Chapter 3

DIGI-Vis: Distributed Interactive Geospatial Information Visualization

3.1 Introduction

Geospatial information systems provide an abundance of constantly updating information for researchers and scientists. Sites such as NASA's OnEarth from the Jet Propulsion Laboratory produce more than 80 GB per day of satellite data [Ple08], accessible to the public via Web Map Service (WMS) [wms]. Unfortunately this type of data can usually only be analyzed a few megapixels at a time, giving researchers a very narrow view into these voluminous data sets.

This Chapter describes a novel system for exploring ultra-high resolution geospatial information on a tiled display environment via a distributed architecture. This work is focused for gathering data through tiled WMS servers, but the proposed methods also work via other transport mechanisms. This system allows scientists to view real-time geospatial information at unprecedented levels expediting analysis, interrogation, and discovery.

A distributed data gathering and visualization system that allows researchers to view these data at hundreds of megapixels simultaneously is presented as shown

Figure 3.1: Researcher viewing the Daily Planet data set from NASA's Jet Propulsion Laboratory on a high resolution display environment.

in Figure 3.1. High resolution displays have been shown to produce much greater level of effectiveness in the analysis of larger data sets compared to that of a pan and zoom environment [BN05]. This may be due in part to the much closer mapping between the resolution of the display system and the resolution of the eye. Furthermore, this massive display real estate empowers multiple researchers to all view the data simultaneously, each creating his or her own pans and zooms through their physical location.

Traditional techniques for tiled display systems generally rely on a streaming approach. For applications such as Chromium [HEB+01], geometry is streamed from a central node to the tiled display environment. For other programs, such as SAGE [JJR+05] [RJJ+06] [JRJ+06], pixel buffers from a central node are streamed to a tiled display environment. Both of these paradigms are unfruitful for dealing with massive dynamically changing data visualization. Applications such as JuxtaView [KVV+04] and Magic Carpet [SLJM08] allow for scalable visualization, but require data to be preprocessed, dismissing the possibility of real-time analysis.

In the presented approach, each node is not only a passive renderer, but is also a data acquisition client. Each node uses its current display viewport to determine what sections of data to request via the Internet. Since the data fetching is done asynchronously and can be done via several open streams for each node, data can be acquired in a massively parallel fashion. This allows for rapid creation and visualization of ultra high resolution image layers. In practice, a 300 Megapixel image can be fully constructed and visualized in a few seconds.

By using data tiling, data can also be stored in manageable segments. This allows sections to be reused, meaning that for small changes, such as translating the viewport, minimal data fetching is required. In turn, the data only need to be fully reconstructed when zoom levels are significantly altered or when the viewport is translated substantially. Otherwise, the system is generally able to predicatively gather the data before it is needed on screen allowing for seamless panning of these real-time immense data sets.

3.2 Data

Two different sources of data were selected for this system, one which is preprocessed and local to the system and a second which is fetched via the network. The localized data provide a reference for geographical location and time-varying differentiation.

3.2.1 Local Data

Because the data pulled from the Web Map Service (WMS) [wms] servers may take some time to acquire, it was important to have a useful reference for the human researcher. The Blue Marble Next Generation [SVS+05] data was selected as it provided a very high resolution backdrop.

Local image layers are first broken up into tiles, a method also used by others [FAJ07] [KUDC07]. By tiling the image, sections of the image can be laid out sequentially in memory allowing sub-sections of the image to be subsequently loaded without massive cache penalties. Image tiling also allows for pre-generation of tile hierarchies, containing different resolutions of the same image.

3.2.2 Remote Data

Remote data could acquired in a variety of ways, but the Web Map Service (WMS) was selected as it provided an efficient client server model. WMS was developed by the Open Geospatial Consortium in 1999 [wms]. The idea of this service was to use a server which could feed georeferenced imagery to clients over the Internet. In this way, geographic information system (GIS) data could be stored in a database which can be quickly addressed from external clients.

The original WMS request structure had the client specify a geographical region of interest along with a set of parameters for a given data layer. The WMS server would then pull these data from a database and return an image for the given request. While this setup minimized the amount of data which needed to be sent via the network, the bottleneck was often on the server end as it had to render a new image for every request.

This problem was only exacerbated when many clients where connected simultaneously. While resultant images could be cached on the server side, this only proved useful for when the queries were an exact match to a previous generated image. As clients had no information as to what queries were currently in the cache, this proved of little use.

Using other models of dealing with large image layers [FAJ07] [KUDC07], many WMS servers which managed a multitude of connections used tiling to reduce the processing required on the server side. Each data layer was segmented into blocks of a fixed size and were prerendered on the server. By giving clients knowledge of how to reference these data tiles, clients could easily compute the address of these tiles. By using a simple caching scheme, the server was able to short-circuit the processing and simply feed the prerendered image back to the client. If the image had not been prerendered, the data request would still be honored, but would require the server to generate the image as before. This allowed the system to be backwards compatible, although non-cached data requests may take substantially more time.

One excellent source of WMS data comes from the NASA's Jet Propulsion Laboratory [Ple08]. This data source presents many interesting and very high resolution data layers which the system can tap into. The most often used is the Daily Planet layer which gives views of the earth which are no older than 24 hours.

3.3 Resource Management

The resource management system is responsible for making sure data is effectively and efficiently read from disk or network, uploaded to the graphics card, and visualized. While the system has to deal with local and remote data, both use the same basic structure with loading threads, data queues, and GPU texture pools as shown in Figure 3.2. The only real difference between these two methods of data acquisition comes from the formating of the data and the speed at which the data could be fetched.

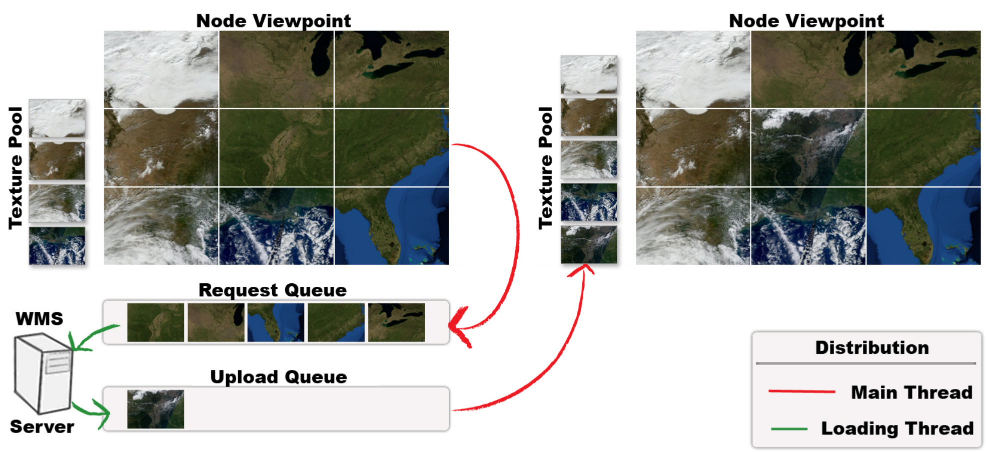

Figure 3.2: Demonstration of the data flow of the system. Data needed to be loaded are first put in the Request Queue. These data are then fetched by the Loading Thread through the WMS server. Once prepared, they are put in the Upload Queue. Between draw cycles the Main Thread can upload these data to the Texture Pool.

3.3.1 Loading Thread

To parallelize the data loading, a second loading thread was created in addition to the main thread. This thread's job is simply to pull data from the Internet and prepare it for the main thread to use. As online requests are extremely variant in the time required to gather information, this thread proved remarkably advantageous as it could work in the background.

Unfortunately, because the loading thread has no OpenGL context, the data can not be uploaded to the graphics card directly. Additionally, for machines which have multiple cores, it is advantageous to have multiple loading threads running in the background simultaneously. This allows loading to continue even if one thread becomes momentarily blocked during data transfer do to network congestion.

3.3.2 Request Queue

The request queue holds a list of what tiles need to be loaded for a given viewpoint. This list is managed by the main thread. On a fixed interval, the list expunges tiles which are in view but are not in main memory or the texture pool add queue. The only other way a tile request can be added to the queue is if a network request fails. In this case, the request is pushed to back of the queue so it can be processed a later time.

3.3.3 Upload Queue

Because the only thread which has access to the graphics card is the main thread, a temporary structure is created to prepare and store data. The loading thread pulls the data from the Internet or reads image tiles from disk. These data are decompressed and stored in a raw format suitable for upload to the graphics card on main memory. Because this task is offloaded to threads, this process happens in the background.

3.3.4 Texture Pool

The texture pool is initialized to a fixed size when the program begins. This allows textures to simply be replaced as opposed to requiring memory to deleted and reallocated during runtime. Each of the texture pool spots stores a pointer to the graphics memory where the image information is stored and meta information of the position and resolution of the tile as while as information about activity.

As the main thread is the only thread which has an OpenGL context, it must be the one to move data to the graphics card. The main thread checks the state of the upload queue in between draw loops to see if something is ready to be uploaded. If there are items in the upload queue, the texture pool needs to be scanned for an open spot. Each spot is initialized as empty, so these spots are taken first. After all spots are filled, the spot which has the oldest time stamp of usage is replaced. This assures that as long as the texture pool is large enough, tiles which are currently being viewed will not be replaced.

3.3.5 Draw Loop

When rendering content to the display space, the viewpoint for each node is first checked. The edges of each node's viewspace is switched into geo-referenced locations. Subsequently, each tile in the texture pool is checked whether it fits into the viewpoint. Additionally it is confirmed that the tile represents the appropriate level of resolution for the viewport. The tiles which are drawn store a time-stamp to mark them as being currently active.

Figure 3.3: Researchers viewing data as it is loaded from remote sources. As tiles are loaded, they are faded in to avoid visual popping effects.

The local data are rendered first with network data superimposed over top of it. To mitigate the popping effect, network data tiles are faded in when they are first loaded as shown in Figure 3.3.

3.4 User Interaction

Currently the system can be run in two ways. Users can interact with the system locally, using a computer with two 30-inch monitors and using the wall as an extended display, or users can manipulate a gyroscopic mouse and interact with the wall directly.

Users can translate the image layers by clicking and dragging with the left mouse button, and can zoom the image layers by clicking and dragging with the right mouse button. Users can also blend between the networked data and the local data. Each of these techniques could be automated without need for continuous user input. One of the most effective interrogation techniques is to slowly pan through the image while flipping between the network and local data. In this way users can follow points of interest or stay in one place while analyzing a wider view.

3.5 Applications

The true potential of this system comes from the ability to construct and visualize massive amounts of data via networked sources. As stated previously, tremendous amounts of data are created daily by NASAs OnEarth from the Jet Propulsion Laboratory. It is an incredibly challenging task to analyze the data through traditional means. By presenting such a massive amount of data in such a high resolution setting, features and differences can efficiently be discerned, a useful attribute for those studying climate change.

A great example of this comes from comparing the Blue Marble Next Generation data set to the Daily Planet data acquired from the JPL WMS servers. Because the Blue Marble Next Generation data set was acquired in 2004, direct comparisons can be made between then and now. One area which was immediately noticeable for researchers came from the area surrounding the Aral Sea. As shown in Figure 3.4, many features have changed significantly in the past half decade.

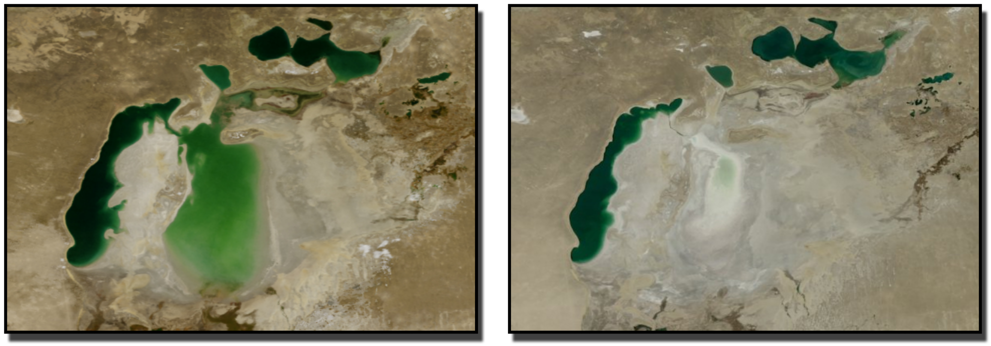

Figure 3.4: The Aral Sea as shown in Blue Marble the next Generation from 2004 on the left and from The Daily Planet on October 25th, 2009 on the right.

Applications may also exist in the field of meteorology. This ultra-high resolution visualization allows scientists to see storms develop from both a micro and macro level. Scientists can use onscreen tools to measure the size of systems and track their progress over time. The power of this system comes from the discoveries which can be made even when they were not searched for specifically.

3.6 Performance Results

Given the highly nondeterministic nature of the Internet, it is very hard to quantify performance. The goal of this system was to maintain user interactivity even while acquiring data.

The challenging scenario was tested, where a node would drive four thirty-inch monitors or sixteen megapixels. As the node may need to load a resolution one step higher than its current viewpoint, the node may need to load 32 megapixels with data from a cold start, meaning none of the data had been previously requested through the network and did not exist on local caches. The data pulled came from the NASA's Jet Propulsion Laboratory WMS server [Ple08]. Each data tile encompassed 512x512 (0.25 megapixels) worth of image information.

As shown in Figure 3.5, this process took a total of slightly more than 5 seconds per node. As 18 nodes worked in parallel during this test, the end result produced an image layer of 576 megapixels. During this entire process the frame-rate maintained a constant performance near 50 Hz.

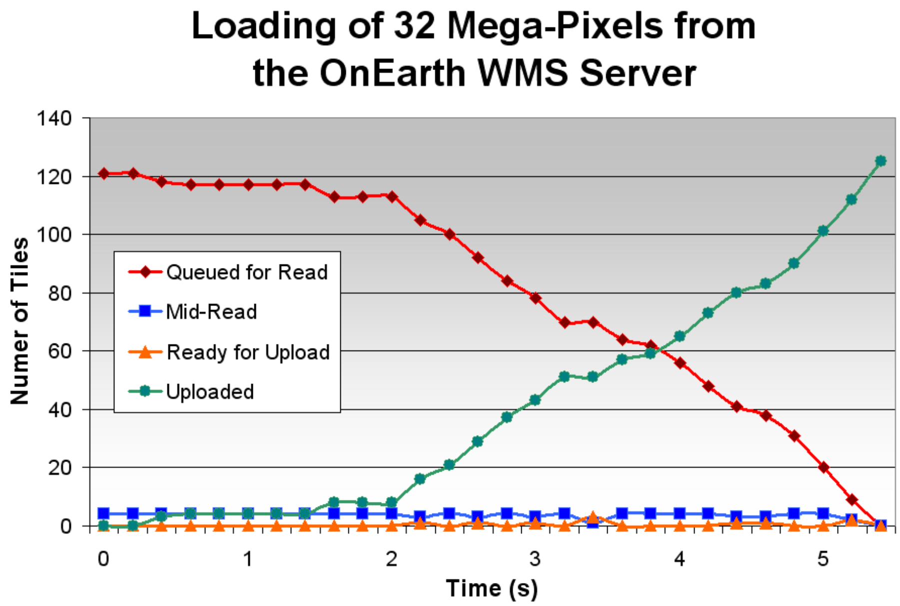

Figure 3.5: Number of tiles in each state for one of seventeen nodes acquiring 32 megapixels from the JPL WMS server.

On average, the system can fully process a tile request in 30 to 50 milliseconds. Because data are reused, slow to moderate pans can be accomplished seamlessly, as only edge data are required to be acquired.

3.7 Conclusion

This chapter presents a system for visualizing massive amounts of geospatial data from local and remote sources. This system uses a distributed approach which allows the data to gathered and visualized in parallel. Through this system, scientists are now able to view real-time geospatial information at an unprecedented resolution, expediting analysis, interrogation, and discovery. Future work will go into new interface technologies to make these systems even more accessible.

3.8 Acknowledgments

©2009 IEEE. This chapter is a reprint, with permission, from “DIGI-Vis: Distributed Interactive Geospatial Information Visualization” as it appears in the

proceedings of IEEE Aerospace Conference 2010, Ponto, K., and Kuester, F. The dissertation author was the primary investigator and author of this paper.

Bibliography

[Ake93] Kurt Akeley. Reality engine graphics. In *SIGGRAPH '93: Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, pages 109–116, New York, NY, USA, 1993. ACM.

[bbc] http://www.apple.com/quicktime/guide/hd/bbc-cfb.html.

[BJH+08] Allen Bierbaum, Christopher Just, Patrick Hartling, Kevin Meinert, Albert Baker, and Carolina Cruz-Neira. Vr juggler: a virtual platform for virtual reality application development. In *SIGGRAPH Asia '08: ACM SIGGRAPH ASIA 2008 courses*, pages 1–8, New York, NY, USA, 2008. ACM.

[BN05] Robert Ball and Chris North. Effects of tiled high-resolution display on basic visualization and navigation tasks. In *CHI '05 extended abstracts on Human factors in computing systems*, pages 1196–1199, New York, NY, USA, 2005. ACM.

[BR98] Uwe Behrens and Ralf Ratering. Adding shadows to a texture-based volume renderer. In *VVS '98: Proceedings of the 1998 IEEE symposium on Volume visualization*, pages 39–46, New York, NY, USA, 1998. ACM.

[BVG05] Stefan Bruckner, Ivan Viola, and M. Eduard Gröller. Volumeshop: interactive direct volume illustration. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Sketches*, page 60, New York, NY, USA, 2005. ACM.

[BXH+09] Leonardo Bonanni, Xiao Xiao, Matthew Hockenberry, Praveen Subramani, Hiroshi Ishii, Maurizio Seracini, and Jurgen Schulze. Wetpaint: scraping through multi-layered images. In *CHI '09: Proceedings of the 27th international conference on Human factors in computing systems*, pages 571–574, New York, NY, USA, 2009. ACM.

[car] http://www.apple.com/trailers/disney/cars/.

[CB04] Xiang Cao and Ravin Balakrishnan. Visionwand: interaction techniques for large displays using a passive wand tracked in 3d. In *SIGGRAPH '04: ACM SIGGRAPH 2004 Papers*, pages 729–729, New York, NY, USA, 2004. ACM.

[CCF94] Brian Cabral, Nancy Cam, and Jim Foran. Accelerated volume rendering and tomographic reconstruction using texture mapping hardware. In *VVS '94: Proceedings of the 1994 symposium on Volume visualization*, pages 91–98, New York, NY, USA, 1994. ACM.

[CEM01] F. Capani, M.H. Ellisman, and M.E. Martone. Filamentous actin is concentrated in specific subpopulations of neuronal and glial structures in rat central nervous system. *Brain Research*, 923(1-2):1–11, 2001.

[Che02] Han Chen. A parallel ultra-high resolution mpeg-2 video decoder for pc cluster based tiled display system. to appear. In *Proc. Int'l Parallel and Distributed Processing Symp. (IPDPS), IEEE CS*, page 30. Press, 2002.

[Che03] Han Chen. *Scalable and Ultra-High Resolution MPEG Video Delivery on Tiled Displays*. PhD thesis, Princeton University, 2003.

[CHS04] Ian Creighton and Chris Ho-Stuart. A sense of touch in online sculpting. In *GRAPHITE '04: Proceedings of the 2nd international conference on Computer graphics and interactive techniques in Australasia and South East Asia*, pages 118–122, New York, NY, USA, 2004. ACM.

[CI05] Alvaro Cassinelli and Masatoshi Ishikawa. Khronos projector. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Emerging technologies*, page 10, New York, NY, USA, 2005. ACM.

[Cor09] Carlos D. Correa. Visualizing what lies inside. *SIGGRAPH Comput. Graph.*, 43(2):1–6, 2009.

[CS02] Hui Chen and Hanqiu Sun. Real-time haptic sculpting in virtual volume space. In *VRST '02: Proceedings of the ACM symposium on Virtual reality software and technology*, pages 81–88, New York, NY, USA, 2002. ACM.

[CSC06] Carlos Correa, Deborah Silver, and Min Chen. Feature aligned volume manipulation for illustration and visualization. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):1069–1076, 2006.

[CSM02] E.F. Churchill, D.N. Snowdon, and A.J. Munro. Collaborative virtual environments: digital places and spaces for interaction. *Educational Technology & Society*, 5(4), 2002.

[CT09] Andrew A. Chien and Nut Taesombut. Integrated resource management for lambda-grids: The distributed virtual computer (dvc). *Future Generation Computer Systems*, 25(2):147 – 152, 2009.

[DC02] James Davis and Xing Chen. Lumipoint: multi-user laser-based interaction on large tiled displays. *Displays*, 23(5):205 – 211, 2002.

[DDS+09] Thomas A. DeFanti, Gregory Dawe, Daniel J. Sandin, Jurgen P. Schulze, Peter Otto, Javier Girado, Falko Kuester, Larry Smarr, and Ramesh Rao. The starcave, a third-generation cave and virtual reality optiportal. *Future Generation Computer Systems*, 25(2):169 – 178, 2009.

[DK10] Kai-Uwe Doerr and Falko Kuester. CGLX: A Scalable, High-performance Visualization Framework for Networked Display Environments. *IEEE Transactions on Visualization and Computer Graphics*, 99(PrePrints), 2010.

[DL01] P. Dietz and D. Leigh. Diamondtouch: a multi-user touch technology. In *Proceedings of the 14th annual ACM symposium on User interface software and technology*, pages 219–226. ACM New York, NY, USA, 2001.

[DLR+09] Thomas A. DeFanti, Jason Leigh, Luc Renambot, Byungil Jeong, Alan Verlo, Lance Long, Maxine Brown, Daniel J. Sandin, Venkatram Vishwanath, Qian Liu, Mason J. Katz, Philip Papadopoulos, Joseph P. Keefe, Gregory R. Hidley, Gregory L. Dawe, Ian Kaufman, Bryan Glogowski, Kai-Uwe Doerr, Rajvikram Singh, Javier Girado, Jurgen P. Schulze, Falko Kuester, and Larry Smarr. The optiportal, a scalable visualization, storage, and computing interface device for the optiputer. *Future Generation Computer Systems*, 25(2):114 – 123, 2009.

[EKCB03] Jr. Easton, R.L., K.T. Knox, and W.A. Christens-Barry. Multispectral imaging of the archimedes palimpsest. *Applied Imagery Pattern Recognition Workshop, 2003. Proceedings. 32nd*, pages 111–116, Oct. 2003.

[EKE01] Klaus Engel, Martin Kraus, and Thomas Ertl. High-quality pre-integrated volume rendering using hardware-accelerated pixel shading. In *HWWS '01: Proceedings of the ACM SIGGRAPH/EUROGRAPHICS workshop on Graphics hardware*, pages 9–16, New York, NY, USA, 2001. ACM.

[Elv92] T. Todd Elvins. A survey of algorithms for volume visualization. *SIGGRAPH Comput. Graph.*, 26(3):194–201, 1992.

[FAJ07] G. Flint, C. Aves, and MT Jones. The gigapxl project. http://www.gigapxl.org, 2007.

[GH91] Tinsley A. Galyean and John F. Hughes. Sculpting: an interactive volumetric modeling technique. *SIGGRAPH Comput. Graph.*, 25(4):267–274, 1991.

[Gra72] R. L. Graham. An efficient algorith for determining the convex hull of a finite planar set. *Information Processing Letters*, 1(4):132 – 133, 1972.

[GRC+07] J.F. Gantz, D. Reinsel, C. Chute, W. Schlichting, J. McArthur, S. Minton, I. Xheneti, A. Toncheva, and A. Manfrediz. The expanding digital universe: A forecast of worldwide information growth through 2010. *IDC white paper*, 2007.

[GSW01] François Guimbretière, Maureen Stone, and Terry Winograd. Fluid interaction with high-resolution wall-size displays. In *UIST '01: Proceedings of the 14th annual ACM symposium on User interface software and technology*, pages 21–30, New York, NY, USA, 2001. ACM.

[HA08] J. Heer and M. Agrawala. Design considerations for collaborative visual analytics. *Information Visualization*, 7(1):49–62, 2008.

[Han05] J.Y. Han. Low-cost multi-touch sensing through frustrated total internal reflection. In *Proceedings of the 18th annual ACM symposium on User interface software and technology*, pages 115–118. ACM New York, NY, USA, 2005.

[Har90] Stevan Harnad. The symbol grounding problem. *Physica D: Nonlinear Phenomena*, 42(1-3):335 – 346, 1990.

[HEB+01] Greg Humphreys, Matthew Eldridge, Ian Buck, Gordan Stoll, Matthew Everett, and Pat Hanrahan. Wiregl: a scalable graphics system for clusters. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 129–140, New York, NY, USA, 2001. ACM.

[Her08] L. Herr. Creation and Distribution of 4 K Content. *Television Goes Digital*, page 99, 2008.

[HKSB06] M. Hadwiger, A. Kratz, C. Sigg, and K. Bühler. Gpu-accelerated deep shadow maps for direct volume rendering. In *Graphics Hardware 2006: Eurographics Symposium Proceedings, Vienna, Austria, September 3-4, 2006*, pages 49–52. Eurographics Association, 2006.

[HLSR08] Markus Hadwiger, Patric Ljung, Christof Rezk Salama, and Timo Ropinski. Advanced illumination techniques for gpu volume raycasting. In *SIGGRAPH Asia '08: ACM SIGGRAPH ASIA 2008 courses*, pages 1–166, New York, NY, USA, 2008. ACM.

[HYB02] T. Hansen, P. Yalamanchili, and H.W. Braun. Wireless measurement and analysis on HPWREN. In *Proceedings of Passive and Active Measurement Workshop, Fort Collins, Co*, pages 222–229, 2002.

[Ini06] Digital Cinema Initiatives. Standard evaluation material (stem), 2006.

[JC06] G. Johansson and H. Carr. Accelerating marching cubes with graphics hardware. In *Proceedings of the 2006 conference of the Center for Advanced Studies on Collaborative research*, page 39. ACM New York, NY, USA, 2006.

[JJR+05] B. Jeong, R. Jagodic, L. Renambot, R. Singh, A. Johnson, and J. Leigh. Scalable graphics architecture for high-resolution displays. In *IEEE Information Visualization Workshop*, 2005.

[JRJ+06] Byungil Jeong, Luc Renambot, Ratko Jagodic, Rajvikram Singh, Julieta Aguilera, Andrew Johnson, and Jason Leigh. High-performance dynamic graphics streaming for scalable adaptive graphics environment. In *SC '06: Proceedings of the 2006 ACM/IEEE conference on Supercomputing*, page 108, New York, NY, USA, 2006. ACM.

[KCC08] D. Kim, K. Cha, and S.I. Chae. A high-performance openvg accelerator with dual-scanline filling rendering. *IEEE Transactions on Consumer Electronics*, 54(3):1303–1311, 2008.

[KKH01] J. Kniss, G. Kindlmann, and C. Hansen. Interactive volume rendering using multi-dimensional transfer functions and direct manipulation widgets. In *Proceedings of the conference on Visualization'01*, pages 255–262. IEEE Computer Society Washington, DC, USA, 2001.

[KSR$^+$06] Matthias Koenig, Wolf Spindler, Jan Rexilius, Julien Jomier, Florian Link, and Heinz-Otto Peitgen. Embedding vtk and itk into a visual programming and rapid prototyping platform. *Medical Imaging 2006: Visualization, Image-Guided Procedures, and Display*, 6141(1):61412O, 2006.

[KUDC07] Johannes Kopf, Matt Uyttendaele, Oliver Deussen, and Michael F. Cohen. Capturing and viewing gigapixel images. In *SIGGRAPH '07: ACM SIGGRAPH 2007 papers*, page 93, New York, NY, USA, 2007. ACM.

[KVV$^+$04] NK Krishnaprasad, V. Vishwanath, S. Venkataraman, AG Rao, L. Renambot, J. Leigh, AE Johnson, and B. Davis. JuxtaView-a tool for interactive visualization of large imagery on scalable tiled displays. In *Cluster Computing, 2004 IEEE International Conference on*, pages 411–420, 2004.

[LBS85] SK Lee, W. Buxton, and KC Smith. A multi-touch three dimensional touch-sensitive tablet. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 21–25. ACM New York, NY, USA, 1985.

[LC87] W.E. Lorensen and H.E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, pages 163–169. ACM New York, NY, USA, 1987.

[Lee84] S. Lee. A fast multiple-touch-sensitive input device. *Master's thesis, University of Toronto*, 1984.

[Leh97] Roy S. Lehrle. Forensics, fakes, and failures: Pyrolysis is one part in the overall armoury. *Journal of Analytical and Applied Pyrolysis*, 40-41:3 – 19, 1997. PYROLYSIS '96.

[LM04] Eric B. Lum and Kwan-Liu Ma. Lighting transfer functions using gradient aligned sampling. In *VIS '04: Proceedings of the conference on Visualization '04*, pages 289–296, Washington, DC, USA, 2004. IEEE Computer Society.

[Mar91] K. Martinez. High resolution digital imaging of paintings: The vasari project. *Microcomputers for Information Management*, 8(4):277–83, 1991.

[MC05] Kirk Martinez and John Cupitt. Vips - a highly tuned image processing software architecture. In *ICIP (2)*, pages 574–577, 2005.

[MCSP02] K. Martinez, J. Cupitt, D. Saunders, and R. Pillay. Ten years of art imaging research. *Proceedings of the IEEE*, 90(1):28–41, 2002.

[MDH+03] A. MacEachren, X. Dai, F. Hardisty, D. Guo, and G. Lengerich. Exploring high-D spaces with multiform matrices and small multiples. In *IEEE Symposium on Information Visualization, 2003 (INFOVIS 2003); 19–21 Oct. 2003; Seattle, Washington*, pages 31–38. Citeseer, 2003.

[Mit97] J.L. Mitchell. *MPEG video compression standard*. Kluwer Academic Publishers, 1997.

[Mor98] H. Moravec. When will computer hardware match the human brain. *Journal of Evolution and Technology*, 1:1–14, 1998.

[MRB05] Shahzad Malik, Abhishek Ranjan, and Ravin Balakrishnan. Interacting with large displays from a distance with vision-tracked multi-finger gestural input. In *UIST '05: Proceedings of the 18th annual ACM symposium on User interface software and technology*, pages 43–52, New York, NY, USA, 2005. ACM.

[MTB03] Michael J. McGuffin, Liviu Tancau, and Ravin Balakrishnan. Using deformations for browsing volumetric data. In *VIS '03: Proceedings of the 14th IEEE Visualization 2003 (VIS'03)*, page 53, Washington, DC, USA, 2003. IEEE Computer Society.

[NR02] S. Navrud and R.C. Ready. *Valuing cultural heritage*. Elgar, 2002.

[PDMDRP08] A. Pelagotti, A. Del Mastio, A. De Rosa, and A. Piva. Multispectral imaging of paintings. *Signal Processing Magazine, IEEE*, 25(4):27–36, July 2008.

[PKS+08] Peter Peltonen, Esko Kurvinen, Antti Salovaara, Giulio Jacucci, Tommi Ilmonen, John Evans, Antti Oulasvirta, and Petri Saarikko. It's mine, don't touch!: interactions at a large multi-touch display in a city centre. In *CHI '08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pages 1285–1294, New York, NY, USA, 2008. ACM.

[Ple08] L. Plesea. The design, implementation and operation of the JPL OnEarth WMS server. In *Geospatial Services and Applications for the Internet*, pages 93–109. Springer US, 2008.

[PSH97] Vladimir I. Pavlovic, Rajeev Sharma, and Thomas S. Huang. Visual interpretation of hand gestures for human-computer interaction: A review. 1997.

[PWFO01] KL PERNG, WT WANG, M. FLANAGAN, and M. OUHYOUNG. A Real-time 3D Virtual Sculpting Tool Based on Modified Marching Cubes. In *Int Conf Artif Real Telexistence*, volume 11, pages 64–72, 2001.

[RBJW01] Meredith Ringel, Henry Berg, Yuhui Jin, and Terry Winograd. Barehands: implement-free interaction with a wall-mounted display. In *CHI '01: CHI '01 extended abstracts on Human factors in computing systems*, pages 367–368, New York, NY, USA, 2001. ACM.

[Rek98] Jun Rekimoto. A multiple device approach for supporting whiteboard-based interactions. In *CHI '98: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 344–351, New York, NY, USA, 1998. ACM Press/Addison-Wesley Publishing Co.

[RJJ+06] L. Renambot, B. Jeong, R. Jagodic, A. Johnson, J. Leigh, and J. Aguilera. Collaborative visualization using high-resolution tiled displays. In *ACM CHI Workshop on Information Visualization Interaction Techniques for Collaboration Across Multiple Displays*, 2006.

[RJL05] L. Renambot, A. Johnson, and J. Leigh. Lambdavision: Building a 100 megapixel display. In *NSF CISE/CNS Infrastructure Experience Workshop, Champaign, IL*, 2005.

[RP00] M. Riesenhuber and T. Poggio. Models of object recognition. *Nature Neuroscience*, 3:1199–1204, 2000.

[Ryd] Thomas Rydell. Virtual autopsy table. https://www.tii.se/projects/autopsy.

[SBD+] J. Schöning, P. Brandl, F. Daiber, F. Echtler, O. Hilliges, J. Hook, M. Löchtefeld, N. Motamedi, L. Muller, P. Olivier, et al. Multitouch surfaces: A technical guide.

[SBdL09] Larry Smarr, Maxine Brown, and Cees de Laat. Special section: Optiplanet – the optiputer global collaboratory. *Future Generation Computer Systems*, 25(2):109 – 113, 2009.

[SC93] D. Saunders and J. Cupitt. Image processing at the national gallery: The vasari project. 1993.

[SGHB07] J.D. Smith, TC Graham, D. Holman, and J. Borchers. Low-cost malleable surfaces with multi-touch pressure sensitivity. In *Horizontal Interactive Human-Computer Systems, 2007. TABLETOP'07.*

Second Annual IEEE International Workshop on, pages 205–208, 2007.

[SGM03] Stacey D. Scott, Karen D. Grant, and Regan L. Mandryk. System guidelines for co-located, collaborative work on a tabletop display. In *ECSCW'03: Proceedings of the eighth conference on European Conference on Computer Supported Cooperative Work*, pages 159–178, Norwell, MA, USA, 2003. Kluwer Academic Publishers.

[SHP+96] Rajeev Sharma, Thomas S. Huang, Vladimir I. Pavlovi'c, Yunxin Zhao, Zion Lo, Stephen Chu, Klaus Schulten, Andrew Dalke, Jim Phillips, Michael Zeller, and William Humphrey. Speech/gesture interface to a visual computing environment for molecular biologists. In *IEEE Computer Graphics and Applications*, pages 30–35, 1996.

[SLJM08] D. Svistula, J. Leigh, A. Johnson, and P. Morin. MagicCarpet: a high-resolution image viewer for tiled displays, 2008.

[SPS48] C. Shannon, N. Petigara, and S. Seshasai. The Mathematical Theory of Communication . *Communication, Bell System Technical Journal*, 1948.

[Sre08] M. Sreenivasan. Microsoft silverlight. 2008.

[STA] C. STANDARD. THE MPEG VIDEO COMPRESSION STANDARD.

[SVFR04] Chia Shen, Frédéric D. Vernier, Clifton Forlines, and Meredith Ringel. Diamondspin: an extensible toolkit for around-the-table interaction. In *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 167–174, New York, NY, USA, 2004. ACM.

[SVS+05] R. Stockli, E. Vermote, N. Saleous, R. Simmon, and D. Herring. he blue marble next generation – a true color earth dataset including seasonal dynamics from modis. *Published by the NASA Earth Observatory*, 2005.

[SW06] Bram Stolk and Paul Wielinga. Building a 100 mpixel graphics device for the optiputer. *Future Generation Computer Systems*, 22(8):972 – 975, 2006.

[SYK+05] H. Shimamoto, T. Yamashita, N. Koga, K. Mitani, M. Sugawara, F. Okano, M. Matsuoka, J. Shimura, I. Yamamoto, T. Tsukamoto, et al. An Ultrahigh-Definition Color Video Camera With 1.25-inch

Optics and 8k x 4k Pixels. *SMPTE Motion Imaging Journal*, pages 3–11, 2005.

[SYS+06] D. Shirai, T. Yamaguchi, T. Shimizu, T. Murooka, and T. Fujii. 4k shd real-time video streaming system with jpeg 2000 parallel codec. In *Circuits and Systems, 2006. APCCAS 2006. IEEE Asia Pacific Conference on*, pages 1855–1858, Dec. 2006.

[TC05] James J. Thomas and Kristin A. Cook. *Illuminating the Path: The Research and Development Agenda for Visual Analytics.* National Visualization and Analytics Ctr, 2005.

[TFM96] S. Thorpe, D. Fize, and C. Marlot. Speed of processing in the human visual system. *Nature*, 381(6582):520–522, 1996.

[Tuf91] E.R. Tufte. Envisioning information. *Optometry and Vision Science*, 68(4):322, 1991.

[TWC+06] Nut Taesombut, Xinran (Ryan) Wu, Andrew A. Chien, Atul Nayak, Bridget Smith, Debi Kilb, Thomas Im, Dane Samilo, Graham Kent, and John Orcutt. Collaborative data visualization for earth sciences with the optiputer. *Future Generation Computer Systems*, 22(8):955 – 963, 2006.

[VBRR02] G. Voß, J. Behr, D. Reiners, and M. Roth. A multi-thread safe foundation for scene graphs and its extension to clusters. In *EGPGV '02: Proceedings of the Fourth Eurographics Workshop on Parallel Graphics and Visualization*, pages 33–37, Aire-la-Ville, Switzerland, Switzerland, 2002. Eurographics Association.

[VL03] H.R. Varian and P. Lyman. How much information. *University of California at Berkeley, School of Information Management & Systems (SIMS)*, 2003.

[VOT] http://www.nationalgeographic.com/field/projects/valley-khans-project.html.

[WAB+05] G. Wallace, O.J. Anshus, P. Bi, H. Chen, Y. Chen, D. Clark, P. Cook, A. Finkelstein, T. Funkhouser, Anoop Gupta, M. Hibbs, K. Li, Z. Liu, Rudrajit Samanta, Rahul Sukthankar, and O. Troyanskaya. Tools and applications for large-scale display walls. *Computer Graphics and Applications, IEEE*, 25(4):24–33, 2005.

[WE98] R. Westermann and T. Ertl. Efficiently using graphics hardware in volume rendering applications. In *Proceedings of SIGGRAPH*, volume 98, pages 169–178, 1998.

[WEH01] W. Westerman, J. Elias, and A. Hedge. Multi-touch: A new tactile 2-d gesture interface for human-computer interaction. In *Proceedings of the Human Factors and Ergonomics Society 45th Annual Meeting*, volume 1, pages 632–636, Minneapolis/St. Paul, MN, 2001.

[Wil83a] Lance Williams. Pyramidal parametrics. In *SIGGRAPH '83: Proceedings of the 10th annual conference on Computer graphics and interactive techniques*, pages 1–11, New York, NY, USA, 1983. ACM.

[Wil83b] Lance Williams. Pyramidal parametrics. In *SIGGRAPH '83: Proceedings of the 10th annual conference on Computer graphics and interactive techniques*, pages 1–11, New York, NY, USA, 1983. ACM.

[Wil04] A.D. Wilson. Touchlight: an imaging touch screen and display for gesture-based interaction. In *Proceedings of the 6th international conference on Multimodal interfaces*, pages 69–76. ACM New York, NY, USA, 2004.

[WK95] Sidney W. Wang and Arie E. Kaufman. Volume sculpting. In *I3D '95: Proceedings of the 1995 symposium on Interactive 3D graphics*, pages 151–ff., New York, NY, USA, 1995. ACM.

[wms] http://www.opengeospatial.org/standards/wms.

[Zha09] Jian-Feng Zhang. Gpu-based direct volume rendering with advanced illumination and deep attenuation shadows. *Computer-Aided Design and Computer Graphics, 2009. CAD/Graphics '09. 11th IEEE International Conference on*, pages 536 –539, 2009.