INTERGRAPH LINE GETS ZIPPY ENGINE, GRAPHICS

Intergraph Corp. is out to save a rapidly growing list of designers a big chunk of money—those people who increasingly need both task acceleration and high-speed graphics-display capabilities in their work stations. The Huntsville, Ala., company is upgrading its work-station line by adding two boards: a floating-point engine that accelerates such tasks as analog simulation, printed-circuit-board layout, and mechanical-solids modeling; and a graphics processor for state-of-the-art graphics display.

The new line starts at less than $30,000. Up to now, designers had to pay a lot more for that combination: a dedicated accelerator would cost anywhere from $40,000 to $250,000, depending on the task, and then it would have to be hooked up to a work station with a high-speed graphics display.

The new work stations serve a broad range of markets: mechanical, electrical, and architectural computer-aided design, and scientific computation. They are expected to be available in September.

The work stations (see fig. 1) retain the names of the three existing Intergraph offerings—InterPro, InterAct, and InterView—but the three now each come in two different versions, the 340 and 360, which have varying amounts of main memory and disk space. The InterPro is for general-purpose use; the InterAct is geared toward mechanical CAD; and the InterView is aimed at mapping. The floating-point engine and graphics processor board (see fig. 2) are what set the new line apart from the previous generation.

The floating-point engine addresses the need for high-performance point acceleration of a variety of diverse tasks, including analog simulation and mechanical modeling. Electrical-engineering point accelerators cannot be used on a mechanical design problem because they are too specialized.

The Intergraph engine can accelerate tasks up to an order of magnitude faster than the 5-million-instruction/s Fairchild Clipper central-processing unit in the Intergraph work station. The engine has a single-instruction, multiple-data-path architecture that can execute as many as five operations concurrently. It has both a dedicated floating-point processor and an integer processor.

The graphics processor board can handle a variety of different applications simultaneously, because it comes equipped with a 1-Mbyte writable control store. Any application—analog simulation, pc-board layout, mechanical-solids modeling—can be programmed into the control store, and calculations can be accelerated at the full speed of the accelerator. The writable control store uses static column random-access memory (SCRAM), where column-address selection is made without the additional address-strobe clock that conventional static RAMs use. The result is static-RAM performance at dynamic-RAM costs.

To address the problem of graphics displays, the company designed its graphics processor with the same single-instruction, multiple-data-path architecture as the floating-point engine. It placed the shading and polygon-fill operations into a 1-Mbyte writable control store on the graphics board to allow real-time wireframe manipulation and shading and polygon fills in less than a second. A special emitter-coupled-logic gate and three additional chips implement the control functions of the board's dual frame buffer, which provides fast screen update.

Both work stations have eight slots, but only four in each are required for a complete system with 16 Mbytes of memory. Two of the remaining slots can be used for additional 32-Mbyte boards of memory, for a total of 80 Mbytes. The wide-band input/output processor offers data throughput three times faster than the previous version. The board uses an 80386 I/O processor board with 2 Mbytes of its own on-board RAM.

The floating-point engine board is based on the Weitek 1066 32-by-64-bit register file and 2264 and 2265 floating-point multiplier/accumulator and multiplier chips. The engine has a peak throughput of 25 double-precision megaflops. It has a 32-bit integer processor and its own 8

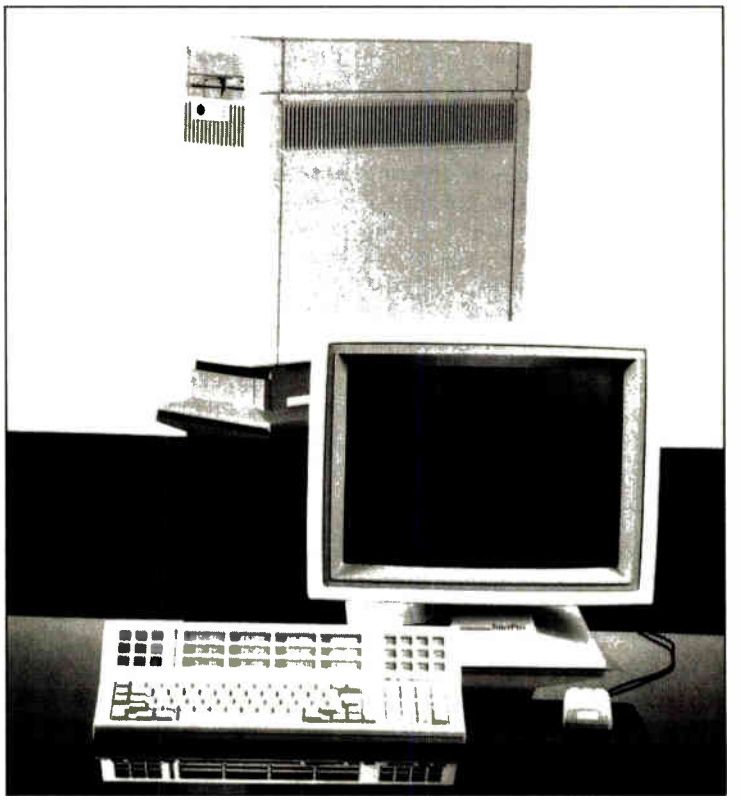

1. QUICK. With its 5-mips Fairchild Clipper chip and new 25-megaflops floating-point engine, Intergraph's work station offers blazing speed.

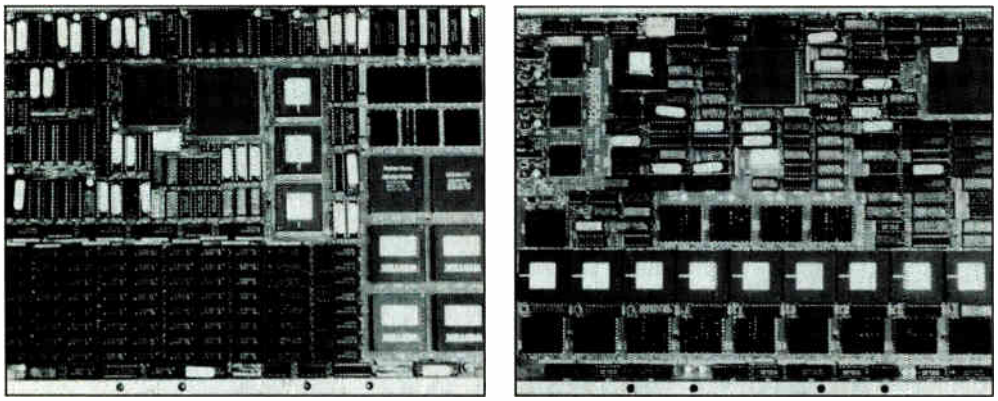

2. POWERFUL. Both the floating-point engine (left) and the graphics processor board use a single-instruction, multiple-data-path architecture to accelerate their processing tasks.

Mbytes of data storage that is separate from other memory in the system. The processor can access any word from the memory with an 80-ns cycle time. Data and instruction transfer in the system is over a high-speed, 64-bit data path. A program sequencer, which accesses 128-bit-wide instructions from the 1-Mbyte writable control store, directs the board's operations.

Each of these operations is specified by individual fields inside the 128-bit instruction word. The operations include multiply and arithmetic logic unit—add, subtract, and compare; floating point, fixed-to-floating-point, and floating-point-to-fixed conversions; loading and storing operations from the processor board's 32-by-32-bit register file to main memory; integer operations performed in the separate 32-bit integer ALU onboard; and branch-control operations.

To speed up operations, users can determine the location of performance bottlenecks in their particular application. They can microcode the analog simulation program into the writable control store of the floating-point engine board and accelerate the operation of the program tenfold.

"We see the engine as a generally programmable supermicrocomputer that can be specialized for any floating-point-type application," says Bruce Imsand, vice president of systems development. The company will provide a standard set of applications, but if the user has an application for which no programs exist, "I have a group of programmers whose job is to convert the user's C and Fortran programs and subroutines into code for the writable control store."

With a megabyte of writable control store, more than one application (such as an analog-circuit simulator) and smaller subroutines (such as fast Fourier transforms) can use the floating-point engine concurrently. Because more than one library or applications program can be loaded into the writable control store, it is not necessary to unload one application to load another. This is in contrast to other implementations of writable control stores, where the system must overlay the previous application to run the current application.

The floating-point engine uses 1 Mbit of SCRAM, which offers nearly the speed of SRAM at the price and power consumption of DRAM. The 1-Mbit SCRAM on the floating-point engine affords the system an 80-ns total memory-cycle time for data accesses within the current page of memory. There is a 96% chance that the next memory access will be in the current page of memory. The SCRAM cycle time can be as low as 60 ns, but this application only requires 80 ns. If the memory access is in the next page, then the access requires another clock cycle to establish another row address.

The graphics processor board is also a single-instruction, multiple-path data processor. To achieve its high-performance color capability, the board contains nine 32-bit data-path processors, each with its own megabyte of control-store memory using SCRAM chips, just like the floating-point engine. All the microcode for performing graphics functions is in the control store.

Shading and polygon-fill operations tend to slow down most graphics processors. The board can handle polygons of any description, rectangular fills, patterning, vector generation of all descriptions, straight shading, or Gouraud shading—a well-defined color-interpolation technique that determines in an image what the interim colors are along a scan-line path on a surface. The board can display 512 active colors from a palette of 16 million and can redraw an image on the screen at a rate of more than 100,000 vectors/s.

Each of the nine data-path processors controls one of the nine video planes of memory to provide the 512 colors out of the palette. Each data-path processor is an enhanced 32-bit bit-slice processor with a 32-bit funnel shifter, video FIFO, and interfaces to two 32-bit buffers—one for its private memory cluster, and the other to a bus that runs throughout the board and with which all 32-bit data-path processors connect. An ECL gate array on the board contains all the logic to move the signal from the nine video planes to the digital-to-analog converters that will drive the video display. The three palette chips—red, green, and blue—that drive the video-display monitor each contain an 8-bit DAC and a 512-by-8-bit lookup table. These parts were developed jointly by Intergraph and Honeywell and are only now becoming commercially available to other graphics work-station makers.

The graphics processor board contains two sets of nine video planes. One is being displayed while the second is being updated. With simple graphics images, 10 frames/s can be produced. A more complex architectural diagram might produce only two frames/s. *–Jonah McLeod*

For more information, circle 481 on the reader service card.