Available for licensing

PixelVizion

An NPU-Embedded Visualization Accelerator for Large Data Sets

- High-frame-rate image composition
- Cost-effective commercial off-the-shelf solution using commodity interconnect technology
- Modular and scalable to fit customer needs
- Solves critical visualization bottlenecks

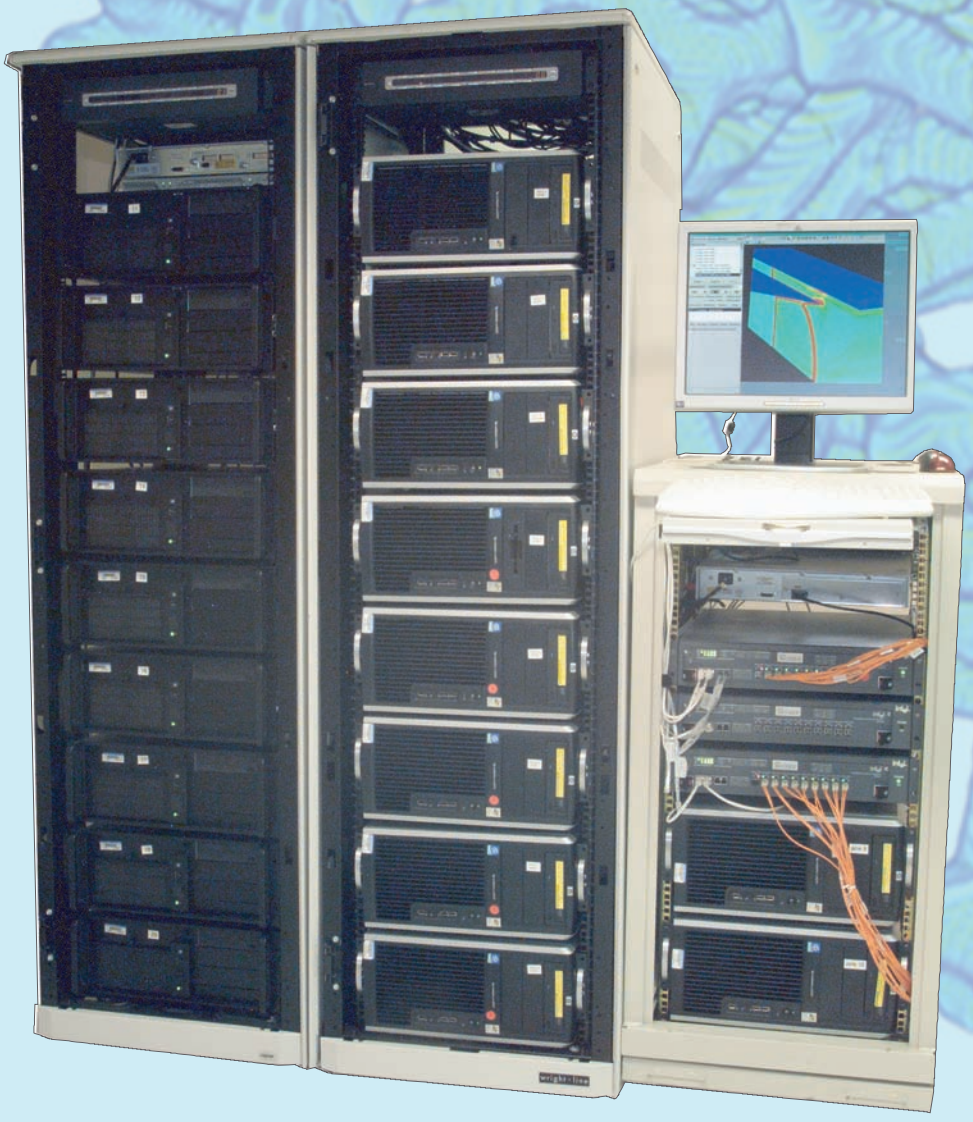

Los Alamos
NATIONAL LABORATORY
EST. 1943

PixelVizion

An NPU-Embedded Visualization Accelerator for Large Data Sets

Technology Transfer Contacts:
John Russell, email: tmt-3@lanl.gov

Technical Contact
Carolyn Conner-Davenport

ABOUT THE COVER

PixelVizion, the first Network Processor Unit (NPU)-based visualization accelerator, can be used to provide high-speed graphics composition for large data sets. Shown is an image from the Foam Crush data set. The image shows the noncompressed foam microstructure, a network of struts and membranes. In front of the image is the PixelVizion system. The screen shows a simulation of a shock wave through a wedge (courtesy of the California Institute of Technology) that is described on the DVD. The two large racks contain 16 rendering nodes. The smaller rack contains three NPUs, which bring single-pass network data transmission and on-the-fly image compositing, plus a display node.

Los Alamos NATIONAL LABORATORY EST.1943

LA-UR-06-1208

Executive Summary

PixelVizion

An NPU-Embedded Visualization Accelerator for Large Data Sets

Features

As imaging and video technology continues to advance, the need to process, analyze, sort, and manipulate large data sets has grown tremendously. The image compositing function has become a visualization bottleneck. PixelVizion is the first Network Processor Unit (NPU)-based computer visualization tool that addresses this bottleneck. It brings single-pass network data transmission and on-the-fly image compositing that yields an order-of-magnitude increase in interactive response times. PixelVizion is a hardware-assisted, lossless, highly scalable, high-frame-rate solution to the visualization bottleneck of image compositing. It composites extremely large volumes of data at rates that are 10 to 20 times faster than those of current compositing technologies.

Applications

- Orthopedics, rehabilitation, and sports science
- Virtual medical training
- Specialized diagnostic imaging
- Virtual skin grafting
- Weather patterns
- Large-scale scientific problems
- Animation and special effects
- Video game graphics
- Film postprocessing

Benefits

- Provides 10× to 20× faster image composition
- Cost-effective, commercial off-the-shelf solution
- Removes the need for an expensive network interconnect
- Flexible programming to enable faster development cycles
- Highly modular and scalable
- Frees processing power for other applications
- Accommodates a variety of software rendering packages

Overview

As imaging and video technology continues to advance, modern society has become increasingly reliant upon electronically generated visual data. Correspondingly, the need to process, analyze, sort, and manipulate visual data (i.e., compositing) has also grown tremendously. Because image and visualization data sets are by nature very large (i.e., terabytes to petabytes, or a trillion to quadrillion bytes), the task of processing all the data can be complex, expensive, and time-consuming. Applications that alleviate the processing burden and allow users to access and manipulate data faster and more effectively are in demand.

Because of the enormity of the data sets, visualization often requires workload distribution across a cluster of powerful computational nodes (i.e., computers networked together to integrate and aggregate computational resources). Workload distribution divides the problem into manageable pieces but results in a cluster of images that must be combined to produce a final image. Traditionally, these nodes both render (i.e., generate the pixels of an image from a high-level description of the image's components) and composite the data.

Compositing is the process of combining multiple images as overlapping layers to produce a single output image, such as was done in the cartoon industry when each artist drew one aspect of the scene as a layer and the multiple layers were combined to form the final frame. Today, image composition is applied in a wide variety of areas such as medicine, computer animation, special effects in movies, video game graphics, bio-informatics, and scientific visualization and modeling. However, compositing is a recognized, resource-intensive, processing bottleneck. PixelVizion implements alpha-blend compositing and Z-buffer ordering of the preliminary images. Alpha blending is used to create the effect of transparency by combining a translucent foreground with a background color. Z-buffering is used to determine which objects are visible and which are hidden.

PixelVizion is the first Network Processor Unit (NPU)-embedded visualization accelerator. It provides a unique hardware-assisted, lossless, highly scalable, high-frame-rate solution to the challenge of image composition, and reduces costs and overall time to solution (product/diagnosis/result) through a simple application programming interface (API). PixelVizion removes the composition task from the rendering nodes, reducing the nodes input/output (I/O) requirements by more than 50%. The NPU hardware does the image composition on-the-fly, thus realizing significant gains in efficiency and performance and freeing potential processing

capability for other applications. Furthermore, many visualization applications are highly sensitive to pixel quality and cannot tolerate losses resulting from compression. PixelVizion provides line-rate (Gbps), lossless visualization acceleration. For further details on the NPU features, see “Network Processor Unit (NPU)-Embedded Graphics Composition for Large-Scale Data Visualization” in the accompanying appendix.

PixelVizion exploits commercial off-the-shelf (COTS) hardware end-to-end. It removes the need for an expensive interconnection network, providing faster compositing and, in turn, realizing new processing potential for large data sets. The use of COTS allows a wide variety of algorithms or algorithm variations to be implemented through the NPU. PixelVizion’s NPU code uses Teja C, a slight variant of the standard C programming language, which makes PixelVizion highly adaptable and easily modifiable to incorporate new or different operations/functionality.

PixelVizion offloads the composition function from the rendering nodes to an NPU device that has been optimized for high-speed processing of data packets (i.e., data packaged for ease of transmission). While NPU technology has been widely applied in the data communications industry, PixelVizion is the first to apply this technology in the field of computer visualization. The NPU compositor receives multiple input streams of data packets from the cluster nodes and performs integer-based numerical processing across the streams to produce a new single-stream packet. It then transmits that packet to an output device (either another NPU or a display device). A single image on the output device is composed of thousands of these packets. PixelVizion allows for single-pass compositing, i.e., data is transferred only once from each rendering node. This single-pass is in contrast to the multiple image-data exchanges across the cluster network required by software-only compositors.

Competition

- Software-only compositing methods
  - VisIt, a Lawrence Livermore National Laboratory (LLNL) interactive parallel visualization software tool for viewing scientific data on Unix and PC platforms, provides representative implementation of software-only compositing. VisIt was designed to handle data sets from the kilobyte to terabyte scale.
- Hardware compositing
  - The Orad VR-X is commercial hardware compositor developed by Orad Hi-Tec Systems, Ltd. The Orad product is based upon

proprietary FPGA (field-programmable gate array) technology and proprietary software and is marketed for television broadcasting applications. Our knowledge of the product is limited to information on the company's web site.

Other competitors are HP's Sepia + VolPro Cards and Mitsubishi's MPC. Sepia is no longer commercially available. To our knowledge, MPC is not available in the United States and information on this product is unavailable to the researcher.

Comparison matrix

Parameters	PixelVizion	VisIt	Orad VR-X	Advantages
Available Alpha-Blend Compositing Solution?	Yes	Yes	Unknown	PixelVizion is a demonstrated, fully functional preproduction system. VisIt is readily available to the scientific research community. Orad VR-X is a commercially available hardware compositor. Orad seems to use a Z-buffering algorithm. It is unknown whether it provides alpha-buffering as well.
High-Frame-Rate Compositing (i.e., >20 fps for 1024 × 1024 pixel image)?	Yes	No	Unknown	PixelVizion's alpha-blend compositing performance allows end users to more quickly and easily gain critical insights into complex scientific/medical data. PixelVizion also helps speed the pre- and post-production of complex movies/games by increasing the compositing from 2 fps to 20 fps. Traditional "software only" solutions such as VisIt are at least 18 fps slower than PixelVizion. Orad did not respond to our request for performance information.
Cost-Effective Comparative Performance?	Yes	No	No	PixelVizion costs $28K per box ($18K for the hardware and $10K for the software), with each box capable of supporting eight rendering nodes. Orad costs between $56K and $72K for an eight-node proprietary solution. PixelVizion uses a gigabit Ethernet that is currently supplied on all computer motherboards. Orad uses a proprietary "pixel bus" to interconnect its hardware.
COTS Compositing Solution?	Yes	N/A	Yes/No	PixelVizion is a commercial off-the-shelf (COTS) solution employing stable technology, available from multiple suppliers. PixelVizion requires no special modifications to the cluster node hardware. Orad is a COTS product with propriety hardware and software and requires modification of the computer graphics card.
Scalable to Larger Clusters?	Yes	Yes	Yes	Each is scalable. PixelVizion is a much more cost-effective solution than Orad's VR-X for systems with more than four nodes.
Graphics Hardware Upgradable?	Yes	N/A	Yes/No	PixelVizion is completely COTS-based, so one can upgrade graphics cards at will and purchase them on the open market. Orad, on the other hand, alters the graphics cards and supplies the altered cards with its system.
Performance Is Independent of Image Complexity?	Yes	No	Unknown	PixelVizion does not require any special data-reduction/compression processing (i.e., background pixel removal or image data compression) to obtain performance. Software-only solutions rely on extra data-reduction processing to enhance performance.
Flexibly Programmable?	Yes	Yes	No	PixelVizion's NPU code uses Teja C, a slight variant of the standard C programming language, which makes PixelVizion easily modifiable to incorporate new or different operations/functionality. The Orad system is far less flexible and more constrained because of its reliance on FPGA technology. FPGAs are much harder to program and take much longer to modify compared with high-level software. In addition, any changes would have to be made by Orad.
Frees Cluster Cycles for Other Needs and Applications?	Yes	No	Yes	PixelVizion removes the compositing task from the cluster nodes, allowing them to work efficiently on rendering and other visualization tasks. Additionally, because PixelVizion is a separate piece of hardware/software, rendering and compositing tasks overlap to increase frames per second of the whole visualization process. Orad, being separate hardware, shares some of these same features.
Adaptable to Other Methods Requiring Sort-Last Compositing?	Yes	Yes	Yes	PixelVizion is rendering-method agnostic, requiring only an image and a view-dependant depth scalar. Orad has a similar function.
Extendable to Multipaneled Facilities?	Yes	Yes	Yes	With PixelVizion, each panel is connected to a display node that is in turn connected to a PixelVizion NPU device. Orad uses a similar scheme to drive a multipaneled display but requires more devices than PixelVizion.
Requires Only Simple, Coarse, Frame-Level Synchronization across Cluster?	Yes	Yes	No	Because of the NPU's built-in, high-performance packet queues, PixelVizion requires only a simple per-frame synchronization scheme, which is what VisIt and other software compositors use. Orad uses a proprietary connection to the graphics card clock to provide synchronization to their FPGAs.
Requires Specialized Hardware Device Drivers?	No	N/A	Yes	Because PixelVizion is connected to the cluster nodes through gigabit Ethernet, no special device drivers are necessary, enabling the user to use the most up-to-date drivers and extensions from the graphics card manufacturer.
Requires Specialized Interface Hardware (e.g., FPGA board)?	No	N/A	Yes	PixelVizion is independent of the cluster nodes and does not require modification of the nodes. Orad uses specialized FPGA hardware that is tightly coupled with the COTS graphics card and requires modification of the graphics card.

Advantages

PixelVizion composites images up to 10 to 20 times faster than its competitors. PixelVizion composites images at 21 to 29 frames per second for a 1024 × 1024 pixel image and at greater than 83 frames per second for a 512 × 512 pixel image. PixelVizion provides single-pass compositing; data is transferred only once from each rendering node. The NPU adds only a small fixed delay in the visualization process. This PixelVizion aspect is unique and powerful in comparison to other compositing technologies.

PixelVizion is a cost-effective, commercial-off-the-shelf (COTS) solution. PixelVizion uses proven technology that is available from multiple suppliers, providing a stable and cost-effective, high-performance compositing solution. As such, when the underlying technology advances, for instance to exploit the 10-gigabit Ethernet, PixelVizion's capability, flexibility, and performance will advance in step.

PixelVizion delivers high performance without an expensive network interconnect. PixelVizion uses a ubiquitous and inexpensive network interconnect (gigabit Ethernet) that is currently supplied on all computer motherboards, removing the need for an expensive cluster interconnect. The software-only approaches require the use of an Infiniband network, while other hardware approaches require expensive and proprietary communication links.

PixelVizion is flexibly configurable. The PixelVizion system can be configured and balanced to meet specific user needs. Design factors including image resolution, number of nodes/subvolumes, functionality, configuration (standard, stereo, tiled, chained), and budget can be balanced to provide the optimal solution to meet user needs. PixelVizion can be configured to drive a stereo display using two NPUs. One NPU and its associated rendering nodes are used to generate the left channel, while the second NPU and its rendering nodes are used to generate the right channel. The results are then displayed on a stereo display. The compositing hardware and software are highly modular. More advanced graphics cards may be plugged into the system at will.

PixelVizion is highly scalable. PixelVizion is completely scalable both in terms of depth (i.e., adding more rendering nodes per display window) and in terms of final display resolution, known as tiling. Depth scaling allows computation on larger data volumes. Tiling, the addition of display windows to produce a larger display image, allows exploitation of power walls for collaborative and interactive data analysis. PixelVizion's unique queuing and data packet synchronization functions allow for robust scalability. It handles large fluctuations between rendering nodes in real time.

PixelVizion is flexibly programmable. The programming model and flexibility of the NPU enables faster development

cycles, increased product/application longevity, and the capability to implement new algorithms and algorithm variations.

PixelVizion will ride the data communications technology curve. As NPU technology evolves, PixelVizion's compositing performance will increase directly with little to no change required to the existing software.

PixelVizion takes full advantage of advances in graphics processing unit technology. The performance of graphics cards is currently accelerating at a speed greater than Moore's Law would predict, and in recent years, programmability has been introduced into the graphics pipeline, allowing such things as programmable shaders. Because PixelVizion does not interfere with the graphics pipeline at all, commericial graphics card drivers can be used, and all current and future capabilities of graphics technology can be accommodated.

PixelVizion enhances its software-only "competitors." PixelVizion is in a unique position. While it significantly outperforms existing technology, in actuality it is not a competitor but rather enhances software-only visualization packages. As an embedded hardware acceleration tool, it seamlessly replaces the software compositor package, operating completely independently of all other software visualization functionality. This replacement is a great advantage for the users, who can fully use their existing software while taking advantage of PixelVizion's unmatched compositing performance.

Primary applications

Image composition is currently used across a diverse set of industries and applications. PixelVizion is particularly powerful because it allows a wide variety of algorithms or algorithm variations to be implemented. Some of its many applications are as follows:

- **Medical Industry**—PixelVizion can be used to visualize human movement for orthopedics, rehabilitation, and sports science. It can also be used to model virtual skin grafting, for virtual medical training and procedure planning, and for specialized diagnostic imaging such as radiology or teleradiology, and x-ray computed tomography (CT) scans.
- **Scientific Research**—PixelVizion can be used to analyze, visualize, and model large-scale scientific problems such as weather patterns, global warming, proteins, and the human genome.
- **Entertainment Industry**—PixelVizion can be used to cut film postprocessing time by 20%. PixelVizion will provide producers and directors with the ability to have more artistic control and make creative modifications more easily. PixelVizion can also be used for animation, special effects, and video game graphics.

Other applications

Any project relying on cluster-based visualization (sort-last parallel volume rendering) of large datasets is a potential beneficiary of this technique. Such applications include the following:

- **Automotive Industry**—PixelVizion can be used to identify potential failure mechanisms in designs for automobile engines and components. It can also be used to assist in developing new design features and to model engine efficiencies, and critical structure time-to-failure .
- **Environmental Modeling**—PixelVizion can be used with groundwater data visualization applications to improve understanding of complex spatial and temporal patterns of multiple analytes and geology. It can also be used with geophysical 3-D seismic renderings of complex forces such as earthquakes, volcanic eruptions, tsunamis, and oceanic currents.
- **Oil and Gas Industry**—Reservoir modeling could benefit from PixelVizion. To predict petroleum reserves, modeling and visualization of very large 3-D data volumes are performed. Much of this data is visualized as voxel data, solids data, and temporal movies. As the contents of a reservoir are depleted, movies are made of the depletion over time. PixelVizion could be used to generate the movies in a more cost-effective manner to improve depletion control and minimize potential ecological events.

Summary

From Leonardo Da Vinci's anatomical drawings to Wilhelm Roentgen's first x-ray of the human hand to today's use of computer graphics and virtual reality to "fly through" three-dimensional data reconstructions, researchers have for centuries used visualization in their quest to understand human and physical complexity. Accurate, physics-based computer simulations are widely used in all branches of science and engineering. As imaging and video technologies have advanced, society itself has become increasingly reliant upon electronically generated visual data. Interactive visualization can greatly enhance the exploration of data and phenomena by exploiting the human brain's ability to process enormous amounts of visual information. Computer simulations are used to forecast the spread of wildfires and contagious diseases, predict the weather, and to identify potential failure modes in complex mechanical systems. Simulations have also become common in the film, entertainment, and advertising industries, as well.

Because visualization data sets are by nature extremely large, the task of analyzing, sorting, and manipulating the data can be complex, expensive, resource-intensive, and time-consuming. The ability to change classification functions and color mappings, animate forward and backward in time, change viewpoints, or zoom in and out on features, all at interactive rates, is essential for

maximizing scientific productivity. Consequently, there is great motivation to increase the speed of visualization technology.

PixelVizion is the first visualization acceleration tool to use the power and flexibility of network processor technology. It provides hardware-accelerated image composition that operates seamlessly within existing visualization software packages. It removes a notorious image-processing bottleneck and provides composited images at speeds greater than 80 frames per second (i.e., from 10 to 20 times faster than currently available technology). PixelVizion is a cost-effective solution that provides unparalleled performance and capability using proven, commercial off-the-shelf technology. It allows users (e.g., scientists, animators, surgeons, etc.) to engage their data faster and more effectively, thus allowing them to solve problems faster, more creatively, and more productively. Furthermore, PixelVizion is flexible and fully scalable, both in terms of higher image resolutions and/or more rendering nodes.

PixelVizion represents a unique, outside-the-box innovation in terms of how it exploits image processing technology, its flexibility and breadth of applicability, and its unparalleled performance in computer visualization.

The appendix includes seven letters of endorsement from diverse fields such as universities, the oil and gas industry, national laboratories, and a network processing unit manufacturer. It also includes further technical details.

A DVD with a short clip showing PixelVizion's capabilities accompanies this entry. The clip shows researchers exploring Foam Crush (the physics of foam mechanics) and asteroid impact data in an immersive visualization environment. The clip then moves to a scientist using PixelVizion to examine a single frame in a material science data set, and a simulation of a shock wave through a wedge of material. It then describes PixelVizion technology and shows the performance comparison of PixelVizion to a software-only compositing solution.