"A New Style of Computing"

Graphics Supercomputers in Meteorology

S D Pass
Stellar (UK) Limited

1. Introduction

Very high performance workstations combining significant computational performance with 3D graphics have now arrived in the market place. These machines offer scientists and engineers an opportunity to break away from the traditional numerical simulation regime of batch computation, probably on a mainframe or supercomputer, followed by post mortem analysis on a graphics workstation, and adopt a more interactive approach with consequent benefits to productivity.

At the forefront of the super-workstation market is the Stellar GS1000 Graphics Supercomputer, a system which integrates 40 MFLOPS of mini-supercomputer performance with exceptional 3D graphics in a single-user package. Such capabilities now make it possible both to simulate a complex time-dependent system and visualize it as it evolves, seeing things as they happen. Moreover, coupling simulation and visualization as closely as this opens the possibility of bringing user interaction directly within the simulation-analysis loop to aid understanding and shorten the time spent within this cycle.

The GS1000 has only been available for a matter of months, but already it has found its way into a wide range of applications where the combination of computational performance, 3D graphics and interactivity is paramount. Not least of these application areas is meteorology where, in the USA, Stellar is working with three different research groups. Indeed, at the recent Cray Symposium in Minneapolis on "Science and Engineering on Cray Supercomputers", a GS1000 was demonstrated interactively visualizing a tornadic storm simulation. More will be said of this later.

2. The Stellar GS1000 Graphics Supercomputer

The GS1000 was designed as a balanced system, offering the right blend of computational performance, graphics performance, memory capacity and I/O performance. From the outset it was obvious that the desired level of simultaneous computational and graphical performance would only be reached if data could be moved about the system extremely rapidly. Stellar's innovative solution to this problem is the DataPath, a set of fast, dedicated pathways which directly interconnects the system's functional units and avoids the inevitable problem of bus contention often experienced in more traditional architectures. Bandwidth into the 1 MByte cache is 1.28 GByte/s, into main memory, 320 MByte/s, and when all pathways are working simultaneously data is being moved around the system at over 3 GByte/s.

The system architecture is shown in figure 1. The Multi-Stream Processor (MSP) can execute four instruction streams concurrently, each at over 6 MIPS, for an aggregate performance of up to 25 MIPS. Stellar Fortran and C compilers will automatically parallelise an application into tightly coupled "threads" (sub-processes) which can use several streams concurrently. In this way, a single application can call upon the whole 25 MIPS when needed. The four streams share cache and main memory to minimise communication overheads and eliminate the problems of cache coherency, and are additionally equipped with concurrency registers for rapid inter-stream synchronisation.

Figure 1. Stellar Graphics Supercomputer Architecture

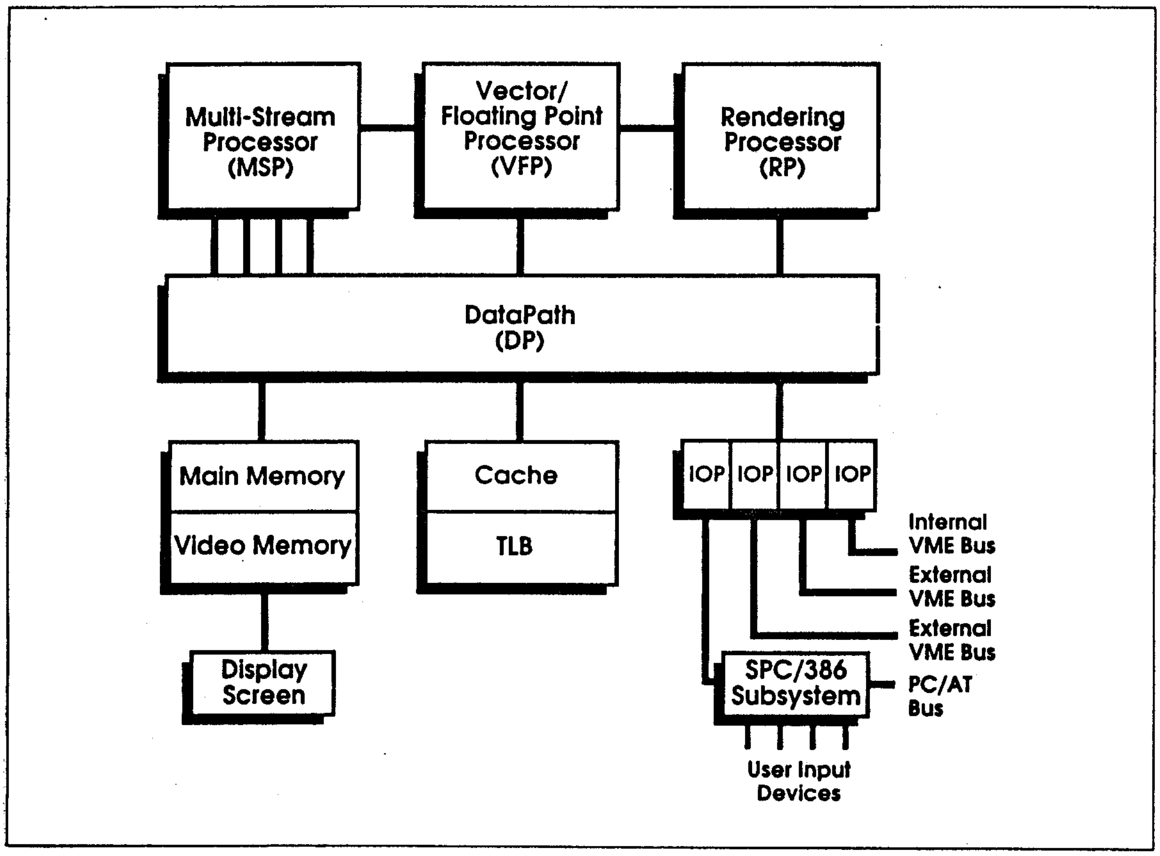

The Vector/Floating Point Processor (VFP) offers a peak performance of 40 MFLOPS (double precision) and can sustain about 34 MFLOPS. It services requests from the MSP for both scalar and vector floating point operations which are performed on registers previously loaded by the MSP. Whilst the floating point calculation is being performed the MSP can operate as normal, the two units working simultaneously.

As with parallelisation, the Stellar Fortran and C compilers will automatically vectorise an application to get the most out of the VFP. Additional hardware support is provided for division and square root to give up to 10 million reciprocals, 2 million square roots and 1 million sines or cosines per second. The VFP also plays an important role in graphics processing, handling coordinate transformations and other associated calculations for the rendering processor.

Up to four I/O Processors (IOPs) handle communication with external controllers and peripheral devices. Each processor has a peak transfer rate of 16 MByte/s. Two are used internally to connect the DataPath to industry standard VME and PC/AT buses. The VME bus supports up to 3 GBytes of disk internally within the system and more than 20 GBytes can be added in expansion cabinets using the other two IOPs. For very high performance a file system can be striped across up to sixteen disk drives, yielding something approaching sixteen times the throughput of a file system on a single drive.

The PC/AT bus connects to the service processor, a PC/386, which performs power-on self-test, boots the system and additionally handles low speed devices such as the keyboard, mouse and serial ports when the system is running. The service processor is also equipped with an extensive suite of diagnostic software to be used by the customer service engineer, either on site or remotely via a modem connection.

The Rendering Processor (RP) draws graphical primitives directly into main system memory, into Virtual Pixel Maps, rather than a frame buffer, as in more traditional graphics systems. The Virtual Pixel Map concept is unique to Stellar and offers several advantages in the visualization environment. Images of almost any resolution and dimension can be created, irrespective of the frame buffer size and monitor resolution, for subsequent display through the windowing system. Smooth scrolling and real-time animation are achieved by rapidly copying data from main memory to the frame buffer, and here again, the bandwidth of the DataPath plays a key role since it provides a pathway through which pixels (32-bit) can be moved at the rate of 80 million per second.

The RP is itself a parallel processor with an Single Instruction, Multiple Data stream (SIMD) architecture. Sixteen identical "toe" processors arranged as a 4x4 "footprint" operate on their own pixel data under the control of a common instruction stream. Each toe processor is in fact a 32-bit integer processor operating at 20 MIPS, giving the RP an aggregate performance of some 320 MIPS. The MSP, VFP and RP together form a graphics pipeline. Graphics instructions are passed from the MSP to the VFP where 3D coordinate transforms, clipping and lighting calculations are performed on the graphical primitives. Packets of graphical data are in turn passed from VFP to RP which "walks" across the primitive within a Virtual Pixel Map, performing shading, hidden surface removal, depth cueing and even transparency as it goes. The graphical draw rates achieved by the GS1000 are outstanding, reaching for example 600,000 3D, 10-pixel lines per second and 150,000 100-pixel, Gouraud shaded, hidden-surface-removed triangles per second.

The RP is a relatively general purpose processor unit and has been micro-programmed by Stellar to perform a number of additional tasks, including Phong shading, drawing spheres directly, and anti-aliasing lines and polygons.

Despite this range of built-in rendering facilities, it is likely that in some visualization applications the desired image can only be created with additional special software techniques. The Virtual Pixel Map offers a real advantage here since it allows the application to easily mix clever software rendering with the available firmware techniques without incurring the large overhead of shuffling pixel data around the system.

Sometimes it may not be possible to render an animated visualization sequence in real-time, particularly if clever software techniques are being employed. However, the sequence can be stored in a set of Virtual Pixel Maps and replayed in real-time as a "flip book" animation. With up to 128 MBytes of main memory a lot of image data can be held ready for immediate access (and any more than that would be pushed out to disk backing store by the virtual memory system). Traditional graphics workstations just do not offer this kind of facility since the frame buffer can only hold one or two screens worth of information, and the bandwidth from main memory to the frame buffer is usually relatively low, perhaps just a few MBytes per second.

The GS1000 has a very innovative hardware architecture which could only be realized as a practical system by making extensive use of VLSI technology. No less than 11 different ASICs were designed by Stellar, and each GS1000 contains 61 of them. Needless to say, some ASICs are used more than once within the system: the DataPath is itself implemented with 32 identical chips, and as has been indicated before, the RP contains 16 identical devices. All in all, over two million gates make up the logic of the GS1000.

Despite the innovation in hardware, Stellar has committed itself to a standard software environment. The operating system, Stellix, is based on Unix System V.3 and is SVID compatible. It has been modified by Stellar to handle concurrent processing (threads), file system performance has been improved considerably, and networking facilities have been added. These include TCP/IP and Sun's NFS distributed file system which has become a de facto industry standard.

Stellar has adopted the MIT X window system and PHIGS+ as its graphics standards. The X system offers a windowing environment for both local and distributed graphical computation. PHIGS+ is an extended version of the Programmer's Hierarchical Interactive Graphics System developed by a group of over twenty hardware suppliers, application developers and end users. It is a 3D graphics system which manages the creation, storage and manipulation of complex graphical objects and which allows sophisticated applications to be created which are independent of a particular hardware system. Along with wire frame modelling, PHIGS+ supports sophisticated solid modelling techniques including lighting, shading, hidden surface removal, depth cueing, and B-spline and Bezier curves and surfaces.

3. The GS1000 in Meteorology

In the USA Stellar has forged close links with two groups which are very active in the use of graphics in meteorology, and a third will take delivery of a GS1000 shortly.

Rob Wilhelmson at the Department of Atmospheric Sciences and National Centre for Supercomputing Applications, University of Illinois is well known for his work on the numerical simulation of severe storms. He has recently collaborated with Stellar both in porting his storm simulation code to the GS1000 and in developing an interactive visualization package. Lou Wicker actually did the port and found it a particularly painless task, the Stellar compiler providing extensive vectorisation and parallelisation automatically. The GS1000 delivers a sustained 15 to 18 MFLOPS to the simulation, which is just one tenth of that delivered to the problem by a CRAY-2 (but at much less than one tenth of the price).

The visualization package was written by Dave Kamins of Stellar and is directed towards visualizing the 3D wind velocity field of the storm. It explores the use of particle advection whereby a cloud of passive particles is injected into the system which then moves in accord with the wind vectors. The particles are drawn as shaded spheres to enhance the 3D effect, a task to which the GS1000 is ideally suited since it can draw spheres at up to 50,000 per second.

At a recent Cray user symposium, held 12-14 October in Minneapolis, Wilhelmson gave a live demonstration of interactive visualization as he presented his paper. A Stellar GS1000 was networked to a CRAY-2 via a T1 link (1.5 Mbit/s) and as the

CRAY machine performed the simulation, the GS1000 presented the current state of the storm system to the audience. It is interesting to note that the GS1000 could actually perform the visualization more quickly when it ran the simulation itself: the T1 link proved to be a significant bottleneck, even though a lot of work was done to minimise the amount of data that had to be passed between the two machines. This is an excellent example of the benefit that can be gained when compute and graphics are tightly coupled in a single package.

The Space, Science and Engineering Centre (SSEC) at the University of Wisconsin is well known for its McIDAS weather data visualization system, and has over 80 such turn-key systems installed in customer sites. McIDAS lets the meteorologist create moving 3D weather displays using perspective, stereo, motion parallax, depth precedence, shading, texture, brightness, colour, shadows, and transparency. After a good deal of careful thought Bill Hibbard chose the GS1000 as the SSEC's "interactive 4D workstation", and even though the machine was delivered just a month ago Hibbard's group has already ported an extensive suite of visualization software.

For a more detailed account of the work now being done on the GS1000 by the SSEC group please refer to Bill Hibbard's paper, "Visualizing Weather Data", presented earlier at this workshop.

Bill Holland at the National Centre for Atmospheric Research (NCAR) in Boulder, Colorado has recently selected a Stellar GS1000 for his work on advanced visualization and compuational methods in ocean circulation modelling, part of NCAR's work on the dynamics of the global climate. He will be taking delivery of his machine shortly.

Trademarks