VolumePro Takes New Approach to 3D

First Single-Chip Volume Renderer Shows Solid Potential

by Peter N. Glaskowsky

A new 3D-graphics chip from the Mitsubishi Electric Research Laboratory (MERL) promises to bring 3D graphics to scientific applications that benefit little from existing 3D accelerators. The company's VolumePro rendering chip, announced by chief architect Hanspeter Pfister at last month's Microprocessor Forum, is the first of its kind to do volume rendering in a single device. In the past, affordable volume rendering has generally been limited to slow software solutions, and only supercomputer graphics systems have had the ability to do volume rendering in real time. VolumePro promises to make interactive volume rendering available to many professional users for the first time.

Volume Rendering Creates New Markets for 3D

How does volume rendering differ from conventional 3D? Today's real-time 3D rendering engines all work in the same basic way. Polygons define the surface contours of objects, while lighting, shading, and texturing effects define how these polygons will appear on a computer screen. A conventional 3D model of a human body can accurately convey the external appearance of the person depicted, but such a model is not particularly useful to a physician concerned with the internal structure of an ill patient.

MICHAEL MUSTACCHI

Hanspeter Pfister, chief architect of VolumePro, explains the chip's volume-rendering pipeline.

Medical imaging through computer-aided tomography (CAT) scans or magnetic resonance imaging (MRI) produces detailed three-dimensional information about the internal structure of patients' bodies. The National Library of Medicine's Visible Human project, available online at *www.nlm.nih.gov/research/visible*, provides an excellent example of this sort of work. These 3D datasets, however, can't be directly displayed or manipulated by conventional 3D engines. Over the past several years, a new field of 3D technology, known as volume rendering, has developed to meet the needs of medical imaging and similar applications in geology, meteorology, and aerodynamics, where 3D datasets are commonly used.

A 2D cross-section of a 3D dataset is easily generated, but a cross-section includes only a fraction of the information in the whole model. It takes physicians years to learn how to interpret the individual layers of a CAT scan and recognize the faint changes that signal the presence of a tumor. A volume-rendering system configured to translate different tissue densities to different colors and opacities can make a tumor more apparent and provide physicians with even more critical information about how the tumor has affected the surrounding tissue.

VolumePro Manages Large Datasets

The richness of 3D datasets is also the primary characteristic that has delayed integrated volume-rendering implementations until now. A 256×256 image in 8-bit gray scale occupies a modest 64K. A 256^3 volume dataset has 16 million voxels (volume elements). VolumePro supports up to 128M of local voxel memory, enough to hold several datasets at once. This capacity allows the operator to switch rapidly among views based on several different characteristics of a single object (density, temperature, composition, etc.) or to view several different objects (such as an object under test and a reference object). Even larger datasets, such as geological scans for oil exploration, can be partitioned into "bricks" and swapped into volume memory under software control.

As you'd expect for a task of this complexity, VolumePro has a complex architecture. The volume-rendering process is very parallelizable, and MERL's first VolumePro chip, the vg500, has four parallel rendering pipelines. The 0.35-micron vg500 has 800,000 gates and 2.2 megabits of internal SRAM. The chip measures 186 mm^2 and operates at 133 MHz. In addition to the four rendering pipelines, the chip has interfaces for PCI and both voxel and pixel memory.

Rendering Uses Ray Casting

The VolumePro architecture relies on a technique known as ray casting to produce a 2D view of a 3D dataset. The image plane may be at any arbitrary angle to the nearest surface of the dataset, known as the base plane. As Figure 1 shows, each pixel in the base plane is used as the origin of a line, or ray, cast perpendicular to the image plane through the dataset. As these rays pass through the dataset, they will not generally pass through the center of each voxel, so trilinear interpolation—a weighted average of the values of the eight nearest voxels—creates a new voxel value for each step along the ray.

Color and opacity values are assigned to each new voxel value according to parameters defined by the application. For example, false color may be used to show changes in density within the volume. In some models, the ray may encounter gradients within the volume that indicate an internal surface, such as the boundary between muscle and bone. VolumePro calculates the 3D angle of the surface at that point, and can apply lighting and shading effects to the surface in a manner similar to that of conventional 3D chips. VolumePro uses Phong shading, a technique that depends on the actual angle of the surface at each point. Phong shading produces better visual results than the interpolated Gouraud shading used in most conventional 3D chips.

Internally, the vg500's rendering pipelines are relatively straightforward, as Figure 2 shows. As the rendering engine steps along each ray and voxel data passes through the pipe, the process generates both alpha (opacity) and color information. The compositing step produces the final result for each pixel, which is then stored in local pixel memory.

The vg500 does not include its own display controller. Instead, pixels are accumulated in a separate bank of local memory, then transferred out through the chip's PCI interface to a conventional 3D chip. The vg500 depends on the separate 3D chip to perform a final image-warp operation to convert from the base-plane perspective created by the vg500 to the image-plane perspective requested by the user.

Volume-Rendering Performance Sets New Mark

According to MDR's estimates, VolumePro will be at least an order of magnitude faster than competing solutions, including some that are far more expensive. Host-based volume-rendering software is very slow, taking a minute or more per frame for 256^3-voxel datasets. Even graphics supercomputer systems provide little hardware acceleration for volume rendering, generating only a few frames per second.

By comparison, VolumePro can display a 256^3-voxel dataset with 8-bit or 12-bit resolution at 30 frames per second in typical applications, performing 500 million trilinear interpolations per second with per-sample gradient-based Phong lighting.

This speed allows the user to rotate and resize the dataset at fully interactive speeds, making it easy to zoom in on details of interest and enhancing the natural ability of the human mind to make sense of complex 3D structures in a way that is simply impossible with slower solutions.

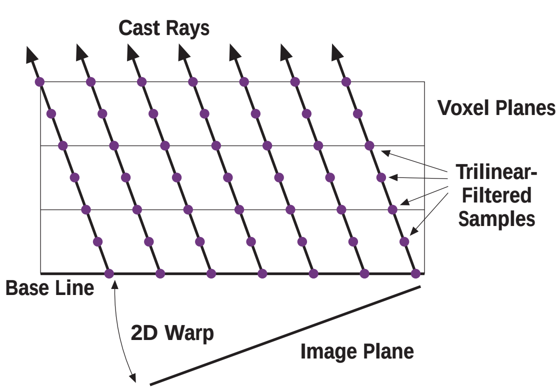

Figure 1. VolumePro uses a technique known as ray casting to generate 2D views of a 3D dataset. Rays cast from each pixel accumulate color and transparency as they traverse the dataset.

Pricing and Availability

In April 1999, Mitsubishi will begin volume production of VolumePro. It expects to offer VolumePro boards with 128M of local memory for $2,995 each. OEM and academic discounts will be available.

More information is available on Mitsubishi Electric Research Laboratory's volume-rendering Web site, *www.3dvolumegraphics.com*.

Mainstream Applications Still Far Away

Though MERL is evaluating mainstream opportunities for volume rendering, such as 3D games and business graphics, we believe it will be several years before volume rendering acceleration comes to mainstream PCs. Polygon-based 3D objects are adequate for current mainstream applications and are much more space-efficient than solid models.

MERL may have found for itself one of the few defensible niches in the 3D market. Of all the existing vendors of conventional 3D acceleration, only SGI and possibly HP appear to have both the resources and customers needed to justify the development of volume-rendering chips.

Over time, we expect MERL to produce more highly integrated VolumePro chips with more rendering pipelines, an internal display controller, and the ability to handle larger datasets. Even $1,024^3$-voxel datasets will be feasible by late 2000 at a cost within the reach of many professional users. Such a product would be able to match the visual quality of today's conventional 3D accelerators and open up even more markets to this extremely interesting technology. M

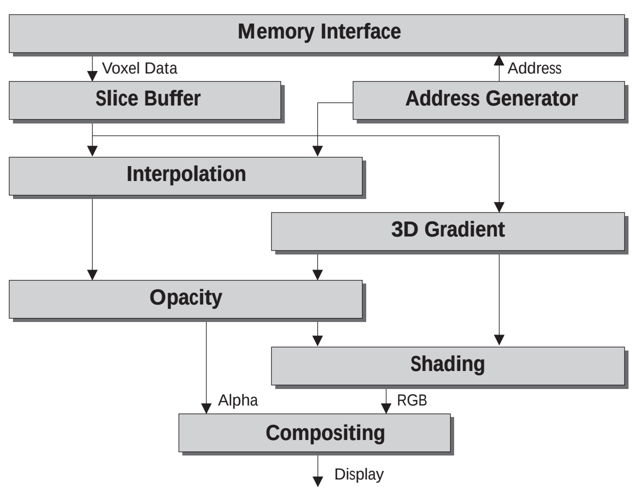

Figure 2. The VolumePro pipeline steps through the volume dataset along each ray, calculating transparency and color values. Finally, these values are combined to create display pixels.