INL/CON-12-24914
PREPRINT

The IQ-Wall and IQ-Station -- Harnessing Our Collective Intelligence to Realize the Potential of Ultra-Resolution and Immersive Visualization

IEEE VR 2012

Eric A. Wernert
William R. Sherman
Chris Eller
David Reagan
Patrick D. Beard
Eric T. Whiting
Patrick O'Leary

March 2012

The INL is a
U.S. Department of Energy
National Laboratory
operated by
Battelle Energy Alliance

Idaho National Laboratory

The IQ-wall and IQ-station -- harnessing our collective intelligence to realize the potential of ultra-resolution and immersive visualization

Abstract
We present a pair of open-recipe, affordably-priced, easy-to-integrate, and easy-to-use visualization systems. The IQ-wall is an ultra-resolution tiled display wall that scales up to 24 screens with a single PC. The IQ-station is a semi-immersive display system that utilizes commodity stereoscopic displays, lower cost tracking systems, and touch overlays. These systems have been designed to support a wide range of research, education, creative activities, and information presentations. They were designed to work equally well as stand-alone installations or as part of a larger distributed visualization ecosystem. We detail the hardware and software components of these systems, describe our deployments and experiences in a variety of research lab and university environments, and share our insights for effective support and community development.

1. Motivation

1.1 The Legacies of CAVERNUS and Chromium

The first virtual reality CAVEs™ were great. They showed the world how a large-format, physically-immersive virtual reality experience could more fully engage the senses and create a heightened state of perception and presence [Cruz-Neira 1992] [Brady 1995]. In an equal but different way, the first generation of ultra-resolution tiled display walls were also great. They illustrated the power of being able to see details within the context of the big picture and created opportunities to collaborate naturally around complex collections of data.

But the first generation of CAVEs was also flawed. They were expensive, complex, and fairly inaccessible to most people, requiring highly-customized and unfamiliar software. Likewise, the first generation of ultra-resolution walls was flawed. They required specialized, cluster-aware versions of software and significant effort to maintain the rendering clusters and projector alignments.

The one thing the promoters of both CAVEs and ultra-resolution walls got right was the fact that hardware standards helped to facilitate the sharing of methods, software, and data sets, and thereby boosted the adoption of these devices and helped to drive the areas forward. ("It's the content, stupid.") In particular, communities were created around the CAVE (the CAVE Research Network User Society — CAVERNUS) and the Chromium project (on SourceForge.net). Such communities allow developers to share techniques that others can use to improve their applications, and for developers and users to download applications from other locations that allow for a greater breadth of exemplar work to explore and learn from. The greater breadth of content can also be a means by which the local community can be encouraged that they too might be able to make use of these advanced displays.

1.2 The promises and pitfalls of off-the-shelf technologies

Periodically, an influx of new commodity technologies leads to a surge of innovation. At the turn of the millennium, the emergence of powerful PC graphics cards and quality, affordable DLP projectors gave rise to cluster-based tiled walls using WireGL [Humphreys 2001] and later Chromium [Humphreys 2002]. Circa 2007, we saw 3D DLP TVs and Nintendo WiiMotes lead to a surge of creativity around "home brew" virtual reality.

With any new area or technology, there is great value in having a large community of people more or less randomly exploring the space of possibilities. Collectively, it allows us to cover the search space more quickly. However, at a certain point, we want those activities to coalesce around a handful of commonly-identified areas of great promise. Otherwise, the community cannot capitalize on these efforts and the activities can seem to devolve into a contest to see whose hack gets the most hits on YouTube or Slashdot.

Over the past few years, we have experienced another explosion of new technologies, including the second generation of 3D televisions, the Microsoft Kinect™, more optical tracking options, stereo cameras, and multi-touch interfaces. This led to another explosion of independent research projects and innovative applications. However, if we fail to coalesce around some standards for hardware and underlying software toolkits, we can't readily share application software and data sets. And without significant sharing, it's difficult to learn from the past, broaden the base of adoption, or mature as a field.

1.3 Identifying your primary interest - VR researcher or VR practitioner?

Before we go any further, we feel it is essential to distinguish between two groups of people interested in virtual reality, especially the off-the-shelf variety. We define a *VR researcher* as someone whose primary interest is in developing new techniques or technologies for

VR and for whom application of VR to any specific field is simply an interesting use case. We define a *VR practitioner* as someone who is primarily interested in helping communities of users to integrate existing VR technologies (developed by VR researchers) into the workflows of those communities in sustainable ways. Naturally the field of VR needs and benefits from both types, and a given individual can certainly embrace both roles. However, we feel it is beneficial to declare one's primary motivation to avoid a mismatch of expectations.

Thus, for purposes of this paper, we declare ourselves to be VR practitioners , functioning as a bridge between current VR technologies and end users with real tasks to perform. As VR practitioners, we are also looking to extend the reach of VR technologies to new user communities inside and outside our institutions.

2. Objectives and philosophy for the IQ project

The IQ project started as an effort to create an inexpensive, interactive, immersive interface (i^4 or i-quadded = IQ). The IQ-station was developed as a low-cost alternative to a CAVE™ or ImmersaDesk™. The IQ-wall followed as a lower-cost and easier-to-maintain alternative to cluster-based ultra-resolution walls.

Given that we were coming from institutions that each already had a CAVE-style display, and that our institutions are responsible for partnering with smaller institutions within our respective states, we gave special thought to the unique role that low-cost technologies can play in the broader visualization and virtual reality landscape. To us, it is much less about saving money than it is about how lower-cost systems can open up new possibilities and serve as bridges between technologies as well as between groups of people.

The goals and guiding principles of the IQ-display initiative include:

- **lower or eliminate the barriers to using immersive technologies.** This includes not only fundamental barriers such as convenient, local access to the physical systems, but also barriers related to usability of the interface, training and support for novice users, and compatibility with software already used in workflows.
- **increase the breadth and depth of utilization.** By explicitly targeting traditional users of visualization and VR (sciences and training) as well as non-traditional communities (e.g. arts and humanities), we can achieve both depth and breadth of usage and further democratize the technology.
- **create a technological bridge between desktop and high-end systems.** We pay careful attention to the software tools and workflows that we promote, looking for those that can transition from the desktop, to the IQ-station or IQ-wall, and then on to higher-end systems like the CAVE, and even out to computational grids. In this way, the IQ-displays are fully integrated into the local and national visualization "ecosystem." They do not become a detour for the users normal workflow, but a "stepping stone" between levels of capability.
- **find a balance between low cost and professional quality**. We aim to constrain costs, but never at the expense of quality. Since we are working to convince "non-believers" to adopt this technology, experiencing a poor quality immersive display or a poor quality software application works against our goal. Thus, we frequently refer to our systems as "*lower*-cost" rather than low-cost.
- **find a balance between good performance and ease of use and maintenance**. Likewise, we make a conscious trade-off between the complexity and performance of a cluster-based solution versus the reliability and ease-of-operation of a single PC solution. We can often look to offset the performance difference by including higher-end components in the single PC. Again, our objective is not to make this another research project but a hardened and reliable recipe.
- **be open and collaborative with the design.** We explicitly sought to design the IQ-station in conjunction with collaborators at other institutions so as to avoid a parochial mindset and to get the best thinking of multiple groups. We evolve our designs to incorporate the best ideas, regardless of where they originate, and we document our recipes through our community portal (http://www.iq-station.org).
- **build communities.** Our ultimate goal is to bring together VR researchers, VR practitioners and communities of end users to share methodologies, expertise, applications, and data sets. The goal of system standards is not to limit creativity, but to enable the sharing which enables the field to mature.

3. IQ Hardware and Software Recipes

3.1 IQ-wall

The IQ-wall is an ultra-resolution tiled display wall that scales from eight up to 24 screens with a single PC. This recipe has evolved through three generations over the past three years. We detail the graphics, display, and software components below.

3.1.1. Graphics and Computation

While the display is the component that grabs most of the attention, the computation and graphics subsystem is the key to the usability and flexibility of the IQ-wall recipe. Whereas most tiled wall installations utilize a cluster of computers with a ratio of computers to displays ranging from 1:1 to 1:4, the IQ-wall utilizes a single PC to drive anywhere from eight up to 24 displays at their full resolution. The key is to utilize multiple graphics cards with multiple outputs along with consumer-level display expansion modules.

Our preferred graphics unit is the nVidia QuadroPlex which marries two nVidia Quadro cards with an SLI (scalable link interface) link into an external unit connected to a host interface card in the PC. nVidia provides driver-level support for “mosaic mode” which gangs together the GPUs and accompanying displays into a single logical GPU/display which is presented to the operating system and applications. Each Quadro card housed by a single QuadroPlex has two outputs and together provide support for four displays. We extend this multiple display capability by adding a Matrox graphics expansion module (DualHead2Go™ or TripleHead2Go™) to each display output of the QuadroPlex. The Matrox devices query the attached monitors and then present the graphics unit with an EDID signal that identifies a single logical display that is two or three times wider than each actual display. In this way, we can expand the four outputs of the QuadroPlex to either eight or 12 displays, each running at their full resolution. Particular computer configurations and motherboards are able to support an additional QuadroPlex, bringing the current maximum number of supported displays up to 24. In more recent drivers, nVidia has also enabled “mosaic mode” for two internal cards, obviating the need for an external QuadroPlex for systems up to 12 displays.

There are a few minor considerations and constraints that one must be aware of when using this single system configuration. To begin, all monitors must be identical in resolution and refresh rate, though not necessarily the same size or manufacturer. In addition, Matrox devices only support side-by-side configurations (not above-below), so it does constrain horizontal layouts to multiples of two or three. Finally, while the Matrox extension modules can be used with most modern graphics cards, designers must pay attention to the overall performance of the graphics subsystem and the increased load on the raster memory and the rasterization operations.

Finally, the computer system should be scaled proportionately with the graphics subsystem. We typically use systems with dual quad- or hex-core CPUs, between 24 and 48 gigabytes of RAM, and plenty of fast disk storage, sometimes employing a PCI-based solid state drive for better performance on large data sets. We have successfully configured the nVidia and Matrox devices under both Windows (WindowsXP and Windows7 64-bit) and Linux.

3.1.2. Displays

In theory, any array of matched display devices can constitute the visual display of an IQ-wall. Projector arrays, popular in early ultra-resolution display wall installations, are less ideal because of space requirements for the projection path, the difficulty with edge-matching and blending, and the considerable alignment maintenance required. Moreover, while blending is supported by nVidia drivers, it is not supported in the Matrox units; therefore, it is necessary to resort to edge-matching which will lead to discontinuities in brightness. One can use arrays of common desktop monitors or consumer level televisions as an affordable or available option (sometimes salvaging screens from surplus when a teaching lab is upgraded.) However the size of the bezels does impact the user’s experience.

For these reasons, we prefer the use of “pro-sumer” level ultra-thin bezel monitors designed for video walls. (We have utilized the UT and UE series of displays from Samsung with great success, but other manufactures make similar products which should work as well.) Not only do these thin-bezel displays eliminate the need for continual realignment and nearly eliminate the discontinuities of bezels, but they also offer control software for synchronized power on/off, timers, and source switching as well as easier individual color balancing and image controls. Some video wall screens also provide a built-in scaling functionality where a single video source can be looped through all of the displays and scaled accordingly. This is ideal for presentations or collaborative sessions using a separate laptop or videoconferencing system.

Finally, there are a variety of mounting options for the displays, ranging from desktop stands to floor stand to flat and articulating wall mounts. We favor those floor and wall mounts that allow the display to be curved towards the user as this makes it easier to see all pixels and adds to the sense of immersion.

3.1.3. Software Options

As stated earlier, the real magic of the IQ-wall recipe comes in the use of a single PC to drive the entire display. This means that almost any application that runs under Windows or Linux can scale to the entire screen resolution. (To date, the only major scalability issues are with Flash applications which have a limited canvas size.) This significantly expands the range of uses and users while significantly reducing the staff effort necessary to support those applications. Use cases include not only scientific, information, and GIS visualization and immersive VR environments, but also collaborative work sessions using common office documents, large group presentations and teleconferencing, interactive art pieces, ultra-resolution media playback, and interactive digital signage.

3.1.4. Representative Costs

Configuration & Use	*Details*	*Approximate Cost*
Desktop 3 x 2 individual user's office or small lab	19-23" 1080p screens; 6-monitor stand; PC with one dual-output graphics card; two TripleHead2Gos, cables, etc.	$7,000
Free-standing 3 x 3 larger lab, conference room, teaming or other informal gathering spaces	46" Samsung UT460 displays with floor stands; PC with QuadroPlex and 3 Matrox TripleHead2Gos, cables and control PC	$45,000
Wall-mounted 4 x 6 large public venues and gathering spaces, galleries, lobbies, or presentation rooms	55" Samsung UE55A displays with extending arm mounts; PC with dual QuadroPlexes and 8 Matrox boxes; video extenders and other cables; control PC	$110,000

3.2 IQ-station

The design of the IQ-station continues to evolve from our initial prototypes [Sherman 2010]. Since tthen we continue to make design improvements and leverage new consumer and "pro-sumer" technology advances that drive down the cost while driving up the usability, capability, and supportability of the core design.

3.2.1. Original IQ-station

The first iteration of the IQ-station integrated a stereo-capable DLP television (typically Samsung or Mitsubishi) with a 6-camera optical tracking system (OptiTrack™ from NaturalPoint), a pair of PCs (a higher-end visualization PC running Windows or Linux and a smaller PC running Windows to support the NaturalPoint TrackingTools™ software), wireless game controllers (Wingman-style or WiiMotes), and an ergonomic stand. Approximate hardware costs ranged from $9,000 to $25,000 depending on features and PC capabilities.

3.2.2. Recent IQ-station enhancements

Building on our experiences with and observed usage limitations of the first generation prototypes, we concluded that great benefits could be obtained from the addition of a second large screen positioned below the primary screen and angled toward the user. The first benefit emanates simply from the increased real estate. This screen provides a more natural location for widgets and controls, freeing the upper screen to focus on the visualization representations. Because the lower screen is primarily for user inputs, we found that including touch capabilities on it enhances the overall IQ-station experience. In our second generation prototypes the lower screen was monoscopic with integrated touch; for the third ("reference") generation we are using stereoscopic screens with an add-on touch system. Occasionally, the system is configured with the virtual world spanning lower and upper screens, so two stereo-capable screens is preferred — although we did find that the mono-stereo discontinuity was tolerable to view.

Another enhancement is a calibration-free tracking system. This is essential for research labs where the research is not about immersive technologies, and where a top goal is robustness and ease of maintenance. Camera-tracking calibration works against both. Yet camera-based tracking tends to be more cost effective. Fortunately, in 2011, two commercial products were released that alleviate the need for calibration. This is done by linking two or three cameras into a precise arrangement by rigidly affixing them into a metal bar. Using the same TrackingTools software, and internal camera designs, NatrualPoint offers the Duo and Trio models with two and three cameras respectively. The company A.R.T. offers SMARTTRACK with two wide-angle cameras. The NaturalPoint tracking bars are less expensive than their six independent camera package, where as the SMARTTRACK is slightly more expensive. Side-by-side comparisons of the tracking quality strongly favors the more expensive SMARTTRACK. Thus for our reference build, this is the system of choice.

One means of lowering the overall cost of the system is to substitute an MS-Windows virtual machine for the physical machine that runs the tracking software. Since the two NaturalPoint options require MS-Windows for operation, and we wish to preserve Linux as an option for the visualization computer, we can run their TrackingTools software in the virtual machine. The downside is that this consumes resources from the visualization computer, and some virtual machine options limit the computational resources of the virtual machine, but the system does function. On the other hand, we can recoup some of the extra expense of the SMARTTRACK system from the fact that it can connect directly to the visualization PC running either MS-Windows or Linux without the need for an additional computer.

Consumer stereoscopic televisions are now available both in active and passive glasses styles. There are clear trade-offs between the two. In particular, active stereo is better because it can provide higher-resolution, whereas passive stereo is better because the glasses are cheaper and do not consume batteries, and there are no synchronization issues across displays. We are presently conducting side-by-side comparisons of the two technologies. We anticipate that for systems deployed in end-user environments, the passive systems will be sufficient, and also far more robust.

3.2.3. Software Options

Immersive visualization software options include a variety of open source and freely-available tools. Domain- or data-type-specific tools include: GPU-accelerated volume rendering and LiDAR visualization tools develop by some of the authors, other VRUI-based tools developed at University of California - Davis [Kreylos 2008], and Visual Molecular Dynamics (VMD) software from the University of Illinois [Stone 2001; 2010]. In addition, a team from INL and IU have been collaborating with Kitware, Inc. to add immersive display capabilities into ParaView [Ahrens 2005], Kitware's scalable, distributed, general purpose visualization tool built on top of their Visualization Toolkit (VTK) library.

Virtual reality and simulation software options included open source development using FreeVR [Sherman 2004] or VRUI [Kreylos 2008], both with options for developing in OpenGL or OpenSceneGraph, or higher-level content development and scripting using the open standard X3D (we prefer the freely available but closed source tools from InstantReality) or VirTools (closed source and commercial.)

3.2.4 Representative Costs

Configuration and Uses	*Details*	*Approximate Cost*
Minimal System student or hobbyist project, 3D video and media kiosk	one 3DTV & glasses, Kinect or other game tracker, $2K computer, basic table and mounting hardware	$5,000 (or less if using re-purposed technology)
Standard System lab system for VR and immersive visualization, development seat for higher-end systems	one larger 3DTV & glasses, mid-level optical tracking system (e.g. OptiTrack Trio), $5K computer with smaller secondary monitor, good table and mounting hardware	$12,000
High-end System shared system for large-data visualization or complex visual simulations	two larger 3DTVs & glasses, bottom screen with touch, high-end all-in-one tracking system (e.g. ART SMARTTRACK), $10K computer, highly adjustable table and mounting hardware	$25,000

4. Our Deployments and Observations

We see a primary benefit of low-cost immersive technologies being the ability to deploy immersive systems directly in labs, classrooms, and studios where scientists, educators, and artists work. Thus, our collaboration team has been pushing IQ-stations and IQ-walls out into the world beyond the computer science research lab.

4.1 IQ-walls

Indiana University (IU) has had a strong response to the prototype IQ-wall and now has a total of five systems deployed across the IU system. Sizes and applications vary between locations and include two 3x3 installations in academic buildings and research centers for visualization and presentation work, a 6x2 installation at the GlobalNOC for displaying a variety of 2D and 3D network monitoring tools, and a 6x4 system in the commons space of new IT building to support presentations, digital art, GIS, and collaborations on programming codes and office docs. Idaho National Laboratory (INL) has a 3x3 system used in support of ultra-resolution visualization and simulation. When not in active use, all systems are able to run our open Web-based digital signage solution based on HTML5 and PHP scripts.

4.2 IQ-stations

INL collaborates with all the major universities in Idaho and thus has made IQ-stations available to many through strategic deployment. The systems at Boise State University, University of Idaho, and Idaho State University are used for a variety of teaching and research purposes in the areas of visualization, computer science, virtual museum repositories, VR simulation, and creative interfaces. The IQ-station at INL's Center for Advanced Energy Studies visualization lab provides a development seat for researchers preparing to using the CAVE, as well as a highly capable station in its own right when the CAVE is busy. IU has deployed its system for uses ranging from visualization to interactive art to a stereoscopic video kiosk. Both institutions have successfully shipped IQ-stations and IQ-walls to national conferences and exhibits.

4.3 Observations

Innovations driven by user needs and creativity. One observation that portends to the great pent-up demand for immersive visualization is how quickly, almost instantly, new end users come up with innovative use cases and feature requests. The ability to interact with scientific data in an immersive visualization format tends to trigger great ideas in the minds of end users and significantly widens the realm of possible analysis and investigation. We have observed several cases where end users have started with some baseline capabilities and expanded or modified the tools as well as their data to do some very interesting work. Viewing data in it's natural 3D representation where the viewer becomes one with the data, tends to expand the view of the end-user and to fundamentally change how they think about their work.

Learning curve of immersive visualization technology. In the two years since the first IQ-stations were deployed, much has been observed regarding the challenges faced by new end-users. In most cases, their first introduction to the technology is both extremely impressive and overwhelmingly complex at the same time. End users have found the most success when they are willing to invest time and effort to learn the basics of the technology and connect with experts who can assist as needed. INL sponsored an immersive visualization workshop in July 2011 where experts conducted training and consulting in a format that brought the regional immersive visualization community together and gave encouragement and real results to all who attended. The most highly utilized INL-sponsored IQ-stations are currently utilizing a combination of local staff and students who have received initial training either through INL internships or other collaborative relationships with experts in the field.

5. Conclusion

The IQ-wall and IQ-station have proven to be valuable, reliable, lower-cost, open recipes for ultra-resolution and immersive visualization. They have been successfully deployed in a variety of environments and are being applied to a wide range of academic, creative, and research pursuits. Continued long-term results will depend on how well the VR community — comprised of both VR researchers and VR practitioners — can coherently adopt hardware and software standards as well as create a collaborative dialog that enables sharing of technologies, techniques, and applications.

References

[Ahrens 2005] James Ahrens, Berk Geveci and Charles Law. "ParaView: An End-User Tool for Large Data Visualization," In *The Visualization Handbook*. Edited by C.D. Hansen and C.R. Johnson. Elsevier. 2005.

[Bohrer 2008] Gil Bohrer; Longo, M.; David Zielinski; Rachael Brady, "VR Visualisation as an Interdisciplinary Collaborative Data Exploration Tool for Large Eddy Simulations of Biosphere-Atmosphere Interactions," *Springer Verlag Lecture Notes in Computer Science (LNCS)*, Vol. 5358, Proceedings of the International Symposium on Visual Computing, pp 856-866, 2008.

[Brady 1995] Rachael Brady, John Pixton, George Baxter, Patrick Moran, Clint S. Potter, Bridget Carragher, Andrew Belmont, "Crumbs: a virtual environment tracking tool for biological imaging ," *Biomedical Visualization, 1995. Proceedings*, vol., no., pp.18-25, 82, 30 Oct-3 Nov 1995.

[Cruz-Neira 1992] Carolina Cruz-Neira, Daniel Sandin, Thomas DeFanti, Robert Kenyon, John Hart, "The CAVE Audio Visual Experience/Automatic Virtual Environment," *Communications of the ACM*, vol 35, no 6, pp 65-72, 1992.

[Humphreys 2001] Greg Humphreys, Matthew Eldridge, Ian Buck, Gordan Stoll, Matthew Everett, Pat Hanrahan, "WireGL: a scalable graphics system for clusters," *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, 2001, pages 129-140

[Humphreys 2002] Greg Humphreys, Mike Houston, Ren Ng, Randall Frank, Sean Ahern, Peter D. Kirchner, and James T. Klosowski, "Chromium: a stream-processing framework for interactive rendering on clusters," *ACM Transactions on Graphics*, volume 21, number 3, 2002, pages 693-702

[Kreylos 2008] Oliver Kreylos, "Environment-Independent VR Development," Springer Verlag Lecture Notes in Computer Science (LNCS) Vol. 5358, *Proceedings of the International Symposium on Visual Computing*, p. 901 ff., 2008.

[Sherman 2004] William R. Sherman, SIGGRAPH 2004 Course Notes, Course 18: "Commodity-Based Projection VR: Software for Virtual Reality", August 9, 2004, Los Angeles, CA.

[Sherman 2010] William R. Sherman, Patrick O'Leary, Eric T. Whiting, Shane Grover, Eric A. Wernert, "IQ-Station: a low cost portable immersive environment," *Proceedings of the 6th annual International Symposium on Advances in Visual Computing* (ISVC'10). pp. 361–372, 2010.

[Stone 2001] John Stone, Justin Gullingsrud, Paul Grayson, and Klaus Schulten. A system for interactive molecular dynamics simulation. In John F. Hughes and Carlo H. Séquin, editors, *2001 ACM Symposium on Interactive 3D Graphics*, pp. 191-194, New York, 2001. ACM SIGGRAPH.

[Stone 2010] John Stone, A. Kohlmeyer, K. Vandivort, and Klaus Schulten, "Immersive molecular visualization and interactive modeling with commodity hardware" *Proceedings of the 6th annual International Symposium on Visual Computing* (ISVC '10), Springer, pp 382-393, 2010.

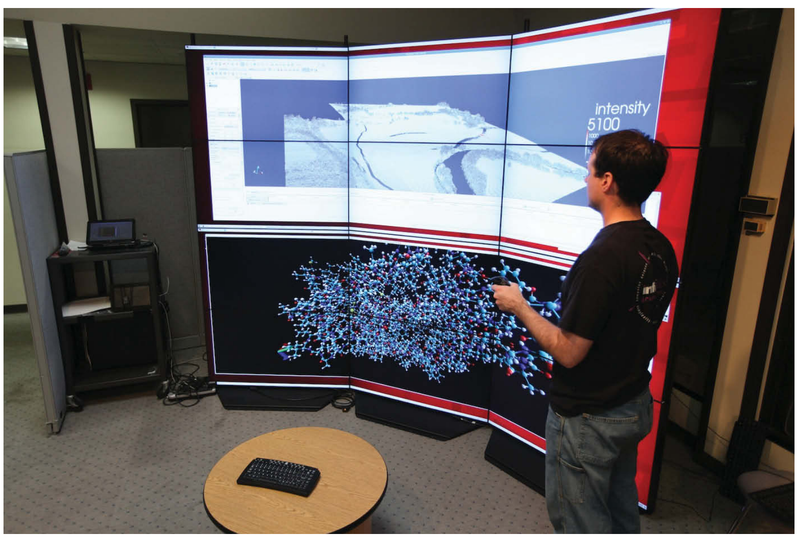

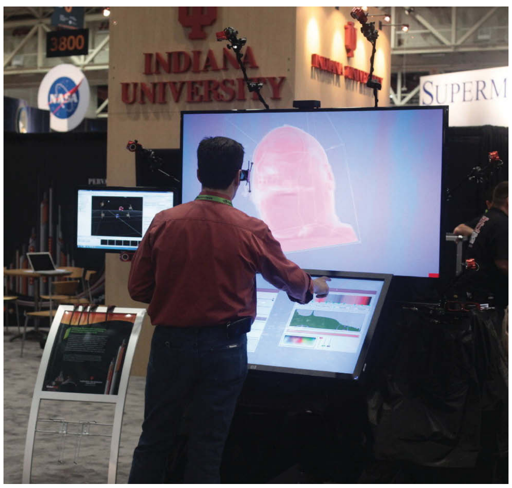
3800
INDIANA