18 node cluster, with 15 nodes displaying to a 60-monitor, tiled-display matrix which spans three walls in a U-shape. The remaining nodes are used for auxiliary computation and coordination.

The wall configuration uses a free-standing monitor-support structure and completely fills the available space. The cluster itself is stored in two 42U server racks in a separate room. Bundles of fifteen-meter DVI cables run through the ceiling from the racks to the displays in cable trays. Each rack-mount workstation houses two graphics cards (Nvidia Quadro FX 5800), and each graphics

card in a display node renders two tiles in the wall, for a total of four displays per node.

We collected quotes for components to be part of the system in March 2009, but it took two years before the cluster would be completely operational. Unanticipated complexities delayed the construction and testing of the two server-racks of workstation nodes. We also needed to have additional power and cooling capacity designed and installed into the space for the wall and the cluster.




Figure 1: *Left to right, top to bottom:* A 3D sketch of the wall layout, the support structure design, a photo of the wall display Google Earth, and a floorplan of the laboratory space occupied by the monitor wall and the cluster that runs it.

2 APPLICATION

Since installation we have constructed several applications to demonstrate the potential of the system. Figure 2 shows a portion of the wall displaying one of our demo applications. Here we have a spherical panorama of the night sky created by Axel Mellinger¹ [4]. The original image is 36000 pixels by 18000 pixels. It is broken into many smaller images and displayed as a pyramid of tiles when displayed in the web application. In our demonstration of the wall, the image is loaded in to every GPU (Using 65% of the available 4GBs) and displayed as one large view using the Equalizer framework [1]. The image is mapped to the inside of a sphere, and the views from each tile are ordinated to look in the outward into that sphere. We also place labeled cursors around objects of interest, just ask Axel Mellinger has done in his web application. The end result is a highly-interactive immersive view of the night sky. A user can use a mobile device to communicate with the head node to rotate the view.

3 LESSONS LEARNED

Over the course of the design, purchase, construction, and installation of this resource, we learned several lessons which we will describe in detail.

3.1 Design with Maintenance in Mind

We constructed a tiled-display matrix which makes the most use of the available space. However, this has made maintenance incon-

¹Axel Mellinger's Milky Way Panorama can be seen on the web at: <http://galaxy.phy.cmich.edu/~axel/mwpan2/krpano/>

*e-mail: mark.livingston@nrl.navy.mil

†e-mail: jonathan.decker@nrl.navy.mil



Figure 2: A photo of the wall display a large-scale spherical panorama of the sky with marked objects

venient. We currently need to disassemble large parts of the display support structure to reach cable ports for troubleshooting. The U-shape of the display gave us one unexpected benefit for maintenance, in that the curve provides a small crawl space in the “corners” of the display that would otherwise only be accessible by dismantling the wall from the nearer end of the U-shape.

3.2 Budget for Redundancy

We only need fifteen nodes to render the display tiles, but we purchased eighteen nodes. This has allowed us to use a separate machine as a master node, and switch components to troubleshooting problems. Additionally, we recommend purchasing additional display and Ethernet cables in case of failure.

3.3 Avoid Separating Displays and Machines

We opted to rack-mount our workstations early in the design. While this provides us with easy access to the nodes, it required running long display cables through the ceiling. We suggest avoiding the installation and maintenance overhead of rack-mount cluster nodes, and instead placing them near the displays with easy access to the back, which was already recommended above. Long runs of display cables are expensive and difficult to replace in case of point failure. The best approach would be to place stand-alone nodes near the displays themselves, and to leave space to get around the display wall. Also consider that newer workstation graphics cards have alternatives to bulky DVI-D cables, such as DisplayPort and HDMI (see: Nvidia Quadro 6000). We recommend these formats because they are much easier to plug-and-play.

3.4 Leverage Open Source Frameworks

We have explored several venues of multi-head application development, including a from-scratch approach. We are currently utilizing Equalizer [1], a GLUT-like C++ framework for parallel rendering with OpenGL. We started out using Rocks Cluster 5.4 [2]. Rocks installs nodes over PXE and walks you through tile-display configuration with its “Vis Roll” add-on. However, it forces you to use older libraries which can make development difficult. It may be worth trying for a trial period, especially if you’re unfamiliar with Linux systems. It also has a number of useful Python scripts which are worth reading and implementing on your own system. We have since installed a custom cluster configuration running Ubuntu Server 10.04 LTS (Long Term Support)

3.5 Study Your Space

The free-standing structure assumes a level floor, which is not entirely true about our actual space. The floor is cement and dips a bit

at the opening and varies across the surface. As a result, the supports don’t sit very level and it was harder to align the parts of the support structure on uneven ground, which led to minor incongruousness among the monitors. While it is only barely noticeable, it could likely have been avoided if we had hired a flooring contractor to raise and level the floor of the space before we built the supports. Post-hoc leveling efforts have provided no improvement.

3.6 Final Thoughts

The bottom line is that you should make the best use of your space, but you will need to exercise restraint in your design to provide room for maintenance and air flow. If you build your machines into server-racks, you will find yourself with all the issues related to owning a server room.

Today, there are a number of companies offering turn-key solutions for tiled-display installations. These are worth comparing to the cost of designing your own system from scratch. Also worth noting, there is driver-level support in AMD Radeon FirePro graphics cards for transparently scaling to a large number of displays without middleware like Equalizer. This feature is referred to as the AMD Eyefinity Technology. Nvidia is just now catching up with driver-level support for scaling to 3+ monitors, but currently AMD shows clear advantage in this space. Avoiding the programmer overhead of designing custom tiled-display applications from the ground up will save your organization the most in the long run.

4 FUTURE WORK

We are currently engaged with demonstrating the cluster with prototypes built with the Equalizer Framework. We plan to evaluate the cluster as a replacement for CAVE applications (although it is not capable of stereoscopy), as well as collaborative presentations and command center applications. We also plan to make use of the space for the study of perceptual effects related to the field of view and spatial memory.

Given additional budget and a larger space for the tiled-display wall, we may consider removing the workstations from their rack-mount cases and placing them in stand-alone towers that could be placed near the wall so that we no longer need to go through the ceiling. We need additional space for easy access to the back of the monitor matrix for maintenance purposes. If the machines could be placed on the floor behind the wall, standard DVI-D cables could reach the monitors. While it is elegant to have the machines all stored together in two full racks, the cables must be longer to reach the monitors, even if the racks were positioned behind the wall. Although, this solution would not be perfect, since the backs of the machines will face the back of the tiled-display wall, and the nest of display cables will make it slightly awkward to access.

REFERENCES

- [1] Equalizer. <http://www.equalizergraphics.com/>, May 2012.
- [2] Rocks Cluster. <http://www.rocksclusters.org>, May 2012.
- [3] T. A. DeFanti, J. Leigh, L. Renambot, B. Jeong, A. Verlo, L. Long, M. Brown, D. J. Sandin, V. Vishwanath, Q. Liu, M. J. Katz, P. Papadopoulos, J. P. Keefe, G. R. Hidley, G. L. Dawe, I. Kaufman, B. Glogowski, K.-U. Doerr, R. Singh, J. Girado, J. P. Schulze, F. Kuester, and L. Smarr. The optiportal, a scalable visualization, storage, and computing interface device for the optiputer. *Future Generation Computer Systems*, 25(2), Feb. 2009.
- [4] A. Mellinger. A color all-sky panorama image of the milky way. *Publications of the Astronomical Society of the Pacific*, 2009.
- [5] T. Ni, G. S. Schmidt, O. G. Staadt, M. A. Livingston, R. Ball, and R. May. A survey of large high-resolution display technologies, techniques, and applications. In *Proceedings of the IEEE conference on Virtual Reality, VR '06*, pages 223–236, Washington, DC, USA, 2006. IEEE Computer Society.