FROM SEIKO, A 3-D GRAPHICS SYSTEM THAT'S 20% CHEAPER

It runs 400,000 vector transformations a second—speedy enough to achieve smooth real-time manipulation; fluid zooms, moves, rotations, and pans are all second nature

by Jonah McLeod

Designers no longer have to shell out up to $50,000 to buy a 3-dimensional color graphics-display system. Seiko Instruments U. S. A. Inc. is introducing its D-Scan GR-4416, which will sell for only $32,000—underselling its market rivals by 20% or more.

The San Jose, Calif., company trimmed costs by going, wherever possible, with VLSI chips rather than standard integrated circuits and discrete components. This decision paid off all along the system's graphics pipeline: a display control module was squeezed onto one board rather than three; a bit-map control module was cut down from three boards to one; and a frame buffer/video output module comes on two boards instead of three. All in all, Seiko cut five boards out of its system.

The Seiko system is fast! It whips through 400,000 vector transformations/s—speedy enough to achieve smooth real-time manipulation. Fluid zooms, moves, rotations, and pans are all second nature to the new product. Fitted with its full complement of 20 frame buffers, all of which can be double-buffered, the 4416 can paint a 1,280-by-1,024-pixel screen with 1,024 colors. At 640 by 512 pixels, it offers a palette of 16 million hues. With this kind of price-performance, Seiko engineers figure their new global product will meet the needs of designers in mechanical computer-aided design, electrical CAD and computer-aided engineering, and scientific modeling, as well as those people involved in medical imaging.

Accounting for the speed of the Seiko system is its pipeline, which has itself been streamlined. The matrix processor in the display control module, for example, is 30 times faster than its rivals. The bit-map control module is home to two data-differential-analyzer chips that share the work usually handled serially by a monolithic processor. Another two DDA chips generate 60 pixels every nanosecond. The module also houses a pair of 4-by-4-pixel buffers that write 60 million pixels/s, ensuring that images on the screen move smoothly. Finally, every module is fitted with a first-in, first-out buffer that helps move data quickly between modules.

Operation begins when a graphics image in the form of a command list ("draw line," "draw polygon," and so forth) is keyed in or sent from a minicomputer or host mainframe. The list first passes through the graphics control module (see fig. 2). A 12.5-MHz 68000 microprocessor within this module converts the command list into vectors and

1. COMPETITIVE. Seiko's GR-4400 series graphics system offers more performance at the same price as other high-end systems.

ships them out to the segment buffers. A design of less than 10,000 vectors can fit into the 0.5-Mbyte segment memory that is standard on the system, but designs typically call for 4 to 6 Mbytes and are stored on the system's optional hard disk.

Once in segment memory, the vectors are loaded into the graphics pipeline. The three modules that make up the pipeline execute the major functions that are associated with manipulating a 3-d image on the screen.

At the head of the pipeline is the display control module. It converts the vector from the coordinate system used within the confines of the CRT screen into world coordinates, a general coordinate system that is unrelated to the display coordinates. It then makes the required calculations on every vector to relocate the image to the position specified by the operator as he moves a mouse.

Inside this module are two processors. The display processor, a 32-bit-wide AMD 29203 bit slice, grabs 400,000 vectors each second and strips them of their display coordinates, assigning world coordinates in their place. The vectors are then shunted to the matrix processor, a 4-by-4-pixel unit assembled from four CMOS 8,000-gate arrays. The matrix-processing unit performs the calculations needed to rotate, position, and scale an image.

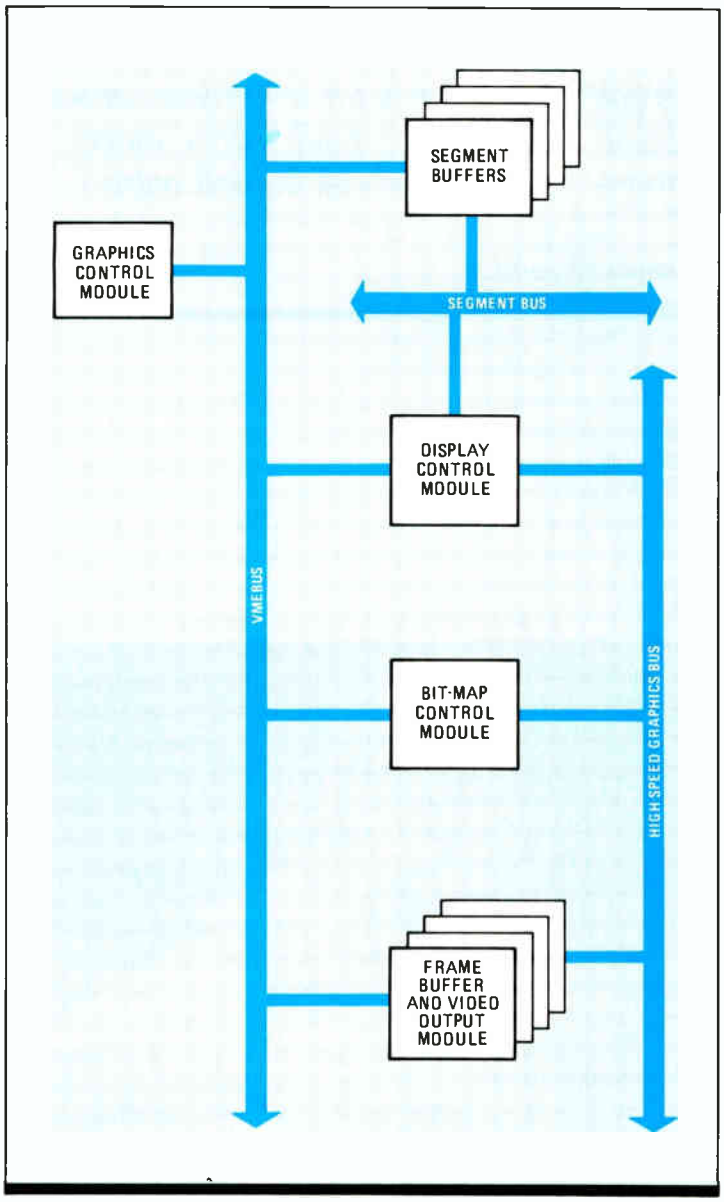

2. STREAMLINED. The GR-4400 graphics pipeline runs without bottlenecks from display control module to bit-map control module to frame buffer.

In many graphics systems, the matrix processor is itself a bottleneck, slowing things to the point that, out of the 400,000 vectors crammed into the system each second, only 100,000 can be handled. Usually, these units work through such trigonometric functions as sine and cosine. But Seiko has chosen to sacrifice floating-point capability for raw speed, designing the processor for matrix algebra only. As a result, it rips through 28 calculations—16 multiplications and 12 additions—in one instruction cycle. Competing systems take up to 28 cycles.

To make certain that the system doesn't bog down, each module is fitted with a D-scan graphics bus controller. The 900-gate array, which acts as a fast FIFO buffer, hustles the data onto and off of the graphics bus.

MAKING THE IMAGE FIT

Once the image has been mathematically positioned within the world coordinate system, the vectors are passed along to the bit-map control module. Here, another 32-bit 29203 clips the image to fit the viewing area on a CRT screen. Then the viewing transformation processor, an 8,000-gate array, converts world coordinates back into display coordinates.

From there, the image data is sent to the first of four DDA chips. The first pair determines the slope of each image vector to be drawn on the screen. The second determines which pixel to illuminate on a bit map of the CRT screen.

On other systems, these data functions are usually handled serially by a monolithic processor. By partitioning the function, the GR-4416 handles the data rapidly and efficiently.

The first data function is fielded by two CMOS 2,600-gate arrays. Each carries an arithmetic unit that calculates such image parameters as drawing direction, transformation values, and slope determination. The second pair of LSI devices, two 3,000-gate CMOS arrays, are 16-bit-wide pixel generators that use the data from the first two DDA chips to determine which pixels to light. They spend a scant 60 ns on a pixel. These four chips together can handle data at 2 ns per vector.

The second team of DDA chips feeds the first of two 4-by-4-pixel buffers, both on one 850-gate CMOS array. Once the first buffer is filled, it dumps its load into a 4-by-4-pixel section of the frame buffer while the second pixel buffer is being topped off. Once the second pixel buffer is loaded, it drops its data into another 4-by-4-section of the frame buffer, and so forth.

By alternating in this manner, the Seiko sys-

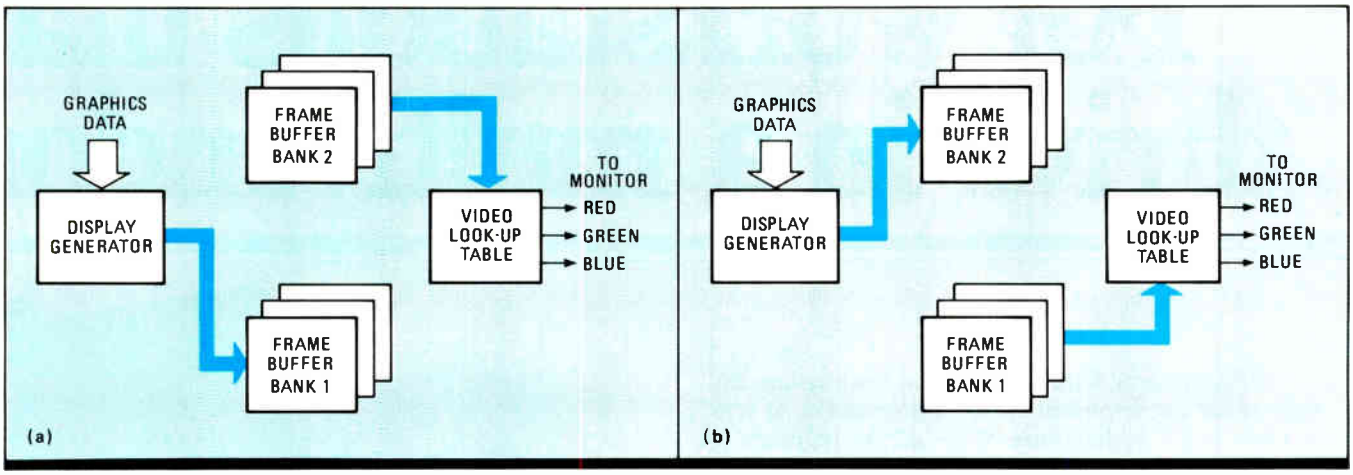

3. DOUBLING UP. With the double buffer on the GR-4416, one frame buffer is filled while a second writes to the monitor (a), then the second buffer is filled with the next image while the first writes (b).

tem makes it is possible to write 60 million pixels/s (54 ns/pixel) to the frame buffer, a 28% improvement over the system's closest competitors, which top out at about 40 million pixels/s (74 ns/pixel). In essence, the system fills the background buffer while the foreground buffer is displayed to the screen. That means an image moves smoothly, without the jerkiness of some real-time graphics systems.

In operation, the pixel buffers pass on vectors to five write-table chips at the mouth of the frame buffer/video output module. These five chips write the actual pixels into the frame buffer and handle other operations. Inside the module are two banks of frame buffers, each of which is four planes deep. Each of the banks can be expanded to 10 planes, for a total of 20 planes, which permits users to take advantage of double buffering. The expanded memory uses the same write and address controller that is standard with the system.

Some competitive systems also offer a double-buffer function by splitting a large memory into two parts. But this approach requires additional logic to coordinate which half of memory is writing to the monitor at any one time. More logic must coordinate the timing for writing into one part of memory from the graphics display generator and reading out to the monitor. And splitting the memory in half also halves the number of colors that can be displayed.

Double buffering ensures a smooth transition between successive images in applications in which a new screen must be written into a buffer by the end of the current vertical retrace interval of the video monitor. Breaking the frame memory into two banks makes it possible to alternately send one and then the other to the monitor for display. That means the display generator can write to its buffer without having to coordinate with the video look-up table, which can read from its own buffer (see fig. 3). Also, since each frame buffer has 10 planes, every plane can maintain a palette of 1,024 colors.

With the high-performance pipeline and double-buffering capability, the GR-4416 can generate graphics images at a rate of 40 frames per second. Movie film is displayed at a rate of 16 frames per second. Such innovative circuit designs as the double buffer, fast DDA pixel buffer and data-differential analyzers, and matrix processor give the GR-4416 a price/performance edge in an increasingly competitive market. □

For more information, circle 484 on the reader service card.

HOW SEIKO GOT ITS PRICE-PERFORMANCE JUMP

In the summer of 1985 the world marketing group for Seiko Instruments Inc. got together at the company headquarters in Tokyo and concluded that the company needed a line of graphics-terminal products if it was to be a full-line supplier of graphics products worldwide. "We had to get ahead of competitors in a marketplace where price/performance was changing on a yearly basis," says Andrew Wei, division manager for the Graphics Devices and Systems Division of Seiko Instruments U.S.A. Inc.

Seiko put together a team that eventually numbered more than 100 hardware and software engineers. Working with Seiko groups in the U.S. and Europe, they developed the D-Scan GR-4416, the company's first high-end display system.

They were able to get the speed up to 400,000 transformations/s and keep the cost 20% below that of competing systems by the judicious use of custom-designed CMOS VLSI chips. "Our corporate management in Tokyo had the foresight to invest in developing 1.2-μm CMOS processing technology inside the company," says Wei.

The world marketing group's contribution was to define a product that would meet global market needs. They pinpointed a three-dimensional system that could provide smooth, fast interaction at a relatively low price. And, Wei says, "we saw a very large market for the kind of price/performance we had in mind."

ANDREW WEI