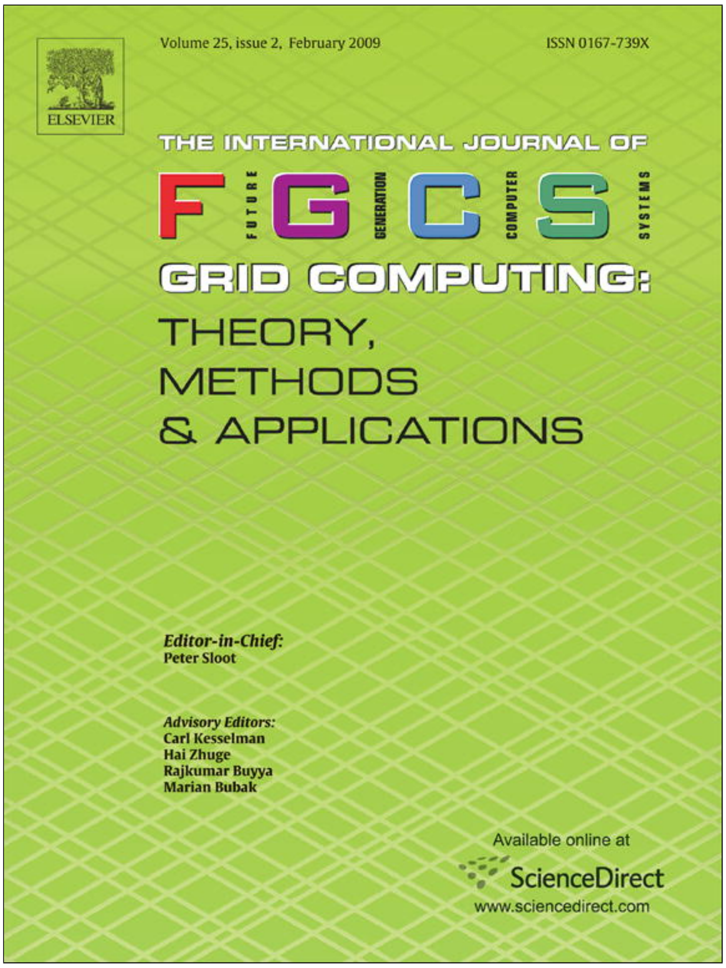
Volume 25, issue 2, February 2009
ISSN 0167-739X
ELSEVIER
THE INTERNATIONAL JOURNAL OF
F FUTURE G GENERATION C COMPUTER S SYSTEMS
GRID COMPUTING:
THEORY,
METHODS
& APPLICATIONS
Editor-in-Chief:
Peter Sloot
Advisory Editors:
Carl Kesselman
Hai Zhuge
Rajkumar Buyya
Marian Bubak
Available online at
ScienceDirect
www.sciencedirect.com

Future Generation Computer Systems 25 (2009) 114–123

Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

The OptIPortal, a scalable visualization, storage, and computing interface device for the OptiPuter

Thomas A. DeFanti [a,*], Jason Leigh [b], Luc Renambot [b], Byungil Jeong [b], Alan Verlo [b], Lance Long [b], Maxine Brown [b], Daniel J. Sandin [b], Venkatram Vishwanath [b], Qian Liu [a], Mason J. Katz [a], Philip Papadopoulos [a], Joseph P. Keefe [a], Gregory R. Hidley [a], Gregory L. Dawe [a], Ian Kaufman [a], Bryan Glogowski [a], Kai-Uwe Doerr [a], Rajvikram Singh [a], Javier Girado [c], Jurgen P. Schulze [a], Falko Kuester [a], Larry Smarr [a]

[a] California Institute for Telecommunications and Information Technology (Calit2), University of California San Diego (UCSD), United States

[b] Electronic Visualization Laboratory (EVL), University of Illinois at Chicago (UIC), United States

[c] Qualcomm, Inc., United States

ARTICLE INFO

Article history:
Received 10 March 2008
Received in revised form
15 June 2008
Accepted 26 June 2008
Available online 3 August 2008

ABSTRACT

The OptIPortal is a tiled display that is the visual interface to the OptIPuter, a global-scale computing system tied together by tens of gigabits of networking. The main point of the OptIPuter project is to examine a "future" in which networking is not a bottleneck to local, regional, national and international computing. OptIPortals are designed to allow collaborative sharing over 1-10 gigabit/second networks of extremely high-resolution graphic output, as well as video streams. OptIPortals typically consist of an array of 4 to 70 LCD display panels (either 2-megapixel or 4-megapixel each), driven by an appropriately sized cluster of PCs, with optimized graphics processors and network interface cards. Rather than exist as one-of-a-kind laboratory prototypes, OptIPortals are designed to be openly and widely replicated, balancing the state of the art of PCs, graphic processing, networks, servers, software, middleware, and user interfaces, and installed in the context of a laboratory or office conference room. Discussed in detail are the design decisions made to achieve a replicable tiled display that can be built by computational science researchers in various disciplines.

1. Introduction

Both display devices and the graphics engines that drive them are currently limited to approximately 4 megapixels (MP) on a single display device. In order to achieve displays of higher resolution, one creates an array of displays, that is, a tiled display. We call our collection of tiled displays "OptIPortals" because they are the visual interface devices for the OptIPuter, a global-scale computer tied together by optical networks using the Internet protocol (IP) www.optiputer.net. Here, we will mainly discuss the hardware, software, and networking issues of tiled displays, leaving discussions of parallel computing and storage aside.

Tiled displays are not new, of course. A large wall of standard TV monitors with means to replicate or enlarge the image over the expanse of monitors has been a feature of museum, public fair, and trade exhibitions for decades. The CAVE [1] is a tiled display in the shape of a 3 m^2 room with 3-4 tiles as walls from the time it was first shown in public in 1992, and it used then much the same configuration of a head node, with several attached workstations for the graphics as we do in tiled displays today. Planar tiled projected displays for high resolution computer graphics were frequently built in the 1990s en.wikipedia.org/wiki/Powerwall, for instance, at Princeton www.cs.princeton.edu/omnimedia/photos.html and Argonne National Laboratory http://www-fp.mcs.anl.gov/fl/activemural/. These devices were all prototypes, one-of-a-kind and, due to cost and size, not designed or supported for wide propagation as laboratory/office work environments. The OptIPuter project, as a computing and communication effort, has put considerable time and effort into designing the OptIPortal as a *replicable* ultra-high resolution display, a compute and storage device for personal and group use, with capacity an order of magnitude or two

* Corresponding author. Tel.: +1 312 996 3002; fax: +1 312 413 7585.
E-mail address: tdefanti@ucsd.edu (T.A. DeFanti).

0167-739X/$ – see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.future.2008.06.016

Fig. 1. As of June 2007, there are 38 OptIPortals installed: above are 12 of the early ones.

more than a PC. The OptIPortal is a (relatively) easy-to-copy[1] and commodity-priced[2] tens-to-hundreds of megapixels (MP) display that is also a parallel cluster computer and host to storage as needed. There are many versions of the OptIPortal http://www.optiputer.net/optiportal/ (Figs. 1 and 2).

A notable early LCD tiled display that was developed in August 2002 is the NASA Ames Research Center Hyperwall (http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1215010) and shown as a 7 × 7 array at the SC02 conference in November 2002. The Electronic Visualization Laboratory at the University of Illinois at Chicago also started building tiled displays in the summer of 2002, and brought a 4-tile display to the iGrid2002 conference in Amsterdam in September 2002. EVL, in 2005, built a 100-MP LambdaVision tiled display, and the 35-panel Varrier [2], a 3-D autostereo version, both constructed using custom-built PCs[3] and NEC monitors.[4] A group at UC Irvine built a 50-tile, 200-MP wall in 2005 out of 30″ Apple Cinema displays and Apple Power Mac G5s. Tiled displays were adopted by the OptIPuter project, and redesigned by the authors to use mass-produced PCs driving 24″ and 30″ displays. This paper's technical content is about the design decisions made to achieve a replicable tiled display that can be built by researchers for researchers in various disciplines, with as little

Fig. 2. Genomics researchers Ginger Armbrust and Terry Gasterland examine data on a 225-MP OptIPortal at Calit2.

[1] The OptIPortal is not a product because there is no recognized mass market of, say $100,000,000/year, to drive commercialization. Custom display fabricators need to charge about 3 times the parts price for a supported, well-integrated product, a cost increment our colleagues at universities will not bear. Instead, the OptIPortal is built by using a collection of specification sheets, software downloads, and wiki documentation. The PCs, graphics cards, network interface cards, HDTV input and sound input/output cards, and software all change as new versions become available.

[2] OptIPortals, with LCD tiles, PCs, graphics, local networking, some disk and framing to hold it all run about US$1,000/megapixel. Thus, a 40-megapixel display is about $40,000. One can reduce the cost down to about $600/megapixel by optimizing pixels over other system costs (like graphics power, disk space, RAM, CPU power, and networking). The OptIPuter project software is open source.

[3] Up until about a year ago, in order to get enough PC-card slots, a big enough power supply, and the latest in motherboard technology, we had to specify our own custom PC motherboard, CPU, etc., and get an integrator to build the system. The resulting systems required time-consuming debugging, and were not easy to recommend to our user community, since the components' availability changed too frequently to achieve stability. Fortunately, PCs with motherboards with enough slots have come along that provide a mass-produced alternative. Major manufacturers change models often too, but the debugging phase is (mostly) done by the manufacturer.

[4] Monitors, like projectors, tend to be available for a relatively short time before they are replaced by a newer model that has a different size, brightness, pixel count/pitch, etc. This means that one should buy several extra displays to use for spare parts, because it is likely that exact replacements will not be available a year hence. Displays are difficult to get repaired. Even warranties do not help—a newer model display that no longer fits physically or visually will be sent to replace it.

assistance from us as possible. We will discuss the choice of display tiles, PCs, networking, and software in the next sections.

Since Rocks www.rocksclusters.org is the software environment upon which the OptIPortal is based, the hardware requirements for the OptIPortal are essentially those for Rocks, once the choice of display is made. Most of the deployments of OptIPortals have been done on commodity hardware, running Intel or AMD processors. Configurations are possible, in which each computer in the cluster can drive one, two or more displays, depending on the performance and capabilities of the chosen graphics interface. OptIPortals can be optimized for specific functionality in terms of processor speed, network bandwidth, storage capacity, memory availability, and cost.

OptIPortals whose primary application include *streaming media and interactive collaboration* are generally built on bandwidth balanced platforms, in which each display has general access to 50% or more of a Gigabit network interface throughput. Such systems typically employ multi-core processors, but need to use the fastest processors available. The goal here is for a balanced bisection between communications speed and processor bandwidth. Computers in these systems computers generally drive one or two 2-MP displays each. It is important to examine where to optimize cost when building OptIPortals, since every saving is magnified 10 to 50 times.

On the other hand, OptIPortals optimized for *maximum pixel count per dollar* are usually driven by fewer computers (that is, each computer may drive 4-8 displays). These PCs are generally based on higher speed multi-core processors with 4 gigabytes of memory each and ideally 10 GigE network interfaces to allow enough streaming data input. These PCs must also have enough slots to hold multiple graphics processing unit (GPU) cards.

2. Choice of the OptIPortal display tiles

The difference between an OptIPortal and a normal OptIPuter PC cluster[5] is mainly that OptIPortals combine compute nodes with display nodes that drive pixels on physically connected display tiles. An OptIPuter PC cluster can be made of 1U rack-mounted PCs; an OptIPortal cannot be, since the graphics cards will not fit in such 36 mm-thick boxes, among other reasons noted below. OptIPortal rendering and display functionality can either be bundled within the same node, or separated and the communication facilitated by the OptIPuter networking architecture. From a user perspective, the most visual part of the OptIPortal is the array of display tiles. Several different options can be considered (in 2008).

- LCD displays for desktop computers are the least expensive, about \$250–\$350 per MP. They have long lives (at least 3 years when operating 24/7), are fairly well color matched out of the box, and require no alignment once put in a frame to hold them. They are made with bezels (frames around the panels) that create mullions (that is, horizontal and vertical strips where there is no visible image) when they are arrayed, which most people find initially distracting but get used to when data is the primary object of display.[6] We use either 4-MP displays (2560 × 1600, 30″ diagonal), or approximately 2-MPs (1920 × 1200 or 1080, 24″ diagonal). LCD displays are lightweight, which matters when designing a framing system to mount them.

Fig. 3. Rapid evaluation of massive, multi-dimensional sea cliff erosion data from: Olsen, M.J., Johnstone, E., Ashford, S.A., Driscoll, N., Young, A.P., Hsieh, T.J., and Kuester, F., "Rapid Response to Seacliff Erosion in San Diego County, California using Terrestrial LIDAR,"ASCE Solutions to Coastal Disasters Conference Proceedings, ASCE, pp. 573–583.

  - 2 MP displays (∼\$700 per tile in 2008) have the advantage that they match the streaming capacity of a modern PC with a gigabit Ethernet (GE) connection (2 MPs × 16 bits/pixel × 24 fps). They can also be driven by very modest and inexpensive graphics cards in the smallest desktop PCs. HDTV LCDs of all sizes have become very inexpensive as well, but tend to have large bezels that are very difficult (but possible) to remove.
  - 4-MP displays are slightly cheaper per pixel (∼\$ 1K per tile in 2008), and have half the mullions per megapixel, of course. They ideally need 10 GE connections for streaming. To date, these 4 MP configurations are used either for local super-resolution 2D images or 3D generated graphics displays (Fig. 3).
  - 8 MP, also called 4K, displays are really four 2 MP displays seamed together 2 × 2. If they were not disproportionately expensive (∼\$50,000 in 2008), and thick-bezeled, they would be preferable to either the 2 MP or 4 MP panels.
- Plasma TVs are bigger (50″ diagonal and larger) and heavier (50 kilos or more) usually have very thick bezels, thus producing large mullions, and are lower resolution (typically 1366 × 768, for instance, although some are higher and some are lower resolution), and cost about \$2000 per MP or more.
- Video projectors have the benefit that the mullions can be made negligible by careful alignment or blending. There are no 4-MP projectors; 8-MP projectors are about \$12,000 per-MP. These projectors (made by Sony and JVC) are driven by four 1920 × 1200 or 1080 inputs since there are no 8-MP direct drive graphics cards as of mid 2008. 4-MP (2560 × 1600) inputs are not accepted by projectors currently sold. 2-MP (HDTV) projectors that can tile well are about \$6000 per MP (and up, depending on light output). Cheaper projectors in the order of \$1000/MP or so can be used, but such projectors have low-cost plastic lenses, which makes edge matching difficult due to non-linear lens distortions. Projectors have the benefit that stereo is relatively easy to achieve. Projectors, however have several problems when used for tiled displays:
  - Projector lamps are expensive (hundreds to thousands of dollars each) and have typical lives of 2000 hours (83 days if left on 24 hours a day). Lamps age unevenly, requiring constant re-calibration and it is labor-intensive to change lamps and recalibrate on a regular basis. Some manufacturers of commercial projected tiled displays offer auto calibration at additional cost.

[5] OptIPuter clusters are similar to any other racked cluster, except for the addition of a second 1GE or 10GE network interface card to all the PCs, and the head node that connects to the private OptIPuter network, a global infrastructure of switched "virtual local area networks."

[6] We do *not* suggest using tiled displays with mullions to enlarge images for videoconferencing or movie watching, both of which are currently at most HDTV (1920 × 1080) resolution and better seen on single displays without mullions. We recommend displaying either many images or super-resolution images that map well to the pixel real estate, and take advantage of the pixel density, rather than simply exploit the overall size of the display.

- Projectors, due to various manufacturing factors, come out with noticeably different color shifts in each projector's output. This can often be mitigated with suitable color balancing on site, if the projector allows it.
- Projectors need projection distance either behind or in front of the screen. If in front, people occlude the projector beam when they get close to the screen. If behind, the shorter the projection distance, the wider angle the lens, and the more off-axis viewing creates luminance shifts in each tile, something hard to correct for unless the viewer is tracked, and then it works well only for one viewer.
- Framing and support for the projectors can be complex and alignment difficult to achieve and maintain over time. If some of the pixels (∼10%) are devoted to edge blending, many of the problems of projecting tiles can be mitigated [3] at the cost of sacrificed pixels.

The StarCAVE is a 34-MP-per-eye stereo projection-based OptIPortal (see StarCAVE paper in this issue of FGCS). It suffers from the problems above (expense, alignment) and surely is not a device for the typical lab or office.

Future technology, such as organic LED displays may provide large, cheap, mullion-less displays. Meanwhile, LCDs and some OLEDs with fashionably narrower bezels are starting to appear on the market.

3. OptIPortal graphics cards

The choice of graphics card, or GPU, is determined by the ultimate use of the OptIPortal and, in turn, puts constraints and demands on the PCs chosen as a host. As noted above, OptIPortals can be cost- and throughput-optimized for either streaming pixels over networks, or for local 3D graphics generation. Streaming can be achieved with one very modest half-height graphics card per tile,[7] which means small, quiet, and low-cost desktop PCs can be used. These half-height graphics cards will not drive the 30″ 4 MP displays, however.

For better 3D graphics, a good choice in 2008 is a board that uses the Nvidia GeForce 8600, GeForce 8800/9800/9900 or Quadro 5600 GPU. These can drive up to two 30″ displays each, and since two cards can fit in a big PC like the Dell 720, a single PC can put out graphics up to 16 MP worth of display. Such a configuration is good for 2D non-streaming graphics and 3D local graphics generation (the latter of which is largely done on the GPU and is not constrained by networking and PC memory speeds as is streaming). However, such a PC needs to have a 1 kW power supply mainly to supply the dual multi-hundred watt GPU boards.

4. Choice of OptIPortal computers

The streaming OptIPortal uses one PC per tile, as mentioned above. This configuration has the benefit that the PC fits completely behind the tile (Fig. 4), which means the computers are not visible to the users, and the cabling distance is minimal. These small PCs are also relatively inaudible, an issue when 10–20 of them are in a working environment. Similarly, one criterion for the bigger PCs that drive one or more graphics cards is that they be relatively quiet, that is, they have large fans in roomy boxes. Hence, one builds OptIPortals out of deskside/desktop machines made for office environments, not rack-mounted PCs built for server rooms, since the latter are optimized for space and have objectionably noisy fans. Of course, one option available to some, is to put the cluster in another room, ideally a server room, but rarely is there a server room conveniently located. The normal DVI cables that connect PCs to their displays are limited to about 30 m in length, so node placement in a remote server room is only possible by using DVI optical fiber extenders, which doubles the cost of the OptIPortal. Optical extension was done with great success with the 100-MP LambdaVision display at EVL since custom-fabricated rack-mounted noisy PCs were (and still are) used.

A good and quite powerful compromise between the smallest and largest PCs is the class of modest desktop machines that can take one 8600//9800 GPU and drive 2 2-MP or 2 4-MP displays in a relatively small, quiet package. As an alternative architecture, there is also a portable 15 2-MP OptIPortal driven by Apple Mac Minis simply cable-tied to the back of the displays. The same display configuration driven by PCs is called the *OptIPortable*. The minime/OptIPortable is housed in a shipping case ideal for sending to conferences for booth exhibits; its screens fold along the vertical axis, and then are pushed down into its box. Once the shipping case top is latched, it is ready to roll into the elevator and off to the shipping company. Deployment at the conference site is fast.

On the extreme end of the OptIPortal spectrum to date is HIPerSpace, a 70 tile, 286-MP display wall using a quad tile per node (16 MP) configuration to maximize pixel real-estate. A 4-MP tile with its associated computer uses about 300–400 W, half for the display and half for the computer, something to be noted when installing a 70-tile wall. 25 kW represents an operational cost of about $2.50/hr in the US (removing the heat generated is an additional cost, not to be ignored).

5. OptIPortal networking

The main point of the OptIPuter project is to examine a "future" in which networking is not a bottleneck to local, regional, national and international computing, an exercise that allows one to discover the other bottlenecks (in buses, network interfaces and protocols, disk drives, etc.). One justification for such high-speed networking is data transfer, disk-to-disk, although to date, in scientific supercomputing, this capability is not well exploited. The majority of data on the OptIPortals is from sensors or scientific instruments (microscopes, satellites, deployed sensor networks with metadata, genome sequencers, SD, HD and 4K video), and these sources are expected to increase exponentially. Most of the labs that have built OptIPortals have geoscience, planetary science, or biomedical image-related data to explore, although some applications in the arts and engineering are emerging. A Gigabit Ethernet (GE) switch is the minimum requirement for inter-nodal communication. We have successfully used a variety of commodity switches for the cluster communication, including those from Cisco, Force10, SMC and NetGear.

The typical OptIPortal PC has two GE network interface cards (NICs), one that is used for normal networking and control, and a second one that is devoted to data transfer on the OptIPuter network, a private network that exists both locally and among widely distributed OptIPuter institutions with dedicated 1–20 GE bandwidth. A 20-tile (40-MP) OptIPortal can saturate two 10 GE links fairly well, when focused on uncompressed streaming.

For an increment of about $1000–$2000 per PC, one can install a 10GE NIC, which, will allow streaming of up to ten times as many pixels per PC, a topic of active research, since it facilitates real-time streaming to the 4-MP displays. High-end PCs in 2008 can handle about 7–9 Gb/s internally [4], which is about 10 times more than their disk drives can deliver. However, optimized disk systems can deliver nearly 10 Gb/s to the switch, and thereby to any PC with a 10 GE NIC; tests to confirm the actual bits delivered are needed.

[7] We have standardized on Nvidia GPUs because Nvidia uses the same driver across its product line, which greatly simplifies the automatic setup and maintenance of OptIPortals as provided by UCSD's Rocks software.

Fig. 4. Small computers can fit behind the display tile, simplifying cabling and hiding the computers.

Currently, A 10 GE NIC is connected to the processor, using PCI-Express 1.0 interconnect over 8 dedicated parallel lanes. This results in 16 Gbps (2 Gbps per lane) dedicated bandwidth between the NIC and the processor. Additionally, GPUs are connected to the processor using PCI-Express 1.0 interconnect over 16 dedicated lanes, thus, resulting in 32 Gbps dedicated bandwidth between them. However, 10 Gbps throughput between the GPU and NIC is still a challenge. A typical transfer from a GPU to the NIC requires 4 memory copies. Thus, we have 8 memory copies (4 copies each at the sender and the receiver) in the screen-to-screen transfer between two OptIPortals. The primary bottleneck in achieving 10 Gbps screen-to-screen throughput between two OptIPortals is thus the available PC memory bandwidth.

A potential solution to sustain 10 Gbps end-to-end graphics streaming between two OptIPortal is the MultiRail Approach [7]. It is a user-space library, consisting of creating parallel "rails" through every aspect of an end-system: from processing on the multiple cores, generation of multiple application data flows, efficiently distributing memory over the memory banks, and streaming over multiple-lanes, multi-wavelength NICs connected via a parallel interconnect. The MultiRail approach would enable a data transfer between the GPU and NIC over an optimized path with lower latency, lower memory contention and higher memory bandwidth.

One could also mitigate the above bottlenecks by using kernel-based streaming and Remote Direct Memory Access (RDMA) www.rdmaconsortium.org/home to reduce the data copies associated with streaming data from the GPU to the NIC. These solutions are complementary to the MultiRail approach. However, these solutions are difficult to deploy [8]. One could also to build a customized OptIPortal node using specialized FPGAs to optimize the data transfer.

Of course, if one wants to stream large numbers of 4-MP tiles, for example, multiple dedicated 10 GE network streams are needed from the source to the OptIPortal. The Calit2 StarCAVE currently has 5 10 GE links into its Myrinet switch, and 10 GE connections to each of its 18 PCs. Some sites (UIC/EVL, Northwestern University/iCAIR, UCSD/Calit2, University of Washington/Research Channel) can, with substantial engineering effort (in 2008), link up wide area 3 10 GE links. Single 10 GE links have become routinely available to several major cities for OptIPuter use. In any event, it is desirable to continue to investigate extreme streaming: a gigapixel display refreshing uncompressed at 30 Hz would require a terabit of networking, that is, 100 10 GE links or 10 100 GE links.

Compression, of course, is a desirable and indeed ultimately necessary mode of operation, especially if it can be done in loss-less fashion, as is possible with various schemes, and easily done with computer graphics (as opposed to sensor-based data) by sending geometry rather than rendered pixels. OptIPortals have enough compute power in their CPUs and GPUs to allow decompression without special additional hardware, especially if a lossy compression is acceptable.

DXT, for example, is a lossy compression technique for texture in 3D graphics. It was designed to reduce the size of textures in video games when video memory was limited. Its main advantages are a fixed compression ratio of 6:1, a wide spread native support on all modern graphics cards, and reasonable quality. DXT compression works as follows: it converts a 4×4 block of RGB pixels into two 16-bit colors and a lookup table of 16×2-bit color indices. In addition to the two selected colors (a, b), two intermediate colors are derived (1/3a + 2/3b, 2/3a + 1/3b), hence the 2-bit color index giving the choice of four distinct colors. So, the initial data of 48 bytes (16 pixels of 3 bytes each) can be represented by 8 bytes. The six-fold data reduction enables lower bandwidth requirements for high-resolution content:

– DXT bandwidth for HD video (1920×1080) is 250 Mb/s at 30 frames per second (fps), reasonable for modest PC disk serving
– DXT bandwidth for 4K digital cinema (3840×2160) is 800 Mb/s at 24 fps, too fast for PC disks today, but achievable with networked connections to optimized servers.

Up until recently, DXT compression could achieve acceptable quality, but was considered too slow, and was mostly used as an off-line process. In an innovative paper, van Waveren re-visited the compression algorithm and proposed an assembly-language implementation using Intel multimedia instructions cache-www.intel.com/cd/00/00/32/43/324337_324337.pdf. However, the implementation described in his paper is not portable. We decided to re-implement the compression functions using C/C++ intrinsic functions mapped to native instructions by the compiler. This allows the code to be ported on wide variety of platforms, in 32-bit or 64-bit mode, while being equivalent to the low-level non-portable assembly code. We can compress HD at 100 frames per second and 4K at 25 frames per second. Higher rates can be achieved by using a multi-threaded implementation with an image-space decomposition on a multi-core processor. DXT compression is now supported by the SAGE environment (see below) as a native pixel format and a SAGE application can stream DXT content to support high-resolution at a moderate throughput [5].

6. HDTV videoconferencing on OptIPortals

The incorporation of videoconferencing in OptIPortals has been an area of considerable focus in the OptIPuter project. Rather than be the point of the tiled display usage, like the Access Grid http://www-fp.mcs.anl.gov/fl/publications-electronic-files/ag-immersive-821.pdf, videoconferencing is an adjunct to the display of high-resolution data. 1280×720 30 fps videoconferencing is now routinely achieved, using commercial devices like LifeSize and Polycom, which use aggressive relatively low-bandwidth compression schemes (as low as 1 Mb/s) that are understandably optimized for human heads and low latency, needed for good communication between people. They support point-to-point and multi-point connections for about $11,000 per site. These systems also incorporate good echo canceling in their audio, a very important feature. It is trivial to insert the unit's outputs (local camera and remote camera) into two of the multi-input displays in an OptIPortal (or provide separate panels, of course), and achieve rather good and instant "high-definition" videoconferencing through conventional network channels. These units also operate at lower resolution to work with legacy H.323 systems.

However, if one tries to communicate detailed data by pointing the LifeSize/Polycom camera at a chart or Powerpoint slide on a screen, the compression is sufficiently extreme that the resulting image is unacceptably lacking in detail, and cannot be read by viewers at the receiving site. Motion is not handled well; quickly moving subjects or feeding a DVI signal of moving computer graphics into these codecs overloads them so they update about 5 fps. Instead, streaming of image data that is not already in digital format can be done with a "prosumer" camera and sent uncompressed at 1.5 Gb/s or compressed (at 145 Mb/s or 25 Mb/s—the effective difference is largely in latency which increases as the compression increases). If the data is already digital, then sending it digitally is far better than pointing a camera at it on a screen. We use all three approaches: LifeSize/Polycom for people's faces, body language, and voice, HDTV[8] in various transmission modalities as needed to show details on real and moving objects, and digital images sent as bit maps or graphics instructions.

The OptIPortal wiki http://www.optiputer.net/optiportal/has details on the specific cameras and video compression approaches found successful. It is important to point out that good lighting is important, and requires care to achieve in a typical laboratory or office setting; bad lighting will make the best camera put out grainy, low-contrast images. Similarly, care must be taken with the audio. High-end commercial videoconferencing systems like the Cisco Telepresence www.cisco.comachieve quality by very careful design of lighting, audio, and environment, a lesson to be noted by most OptIPortal users whose normal work environment is noisy, lit by harsh overhead fluorescent fixtures, and scenically furnished with large cardboard boxes.

7. Framing

We use extruded aluminum framing elements from 80/20 www.8020.netfor the support frame. The LCD tiles are attached using standard VESA mounts. The framing was designed such that displays can be arranged in columns (subject to room height) and an arbitrary number of columns can be connected with hinges so that the columns can be arranged into a flat, faceted or curved wall. The OptIPortal wiki has spreadsheets to generate the parts list for the 24″ and 30″ Dell panels, since it is non-trivial to get all the little parts right. Adapting to a different panel simply requires accurately measuring and ordering the horizontal size of the panel.

[8] HDTV exists in a variety of formats, and is casually used to describe any video signal that is appreciably better than 480 vertical lines at 30 fps in the US, Canada, and Japan and 25 fps in most other countries.

8. OptIPortal software

Middleware and applications leveraging OptIPortal technology can be grouped into three major categories, consisting of stream-centric techniques, parallel distributed rendering techniques and hybrid systems combining distributed real-time rendering and streaming within the same context. These in turn can scale from low-level visual content distribution approaches to high-performance parallel real-time rendering engines with multithread CPU support and GPU-based hardware acceleration.

8.1. Stream-based systems

The proxy-based DMX (Distributed Multi-head X Project) http://dmx.sourceforge.net/operates on the assumption that a single front-end X server will act as a proxy to a set of back-end X servers. Rendering requests will be accepted by the front-end server, broken down as needed and sent to the appropriate back-end server(s) via X11 library calls for actual rendering. This architecture requires that the front-end server manages/renders the visual content of all nodes in a visualization grid. DMX is therefore limited to a smaller display array and not scalable without dramatic performance penalties. Although DMX is not able to take advantage of the hardware acceleration on the rendering nodes, it renders a viable solution when less dynamic graphical content has to be distributed.

SAGE (Scalable Adaptive Graphics Environment) [6] (Fig. 5) targets especially high-resolution tiled display systems that could potentially cover all the walls in a room. It operates on the assumption that as wall sizes increase, multiple users will naturally find a need to make full use of the available resolution to display multiple visuals and interact with them at the same time. It also assumes that it is possible for any type of application, given the appropriate middleware, to send a pixel stream to the SAGE tiled display. SAGE middleware instructs each of the incoming pixel streams from an application to the correct portion of a tiled wall allowing the system to scale to any number of streams and tiles. More importantly, it allows multiple applications on multiple distributed rendering clusters to run simultaneously, and be viewed simultaneously on the tiled display, in essence, a true multi-tasking operating system for tiled displays. Anything from a parallel OpenGL http://www.opengl.org/ application to a HD/4K video stream to a remote laptop can be displayed on the tiled display as long as the pixels from their image buffers can be extracted. SAGE also features a capability called *Visualcasting* whereby dedicated clusters can be placed at high speed network access points to replicate incoming pixel streams and broadcast them to multiple tiled displays at the same time enabling users on distributed OptIPortals to look at the same visuals and therefore work collaboratively. The number of Visualcasting nodes can be adjusted to suit the anticipated number of streams. This capability has been successfully demonstrated over transoceanic links. Addition of trackers or cameras for gesture input allows for richer control and interaction (Fig. 6).

8.2. Parallel distributed rendering

Many software packages distribute visual context exclusively to multiple rendering engines in a parallel master slave or client server approach. A common shortcoming by most packages is scalability across multiple-display tiles connected to a single machine, when the combined tile resolution exceeds the supported OpenGL display context size.

Fig. 5. Experimental classes at EVL are conducted using SAGE on the 100-MP OptIPortal. Content from students' laptops are wirelessly pushed onto the screen for discussion.

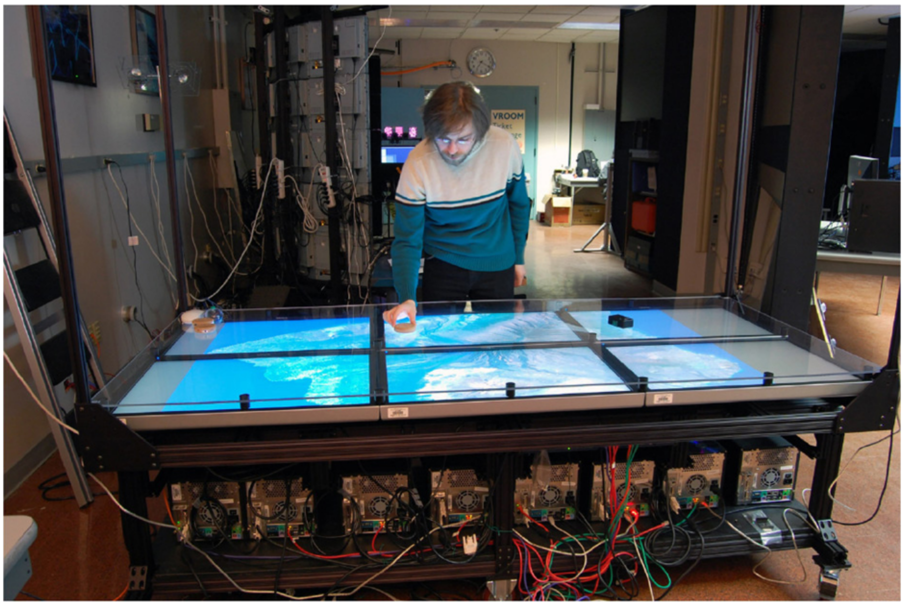

Fig. 6. The LambdaTable combines the OptIPortal resolution and SAGE software with gesture recognition using overhead cameras.

8.3. *Hybrid systems*

Chromium [9] can take advantage of the hardware acceleration on the tile nodes by using tile-sorting processes to determine which node in the cluster needs to draw which sections of the OpenGL content. Chromium splits the OpenGL commands and sends them in form of a network stream to the corresponding nodes in the cluster. Stream Processing Units on these nodes will read the received "OpenGL Streams" and pass them directly to the local graphics card on the nodes.

CGLX http://vis.ucsd.edu/ftp-site/pub/cglx/ explores an approach where high performance real-time parallel rendering and streaming of visual content from other applications can be combined. The middleware is based on the assumption that the rendering nodes in a cluster have sufficient CPU and GPU resources at their disposal. The framework can leverage from these resources, by utilizing classical work distribution strategies in cluster systems such as culling and multi-threading for OpenGL applications and provides a freely programmable API in combination with a native container-based distributed desktop management application which accepts multiple pixel streams. To maximize the availability of network resources for data transmission related to the visualization content, CGLX implements its own lightweight network layer and message passing environment. CGLX provides users with access to parallel hardware accelerated rendering on different operating systems and aims to maximize pixel output to support high resolution tiled display systems. Natively, CGLX maps an OpenGL context to each display tile, resulting in multiple contexts when multiple displays are connected per node. This attribute makes CGLX the only fully scalable OptIPortal interface currently available.

Crucial for all distributed rendering approaches is the availability of a reliable high performance network to retrieve massive data content, or to control the visualization system itself. An OptiPortal features a network solution that can provide data transfer rates up the 10 Gbits/s. These maximum values can be maintained due to dedicated high performance local networks or a high speed network grids such as OptiPuter, combined with interconnection technology such as the Myrinet (Myri-10G) http://www.myri.com/Myri-10G/overview/we use on the StarCAVE. The access to vast amounts of distributed storage and computational resources on an OptiPortal and the additional network bandwidth enables stream-based approaches to dramatically increase their achievable performance. High performance real-time parallel visualization systems, which can also act as rendering back ends for stream-based approaches, can leverage these network resources to load and process data at remote sites and to simply stream the final results at interactive rates. This attribute of OptIPortals allows users to share, exchange and manipulate remote data sets interactively in distributed cooperative workspaces spanning the globe.

While OptIPortals leverage a ROCKS-centric approach for the display cluster, middleware developed by our team scales across different operating systems, operating system flavors and heterogeneous clusters. The middleware hides OS specific aspect and provides a cross-platform API. Locally available resources, such as the number of available graphics cards, displays and associated capabilities (resolution, swap and frame synchronization, etc.) can be probed at the device driver or the window manager level, allowing the middleware to report and adapt to hardware capabilities. Considering the number of PCs in a typical OptIPortal, mean time to failure becomes an important parameter when selecting cluster management strategies. From a system administrator's perspective, ROCKS-based systems are easier to manage, largely by pruning system management overhead down to a single node. Under Mac OSX, Apple Remote Desktop and AppleScript can be used to streamline some of the system maintenance and management tasks, however in general, each OSX node needs to be manually maintained.

We have also ported core OptIPuter software to the Windows XP, 32-bit platform. A 12-tile, 27-MP OptIPortal was built using Windows XP nodes. Software ported includes SAGE 2.0. LambdaCam http://www.evl.uic.edu/cavern/lambdacam/was ported to enable real-time visual monitoring of remote tile-displays. During extensive tests at SC'07 in Reno, Nevada, when connected at 10 GE to San Diego and Oxford, UK, we observed that the gigabit networking performance of Windows XP nodes was typically 45% of that of Linux machines using the exact same hardware. Since SAGE is heavily dependent on bandwidth for raw pixel streaming, it means we can only support $\sim$half of the pixels that could be streamed to a similar Linux-based OptIPortal. Although distributing tested Windows binaries for OptIPortals makes deployment of software effortless, the lack of any automatic cluster building tools, such as ROCKS, makes it difficult to build large Windows clusters. Every node in our Windows-based cluster was manually configured; this does not scale well.

9. OptIPortal future work

In order for OptIPortals, or something like them to become ubiquitous, several greatly improved technologies are called for:

- Panels with zero-width bezels or some sort of continuous display technology that offer large-scale, print-quality resolution
- PCs or other node design that can handle 10 Gb/s and higher throughput from NIC to GPU
- Low power displays and processing (100W per MP is not "green" enough).

As predicted in 1968, http://blog.modernmechanix.com/2008/03/24/what-will-life-be-like-in-the-year-2008/, it is conceivable that our walls of the future will be covered with screens much as they are covered by paint today. Improvements in parallelism in networking, computing, graphics, and storage will be needed for these future tiled displays as well.

Acknowledgement

UCSD receives major funding from the National Science Foundation (NSF) for the OptIPuter, cooperative agreement OCI-0225642, Larry Smarr, principal investigator. UCSD also receives funding for OptIPortal and application development from the Gordon and Betty Moore Foundation for CAMERA, the Community Cyberinfrastructure for Advanced Marine Microbial Ecology Research and Analysis, and from the State of California for StarCAVE. UIC receives NSF support for OptIPuter via UCSD, as well as NSF award CNS-0420477 for LambdaVision/LambdaTable and NSF award OCE-0602117 for CoreWall. UIC also receives funding for OptIPortal and application development from the Science Museum of Minnesota, the Adler Planetarium & Astronomy Museum and the Office of Naval Research on behalf of the Technology Research, Education, and Commercialization Center (TRECC). Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the funding agencies and companies.

References

[1] C. Cruz-Neira, D. Sandin, T. DeFanti, R. Kenyon, J. Hart, The CAVE®: Audio visual experience automatic virtual environment, Communications of the ACM (1992) June.

[2] Daniel J. Sandin, Todd Margolis, Jinghua Ge, Javier Girado, Tom Peterka, Thomas A. DeFanti, The Varrier Autostereoscopic Virtual Reality Display, ACM Transactions on Graphics, in: Proceedings of ACM SIGGRAPH 2005, Los Angeles, CA, July 31 August 4, 2005, vol. 24, No. 3, pp. 894–903.

[3] Aditi Majumder, Rick Stevens, Color non-uniformity in projection based displays: Analysis and solutions IEEE Transactions on Visualization and Computer Graphics 10 (2) (2004).

[4] V. Vishwanath, J. Leigh, E. He, M.D. Brown, L. Long, L. Renambot, A. Verlo, X. Wang, T.A. DeFanti, Wide-Area experiments with LambdaStream over dedicated high-bandwidth networks, in: IEEE INFOCOM 2006.

[5] L. Renambot, B. Jeong, J. Leigh, Real-Time Compression For High-Resolution Content, Proceedings of the Access Grid Retreat 2007, Chicago, IL 05/14/2007 http://www.evl.uic.edu/files/pdf/ag2007-renambot.pdf.

[6] B. Jeong, L. Renambot, R. Jagodic, R. Singh, J. Aguilera, A. Johnson, J. Leigh, High-performance dynamic graphics streaming for scalable adaptive graphics environment, in: ACM/IEEE Supercomputing 2006 (2006) November 11–17.

[7] V. Vishwanath, T. Shimizu, M. Takizawa, K. Obana, J. Leigh, Towards Terabit/s Systems: Performance Evaluation of Multi-Rail Systems, in: IEEE INFOCOM 2007, Anchorage, Alaska, 05/08/2007 – 05/11/2007.

[8] Geoffray Patrick, A critique of RDMA, http://www.hpcwire.com/hpc/815242.html.

[9] Greg Humphreys, Mike Houston, Ren Ng, Randall Frank, Sean Ahem, Peter D. Kirchner, James T. Klosowski, Chromium, ACM Transactions on Graphics 21 (3) (2002) 693–702.

Thomas A. DeFanti, Ph.D., at the University of California, San Diego, is a research scientist at the California Institute for Telecommunications and Information Technology (Calit2). At the University of Illinois at Chicago, DeFanti is director of the Electronic Visualization Laboratory (EVL), and a distinguished professor emeritus in the department of Computer Science. He has researched computer graphics since the early 1970s. His credits include: use of EVL hardware and software for the computer animation produced for the 1977 "Star Wars" movie; contributor and co-editor of the 1987 NSF-sponsored report "Visualization in Scientific Computing;" recipient of the 1988 ACM Outstanding Contribution Award; he became an ACM Fellow in 1994.

Jason Leigh is an Associate Professor of Computer Science and director of the Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago. Leigh is a co-founder of VRCO, the GeoWall Consortium and the Global Lambda Visualization Facility. Leigh currently leads the visualization and collaboration research on the National Science Foundation's OptIPuter project, and has led EVL's Tele-Immersion research since 1995. His main area of interest is in developing collaboration technologies and techniques for supporting a wide range of applications ranging from the remote exploration of large-scale data, education and interactive entertainment.

Luc Renambot received a Ph.D. at the University of Rennes-1 (France) in 2000, conducting research on parallel rendering algorithms for illumination simulation. Then holding a Postdoctoral position at the Free University of Amsterdam, till 2002, he worked on bringing education and scientific visualization to virtual reality environments. Since 2003, he joined EVL/UIC first as a PostDoc and now as Research Assistant Professor, where his research topics include high-resolution displays, computer graphics, parallel computing, and high-speed networking.

Byungil Jeong received the BS and MS degree in electrical engineering in 1997 and 1999 from the Seoul National University, South Korea. He worked as a researcher at Imaging and Media Research Laboratory, Korea Institute of Science and Technology until 2002. He is a Ph.D. candidate in computer science and working as a research assistant at Electronic Visualization Laboratory, University of Illinois at Chicago. His research interests include scalable graphics architecture, high performance graphics streaming and tiled high-resolution displays.

Alan Verlo is an Associate Director of the Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago, responsible for EVL's high-performance computing and communications infrastructure. He is also a member of the StarLight network engineering team, and supports international, national, regional and local research and education network connections. For many years, Verlo has been a volunteer member of the Supercomputing Conference (SC) SCInet networking crew, supporting StarLight's involvement in research demonstrations. Verlo is active in the Technical Engineering Working Group of the Global Lambda Integrated Facility (GLIF), and the US National Coordinating Office's Large Scale Network (LSN) Joint Engineering Team (JET).

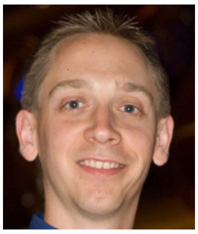

Lance Long is a Senior Research Programmer for the Electronic Visualization Laboratory at the University of Illinois at Chicago (UIC). Long eceived his MS degree in Computer Science in 2003 from UIC. Long upports and develops software/hardware solutions for UIC/EVL grant- unded research initiatives, including: a scalable 105-Megapixel tiled-display onnected to a versatile visualization cluster with multi-terabyte storage and high-bandwidth capabilities; a portable autostereoscopic display system using magnetic and optical tracking; a horizontal tiled display system with overhead optical tracking; a 4K digital cinema (CineGrid) system. Long ssists with EVL research experiments and provides technical support at major professional conferences. Along with EVL tech transfer and outreach, he collaborates with researchers in the US, Europe and Asia.

Maxine Brown is an Associate Director of the Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago (UIC). Brown is the project manager of the NSF-funded OptIPuter project, a member of the Pacific Rim Applications and Grid Middleware Assembly (PRAGMA), a founding member of Global Lambda Integrated Facility (GLIF), and is co-chair of the GLIF Research & Applications working group. She is also co-principal investigator of the US National Science Foundation's International Research Network Connections Program's TransLight/StarLight award. Brown has been active in the ACM SIGGRAPH organization as well as SIGGRAPH and ACM/IEEE Supercomputing conferences. She is a recipient of the 1990 UIC Chancellor's Academic Professional Excellence (CAPE) award; the 2001 UIC Merit Award; and the 1998 ACM SIGGRAPH Outstanding Service Award.

Daniel J. Sandin is director emeritus of the Electronic Visualization Lab (EVL) and a professor emeritus in the School of Art and Design at the University of Illinois at Chicago (UIC). Currently Sandin is a researcher at EVL at UIC and at CALIT2 part of the University of California at San Diego. Sandin's latest VR display system is Varrier, a large scale, very high resolution head tracked barrier-strip autostereoscopic display system that produces a VR immersive experience without requiring the user to wear any glasses. In its largest form it is a semi-cylindrical array of 60 LCD panels.

Venkatram Vishwanath is a Ph.D. candidate in the Electronic Visualization Laboratory (EVL) and the Department of Computer Science at the University of Illinois at Chicago (UIC). His research interests are in high performance networking, high-speed transport protocols, data Intensive computing over optical networks and petascale systems.

Qian Liu systems integrator at the California Institute for Telecommunications and Information Technology (Calit2). Prior, He worked as system engineer for National Center for Supercomputing Applications (NCSA). His research interests include cluster technology for scientific visualization in virtual environments, high performance scientific computing cluster, Super/High Definition video streaming and conference tools, Client/Server network applications in WAN/SAN and systems integration of multi-platform computing and visualization.). He received his MS degree from University of Illinois at Urbana-Champaign (UIUC).

Mason J. Katz is currently the Group Leader for Cluster Development for the San Diego Supercomputer Center (SDSC) at the University of California (UCSD). Mr. Katz received his BS in Systems Engineering from the University of Arizona. He worked for five years as an embedded software engineer on networks of lightning detection sensors. He then spent three years working at the University of Arizona on network security protocols (IPSec), and operating systems (x-kernel, Scout). He has spent the last six years working on Windows and Linux commodity clustering (HPVM, Rocks). The focus of his current work is on the Rocks Clustering Distribution, a complete software stack building high performance computing clusters. In addition, he actively involved in Pacific Rim Applications and Grid Middleware Assembly (PRAGMA).

Philip Papadopoulos received his Ph.D. in 1993 from UC Santa Barbara in Electrical Engineering. He spent 5 years at Oak Ridge National Laboratory as part of the the Parallel Virtual Machine (PVM) development team. He is currently the Program Director of Grid and Cluster Computing at the San Diego Supercomputer Center. Dr. Papadopoulos is deeply involved in key research projects including the Biomedical Informatics Research Network (BIRN), OptIPuter, the Geosciences Network (GEON), the NSF Middleware Initiative (NMI), The National Biomedical Computation Resource (NBCR), and the Pacific Rim Applications and Grid Middleware Assembly (PRAGMA). He is also well known for the development of the open source Rocks Cluster toolkit, which has installed base of 1000s of clusters. Dr. Papadopoulos is a Co-PI of the CAMERA Project.

Joseph P. Keefe is a Research Project Manager for the California Institute for Telecommunications and Information Technology (Calit2) at the University of California, San Diego (UCSD). He is the project manager assigned to the OptIPuter Project at Calit2, UCSD. He has a B.S. in Applied Mathematics from the California Polytechnic State University, San Luis Obispo, with graduate study in Applied Mathematics at the University of British Columbia in Vancouver, Canada and at the University of Washington in Seattle, WA. His career at UCSD, includes: Scientific Programming for the Ocean Engineering Research Group (OERG) at the Scripps Institution of Oceanography (SIO); Resarch Computing and Systems Management for various departments and research units at UCSD, including the UCSD Cancer Center, Department of Mathematics and the Department of Chemistry and Biochemistry.

Gregory R. Hidley, Ph.D., at the University of California, San Diego, is a technical lead for a number of Cyberinfrastructure projects including CAMERA (http://camera.calit2.net) and OptIPuter (http://www.optiputer.net). He was the first CIO of the UCSD Division of CALIT2 (http://www.calit.net) and has participatied in campus IT infrastructure planning and implementation for UCSD and for the UC system for the past 25 years.

Gregory L. Dawe, MFA, is a Principal Development Engineer at the California Institute for Telecommunications and Information Technology (Calit2). Prior, at the University of Illinois at Chicago, Dawe was Manager of System Services at the Electronic Visualization Laboratory (EVL). He holds a BA in Design from the University of Illinois at Chicago and a MFA from the School of the Art Institute, Chicago.

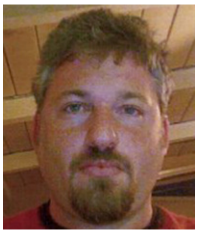

Ian Kaufman is a Research Systems Administrator at the Jacobs School of Engineering at the University of California, San Diego. Prior to that, he was a Computer Systems Engineer at Lawrence Berkeley National Laboratory's National Energy Research Scientific Computer Center (NERSC). His interests include cluster technology, visualization, audio/video streaming, security, multi-platform systems integration and wireless communications. He received his BA in Theatre/Acting from the University of California, San Diego.

Bryan Glogowski graduated with a BS from the University of California San Diego in Ecology, Behavior and Evolution while working at the Center for Reproduction of Endangered Species where he was studying Giant Panda behavior. His independent research project, conducted at the Scripps Institute of Oceanography focused on the influence of size as it relates to gender roles in the hermaphroditic bubble snail, Bulla gouldiana.

Subsequent to this, he attended Penn State University where he received a professional certification in UNIX/C/C++ and Object Oriented Programming. After completing the program, he worked as a second tier Sun Microsystems support engineer, and became a certified Solaris systems and network administrator.

After returning to San Diego, he worked at CERFnet supporting production UNIX systems. After being acquired by AT&T he was promoted to AT&T Labs while attending San Diego State University where he studied Computer Science. He later left AT&T Labs and went to Sony Online Entertainment where he was involved in supporting the global production infrastructure for their massively multiplayer online role-playing games such as Starwars Galaxies, Planetside, and Everquest II.

Bryan now works as a research systems administrator for the Office of Engineering Computing at the UCSD Jacobs School of Engineering.

Kai-Uwe Doerr received his Ph.D. degree from the Darmstadt University of Technology, Germany, in 2004. His expertise includes virtual cockpit simulation, virtual prototyping, computer vision and 3D database generation. Currently he is a project scientist working at the California Institute forTelecommunications and Information Technology (Calit2) at the University of California, San Diego. His work focuses on image-based tracking algorithms, cluster-based large scale data visualization and human factors research for interactive 3D visualization technologies.

Rajvikram Singh, is a computer scientist with the National Center for Microscopy and Imaging Research at the University of California San Diego. He holds an MS from the University of Illinois at Chicago and Bachelors in Electrical Engineering from Mumbai University, India. His research interests include video streaming, scalable graphics systems, distributed computing and high-speed networking.

Javier Girado, Ph.D., is a Staff Engineer at Qualcomm (Graphics team). He earned his MS degree from the Buenos Aires Institute of Technology (ITBA), held a research fellowship at the Industrial Technology National Institute (INTI), Argentina. He taught at the National Technology University (UTN) and the ITBA. He completed his Ph.D. from the University of Illinois at Chicago (UIC) in 2004. He worked as a Postdoctoral Researcher at the California Institute for Telecommunications and Information Technology (Calit2), University of San Diego, California until 2007. His research interests include virtual realty (VR), auto-stereoscopic displays, computer vision, and neural networks. He specializes in camera-based face detection and recognition to support real-time tracking systems for VR environments, and video conferencing over high-speed networks.

Jurgen P. Schulze, Ph.D., is a Project Scientist at the California Institute for Telecommunications and Information Technology in San Diego, California. His research interests include scientific visualization in virtual environments, human-computer interaction, real-time volume rendering, and graphics algorithms on programmable graphics hardware. He holds an MS from the University of Massachusetts and a Ph.D. from the University of Stuttgart, Germany.

Falko Kuester, is the Calit2 Professor for Visualization and Virtual Reality and an Associate Professor in the Department of Structural Engineering at the Jacobs School of Engineering at UCSD. He received his MS degree in Mechanical Engineering in 1994 and MS degree in Computer Science and Engineering in 1995 from the University of Michigan, Ann Arbor, and the Ph.D. from the University of California, Davis, in 2001. His research is aime at creating intuitive, collaborative digital workspaces, providing engineers and scientists with a means to intuitively explore and analyze complex, higher-dimensional data. In support of this research, he is developing new methods for the acquisition, compression, streaming, synchronization and visualization of data, including the ultra-high resolution HIPerWall and HIPerSpace visualization environments.

Larry Smarr is the founding director of the California Institute for Telecommunications and Information Technology and Harry E. Gruber professor in the Jacobs School's Department of Computer Science and Engineering at UCSD. Smarr received his Ph.D. from the University of Texas at Austin and conducted observational, theoretical, and computational based astrophysical sciences research for fifteen years before becoming the founding director of the National Center for Supercomputing Applications (1985) and the National Computational Science Alliance (1997). He is a member of the National Academy of Engineering and a Fellow of the American Physical Society and the American Academy of Arts and Sciences. Smarr is Principal Investigator on the NSF OptIPuter LambdaGrid project and is Co-PI on the NSF LOOKING ocean observatory prototype.