

THE GRAPHICS SUPERCOMPUTER: A NEW CLASS OF COMPUTER

Gordon BELL and William S. WORLEY Jr.

Arden Computer Company,
Sunnyvale, California, U.S.A.

Invited Paper

In 1988 a new class of computer, the graphics supercomputer, was introduced by two start-up companies. As a member of both the supercomputer and workstation families, the graphics supercomputer enables both high-performance computation and high-speed, three-dimensional, interactive visualization to be performed on a single, integrated system. In Arden's TITAN system, high floating point operation rates are provided by combining the fastest RISC microprocessors with traditional supercomputer components such as multiple, pipe-lined vector processors and interleaved memory systems. This has resulted in much more cost-effective computation and a clear dis-economy of scale over supercomputers and minisupercomputers. For users of workstations that lack the benefits of supercomputer components, the graphics supercomputer can be used simply as a workstation which provides five-to-ten times more performance for important classes of applications.

Arden has introduced a next generation, extendible, object-oriented dynamic rendering environment, Doré, as a de-facto standard for the graphics supercomputer computing paradigm, Borden [1]. The graphics supercomputer enables new applications in computational chemistry, computational fluid dynamics, real time simulation, animation, mechanical design and engineering, interactive visualization, and image processing as typified in medicine, surveillance, and petroleum exploration.

Machine Classes

Supercomputers, mainframes, minicomputers, super-mini-computers, mini-supercomputers, workstations, and personal computers have formed and evolved as distinct computer classes, differentiated primarily by price. Each new class that formed has been based upon key new technologies; a market of users with corresponding application and style-of-computing needs; and start-up companies who built and introduced the new type of computer. While the record shows that established companies do not innovate fundamentally new products, they frequently are a source of key people and serve as initial market outlets for those companies that do innovate. Ultimately, if the market is a large one, they adopt the new class idea from a start-up. Examples include IBM's adoptions of the minicomputer from DEC and the PC from Apple, and DEC's recent adoption of engineering workstations from Apollo and Sun.

Figure 1 shows how various members of established computer classes have occurred in time, following two "main lines" of development:

- general purpose processing for both commercial and scientific use, as characterized by the thread of development from the Univac I "mainframe" in 1950, to the IBM S/360, S/370 ... 3090 evolution (general complex instruction-set, large memory, microprogrammed, cache memory, virtual memory operation, multiprocessors for multiprogrammed use, and, finally, vector processing); to the VAX "minicomputer"; and to the truly interactive "personal computers" and "workstations."
- peak computational power as measured in millions or billions of floating point operations per second for scientific use. This line has been characterized by the Seymour Cray designs which employ the fastest clocks, hardwired logic, and extensive parallelism through pipelining and multiple functional units. These systems evolved to vector processing, multiple processors for multiprogrammed use, and finally to the parallel processing of a single computation. Lower cost supercomputer lines based on extensive parallelism were extended in the early 80's with the "minisupercomputer" by Alliant and Convex, and with the introduction this year of the "graphics supercomputer" by Arden and Stellar.

The importance of the UNIX operating system to the evolution of computing should be emphasized. It has permitted users of one class of system to have compatibility with machines of other classes of systems. We believe that both horizontal (machines within a class) and vertical (machines of different classes) compatibility are essential for accelerating the growth of computer use.

How Is The Graphics Supercomputer Used?

Unless a new class of computer engenders a new type, style, logistics, or economics of use which differentiates it from other classes of computers, it is unlikely to be successful. Although the graphics supercomputer has been on the market for less than one year, radical new uses are emerging in science, engineering, medicine, finance, arts, training (real time simulators), and intelligence/surveillance based on coupled high computation and interactive visualization capabilities.

At the same time, many of the uses of a graphics supercomputer are simply evolutions of the uses of structures that have been in operation on ordinary workstations. In effect, a graphics supercomputer can be viewed as a very high performance workstation—the quantitative differences in its performance capabilities result in qualitative differences in the applications and manner in which it is used. In this regard, compatibility at many levels is essential for training and the preservation of software investments. This includes common hardware busses, adapters, and protocols; Local Area networks; data formats; programming languages and compilers; operating systems, including the system, windowing, and user interfaces; graphics languages; and end-user applications.

EVOLUTION OF COMPUTING STYLES

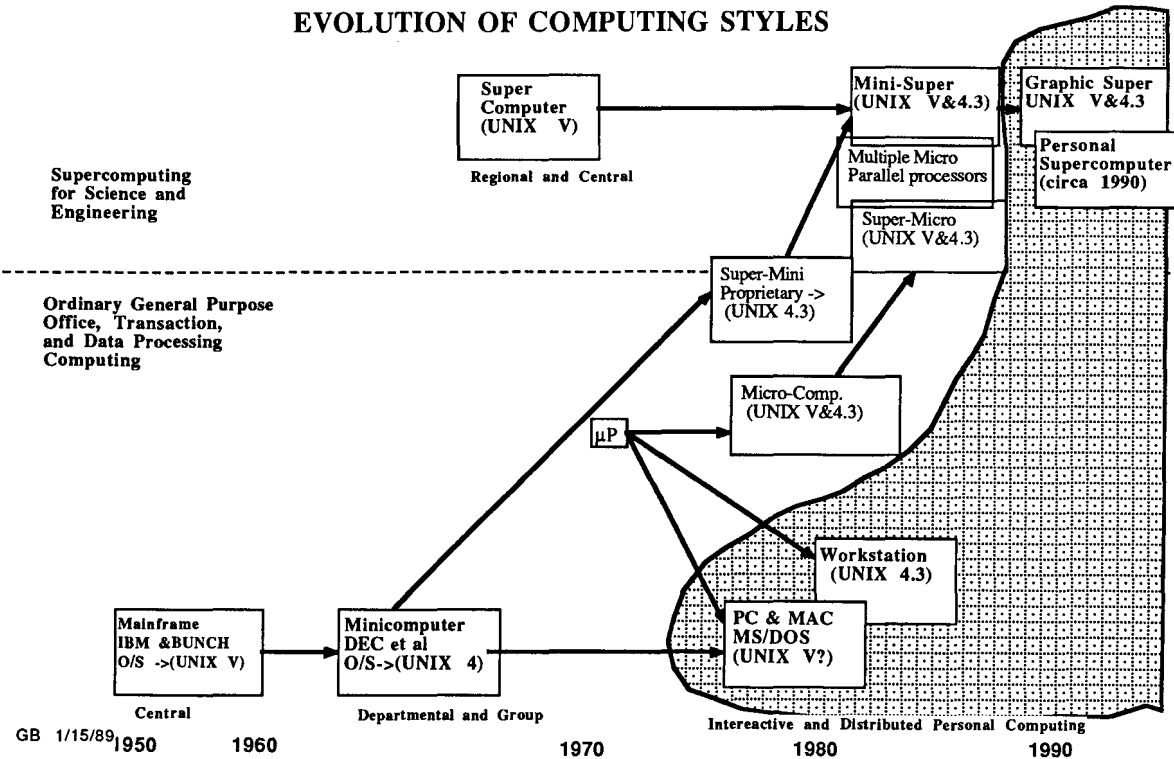


Figure 1

In other cases, the graphics supercomputer simply can substitute for an ordinary computer of less power, such as a super-minicomputer (e.g. VAX 8600), a mini-supercomputer with equivalent power and higher price (e.g. Alliant or Convex), or a mainframe (where the performance is often an order of magnitude worse than a supercomputer). In some instances, the graphics supercomputer is used together with a supercomputer, as a workstation for visualization and for ancillary high speed computations.

The most interesting and perhaps most important aspect of the graphics supercomputer is that it is the *avant garde* for personal supercomputing. We believe that the productivity of scientific researchers can significantly be augmented by the availability to every scientist and engineer of his own, personal supercomputer. A graphics supercomputer of the power of the Ardent Titan System, which is usually sufficient to substitute for a share of a central supercomputer, at a cost below \$50,000, could make this possible. The next generation graphics supercomputers are expected to reach this performance/price, enabling the transition to full distributed supercomputing where every scientist and engineer has one on a desk, instead of timesharing a large, single system.

New Technologies Have Enabled the Graphics Supercomputer Class to Form

Each of the computer classes cited earlier was formed initially, and evolved subsequently, based on new circuit, and software technologies. The large scale integrated circuit is the basic hardware technology that has enabled the construc-

tion of graphics supercomputers. Very large scale integrated CMOS circuits can be used to build both large and fast memories; fast, simple, pipelined scalar processors (i.e. RISC microprocessors); pipelined floating-point arithmetic units; and large, complex application-specific ICs. Large gate arrays provide the control for the complex vector units and graphics processors which compose the heart of the graphics supercomputer. The main architectural ingredients of graphics supercomputers are RISC microprocessors, vector processors, and specialized graphics processors. Components realizing each of these architectural elements became possible due to large scale integrated circuits.

Few now debate the pro's and con's of RISC architectures, the RISC revolution is over. RISC won. Today, RISC microprocessors set the performance/price standards for workstations, technical computation, and even some commercial systems. They simply cannot be equaled by CISC microprocessors, mainframes, or minicomputers.

The power of vector processors derives from the vector paradigm that is very generally applicable and equivalent to concurrent issue and execution of many scalar instructions. The Ardent Titan vector processor, for example, simultaneously executes two load vector register pipelines (pipes), one store pipe, and a one-to-three input operation pipe. Each completed pipe operation includes executing its main operation, advancing register and memory addresses, decrementing a vector length, testing for completion, and continuing operation until complete. In sequential form, if each of these elements were done by a single instruction, and

assuming self-address-incrementing load/store instructions, each Titan vector unit cycle would require approximately sixteen scalar instructions. In effect, specialized vector hardware permits sustained, pipelined execution of sequences of highly generic operations. The only fundamental differences between a vector unit and a more specialized pipelined hardware component, such as a graphics pipeline, are that though the vector unit may not execute fully as rapidly on the particular, specialized task, it can provide its high execution rate to a much, much broader class of problems. As a result, workstations which provide both scalar capabilities and vector capabilities will far outperform those without vector capabilities for important classes of problems, Savage, et al [2]. It also has been found that even within specialized areas, such as traditional line and polygon graphics, the Titan vector unit is a practical tool not only for basic operations, but also is more readily adaptable to emerging forms of image and combined image/polygon processing than is the case for more specialized graphics hardware. At the pixel level, graphics processing requires parallel computation and high bandwidth. Multiple, specialized processors and partitioned memory organizations provide the additional means to sustain the required rendering and drawing rates.

The most important enabling software technology for the graphics supercomputer is that employed by a vectorizing and parallelizing compiler. This technology emerged from supercomputer and mini-supercomputer developments, and is now in widespread use.

Providing and exploiting parallelism is the major challenge both to the designers and to the users of the graphics supercomputer. In a graphics supercomputer environment, one can observe the following forms of parallelism:

<u>Technique</u>	<u>Support and Use</u>
multiple computers	LAN interconnection, message-passing.
multiprogramming	multiple jobs running on multiple processors.
multi-tasking	multiple processes composing a job, message-passing.
multi-threads	micro-tasking a job by multiple, cooperating processors. (code generated by a parallelizing compiler)
co-execution	integer, scalar, vector-ops scheduled for concurrent operation.
vector processing	multiple operations with load, store, execute, pipelines and chaining.
RISC	multiple vector processors.
rasterization	pipelining of multiple integer and scalar operations.
concurrent I/O	multiple pixel and polygon processors to transform polygons into shaded images.
	multiple, independent, concurrent I/O streams.

In all but the case of explicitly coded parallelism through message passing, the parallelism is transparent. Ideally, users take existing Fortran and C programs written for traditional, scalar computers, and simply recompile and run them on vector multi-processors of the type which characterizes today's "mainline" supercomputers.

Applications Needs are Critical for Market Formation

While technology lies at the root of a new computer class, corresponding applications' needs must exist and be in balance in order for a market to develop which will sustain the new class. These needs for the the graphics supercomputer class are two-fold:

- Solving the computational/visualization power paradox to unlock new use. Regional and central supercomputers, capable of supplying a large number of floating point operations per second, have no visual capabilities. Workstations and personal computers have visual ability, but little computational ability. Modern scientific and engineering users require both for a variety of existing and new applications. Given the limited capability of today's wide area and local area networks, graphics must be an integrated part of the computer. Note that entirely new applications emerged with the PC and workstation because of the intrinsic interactive capability. This occurred particularly for office automation with word processors and spreadsheets. The greatest change enabled by the graphics supercomputer will be the use of the computer as an interactive aid in many domains of iterative analysis and design. Specifically:

Computer Aided x in domain y , i.e. $yCAx$. Where

$x = \{ \text{Design, Engineering, Software Engineering, Publishing} \}$ and
 $y = \{ \text{Electrical, Mechanical, Chemical} \}$.

The changes which will aid scientists in discovery already have been equally pronounced.

- Computers in one class may substitute for computers in another class provided the new class has price, solution time, and ease of use (e.g. low cost to buy, operate, and distribute) advantages. Any new computer class capable of supplying power at a significantly better level of performance (supers allow pioneering new applications requiring computation), or performance/price or price (provided the price and solution time are acceptable) level has a market proportional to the advantage.

Interactive Visualization, the new paradigm for supercomputing

The first need above is equivalent to eliminating the impediment of not being able to visualize and control results at the prodigious rates at which current supercomputers generate results. Even with the emerging faster local area networked environments, the network is a significant impediment to use—one simply can not generate animated results in real time. This requires users to make movies in order to view results. This is actually a reversion to the batch processing days, circa 1970, when users computed the results and then went off-line to look at them. Visualization goes beyond the ability to have a system that looks through a large database, plotting results (including 3D images) in interesting ways. It must include the ability to interact dynamically with a computation, both for purposes of perceiving what's going on and of directing the process.

The new model provided by a graphics supercomputer is that of interactive visualization. Results are viewed concurrent with their generation. This takes place either solely within the graphics supercomputer, or with a tight coupling to a supercomputer for problems requiring higher computation rates.

Substitution for improved performance or performance/price.

To understand the advantage to a user of seeking better performance or performance/price by using a graphics supercomputer as a substitute for a computer of another class, we can look at what each of the machine classes provides. Figure 2 plots key machines in the classes using the LINPACK (100 x 100) benchmark. The art of deciding on an appropriate benchmark to characterize a computer is well beyond the scope of this paper. A user can obtain compare the results of various simple, stand-alone benchmarks, but actual workload results may bear no correspondence to any of them. We use LINPACK (100 x 100) because it happens to correspond to the average speed that supercomputer installations operate. Similarly, it is strongly correlated with a number of programs based on linear algebra such as ANSYS, Nastran, and Patran. Note that the positioning in the leading performance/price band of Figure 2 in effect ranks machines based upon whether they have vector units or not. Note also the clustering of classes of machines in distinct performance/price regions.

Table 1 gives the purchase price, performance, and performance/price for several benchmarks, run both sequentially and in parallel, for a Cray Y/MP and a Titan Graphics Supercomputer. Observe that for the purchase price of the Y/MP, one could buy 166 distributed graphics supercomputers for

personal, project, or departmental use. This simple model ignores operating costs, which in the case of a central, shared computer are quite large and quite visible. In the distributed approach operations costs (e.g. file backup) can be buried in the organization. Similarly, the costs of support, software acquisition and maintenance, and file maintenance also vary between the two approaches.

Obviously, each of the benchmarks runs for a longer time on the slower machine. The "stretch factor" is the ratio of the time a program runs using the Titan to the time the same program runs on the Cray YMP. Also, associated with each benchmark is the cost-effectiveness or performance/price (e.g. megaflops/sec/\$) of the YMP versus the Titan.

The range of results are comparable with an analysis of Titan and the Cray X/MP by Misaki and Lue at the Institute for Supercomputer Research, in which, for scalar and vector loops, the Cray was approximately a factor of five-times and ten-times faster respectively. The Whetstone benchmark is indicative of such use. For simple, integer-oriented benchmarks, such as those encountered during editing, compiling, and running operating systems, the YMP offers little advantage, since it is about the same speed as the Titan. This indicates that although a YMP is still faster for many utility programs, it is less cost-effective by almost two orders of magnitude. Sharing a YMP for such work makes no economic sense.

For a single program, it takes 12 times longer to complete the same amount of work using a distributed graphics supercomputer. But the distributed approach is almost three times more cost effective. In principle, users spending the same amount of money could buy three times as much computing. At a typical supercomputer center very large projects receive

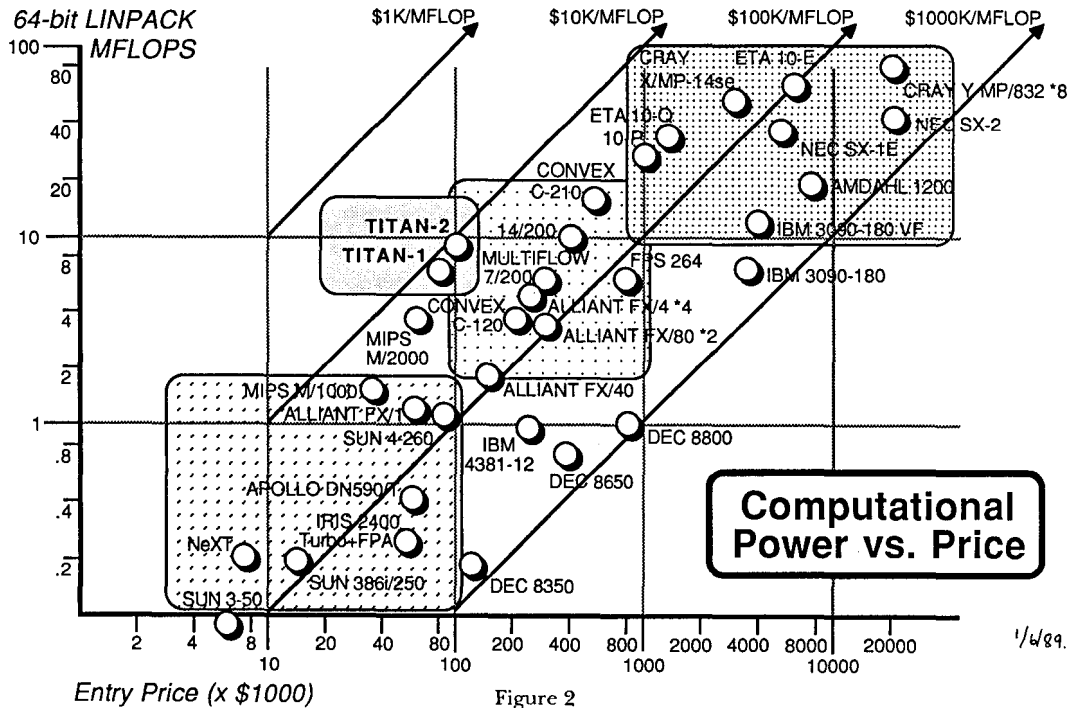


Figure 2

only a few processor hours per day, or about 1000 hours per year. Large users of supercomputer centers get about an hour per day. The average user consumes approximately an hour per week. Thus, provided the time need not be concentrated into a brief calendar period, a single graphics supercomputer can provide essentially the same amount of computation as a shared supercomputer, even for a large user.

By automatically parallelizing LINPACK, even for the small case, the Cray Y/MP runs about 2.5 times faster using all eight of its processors. This again establishes the Cray as the world's fastest computer. This also shows the importance of parallelization to increase speed. Since the 100 x 100 LINPACK benchmark is too small to run at top efficiency in parallel, the cost-effectiveness of the approach decreases by over a factor of three. Since the Titan has only two, relatively slower processors, the effect of parallelization on cost-effectiveness is not as great. However, stretch times in the order of 10-20 for the distributed, dedicated graphics supercomputers show that even large users can get about the same amount of work done as if they had used a centralized supercomputer.

By computing at the peak speeds, which only can be reached by running each of the processors at peak speed and in parallel, the potential differences in speed between the Cray YMP and the Titan finally are apparent. However, even though the times stretch by almost a factor of 90 (i.e. to do an hour of computing on the YMP requires almost 90 hours on the Titan), the cost-effectiveness of the Titan still remains greater, but only by a factor of two.

Table 1. Central Cray YMP versus distributed Titan Graphics Supercomputer

	Cray YMP 832	Titan 24
Price (in millions\$)	20.	.12
Processors	8	2
Megawords of memory	32	4
Dhrystones (integer-oriented)		
KDhrystones/sec single processor	25	23
Dhrystones/s/\$ multiprogrammed	.01	.383
Whetstones (scalar floating point)		
MWhetstones/sec single processor	65	6.5
Whetstones/sec/\$ multiprogrammed	26	108
Linpack (100x100) typifies actual use		
Mflops/sec single processor	79	6.5 (12)*
flops/s/\$ multiprogrammed	31.6	108
Mflops/sec parallel	195	9.4 (21)
flops/s/\$ parallel	9.8	78
Peak performance (1000 x 1000 Linpack, theoretical peak)		
Mflops/sec Linpeak	2144	25 (86)
Mflops/sec/\$ Linpack	107	208
Mflops/sec peak op rate	2667	32 (83)
Mflops/sec/\$ peak op rate	133	267
Millions of pixels rendered/sec	40**	50

*() the time stretch factor for Titan relative to the Cray YMP
 ** with a frame buffer directly connected to Cray channel

Finally, with the graphics supercomputer visualization is implicit in the system. Each computer has a significant amount of power which can be used to render and display data. The Cray Y/MP and Titan transform and render at roughly the same rates. Modern supercomputing requires additional systems, such as graphic supercomputers or high performance workstations, just to handle the preparation and drawing of displays based upon the computed data. We recommend that future supercomputers include embedded rendering hardware to provide cost-effective and truly interactive graphics. Networks alone will not be capable of providing sufficient levels of interconnection bandwidth to obviate such hardware.

Titan: The first graphics supercomputer

Titan, introduced March 1988, combines the essential architectural components of the Cray scientific computing mainline with a very high performance graphics rendering processor and display that is used to visualize the results of computations. Figure 3 shows the relative computation and graphics performances for Titan compared with other workstations and computing environments. Many interesting applications become possible, ranging from the modeling and animated display of molecular structures, as required in computational chemistry; computational fluid dynamic modeling and display which permits interactive design; geological and medical imaging and modeling; and finite element modeling for interactive design of physical objects, displaying force, heat, fluid flow fields, electric and magnetic fields, etc.

Titan is a multi-processor computer with up to four vector processors, each of which provides a peak rate of 16 million instructions per second and 16 million, 64-bit floating point, operations per second, operating from a common, shared memory of up to sixteen million 64-bit words (128 megabytes). The graphics hardware unit is capable of shading 50 million pixels per second. A single processor can transform and display at the rate of 25,000 Gouraud shaded triangles and 55,000 depth-cued, Gouraud shaded vectors per second. Titan operates in a distributed UNIX environment via Ethernet. This provides a common, compatible, transparent environment. Fortran and C compilers provide both vectorized and parallelized compilation. A more complete description is found in Bell, et al [3].

Figure 4 gives a system block diagram of the Titan hardware. The central communications mechanism is the unified, high-speed bus for interconnecting 10 modules. This bus can operate at high levels of utilization and interconnect the processors, graphics, I/O subsystems, and memory system at up to 32 megawords per second (or 256 megabytes per sec). Data is transferred at a synchronous rate of 16 megahertz. Nearly all other multiprocessor systems being introduced (e.g. Apollo, Digital, Silicon Graphics) use similar but substantially slower, central busses. Similarly, while the multiple processors of these systems provide more power than a single processor, they simply cannot achieve the very high peak processing rates that characterize a supercomputer. For example, using all four processors and the best formulation of linear algebra (i.e. LAPACK), Titan operates at 47 megaflops, or about 5-8 times the speed of ordinary, multiprocessor workstations.

Graphics versus Computational Performance




Figure 3

The memory subsystem provides up to 16 megawords (128 megabytes) on up to four boards, with four megawords per board. Each board is interleaved eight ways by using independent memory banks. Two identical boards operate together to provide 16-way interleaving. A variety of I/O devices can be accommodated either directly or through an adapter to an internal VME bus which holds two VME boards. Two SCSI controllers drive three internal disks, providing over two gigabytes within the base enclosure. Additional external mass storage and I/O can be attached. Using the fast file system, data transfer occurs at a rate of over one megabyte per second for each SCSI channel, and over two megabytes per second for each channel when VME attached SMD peripherals are used. Titan consists of a base system enclosure which is 50.5" x 22.5" x 22", weighs 170 Kg, and consumes a maximum of 1692 watts. The Desktop component of the system consists of a 19" CRT, keyboard, mouse, and user console which can be located up to 200 feet from the base enclosure.

Doré—Graphics for Interactive Visualization

Standards are the key to providing powerful graphics. While PHIGS is the current standard dealing with 3D graphics, it is acknowledged to be functionally inadequate and difficult to use. PHIGS+, which includes some of the needed extensions to PHIGS, has not yet become a standard. As a result of delays in standardization, many system suppliers are making unilateral extensions to PHIGS which will lead to confusion and incompatibilities between manufacturers. However, in our opinion, PHIGS+ also is insufficient. It is an instance of an

earlier genre in which it is difficult to achieve the best system performance. Users also find it relatively difficult to use, requiring an entire extended software environment to be applicable, and it lacks capabilities important for advanced scientific applications. Therefore, we have developed and offered Doré (Dynamic Object-oriented Rendering Environment) as a de-facto standard for the entire community requiring interactive, dynamic 3D graphics. Doré has a fundamentally different structure than PHIGS+. This is shown in Figures 5 and 6. Doré has been used for many applications for over a year, is object-oriented and is licensable to other system suppliers. Doré is user extensible without the need for source code. Its design objectives included:

High performance. Graphics hardware should be driven at the highest possible speed.

Advanced rendering. Not only must the library provide interactive dynamics, it must provide advanced, photo-realistic rendering, including shadows, transparency, reflections, and textures.

Orthogonality of functions and attributes, through description versus procedure. Rather than describing steps to display a picture, a graphics environment permits specification of orthogonal attribute types for the items being displayed. These attributes specify the desired rendering, grouping for control (e.g. picking) and perspective for viewing.

Extensible. Given the nature of constant change in capabilities of graphics, including the inclusion of image processing primitives, an environment must be highly modular and user-extensible, permitting graceful evolution including image processing. Although written in C, Doré is written

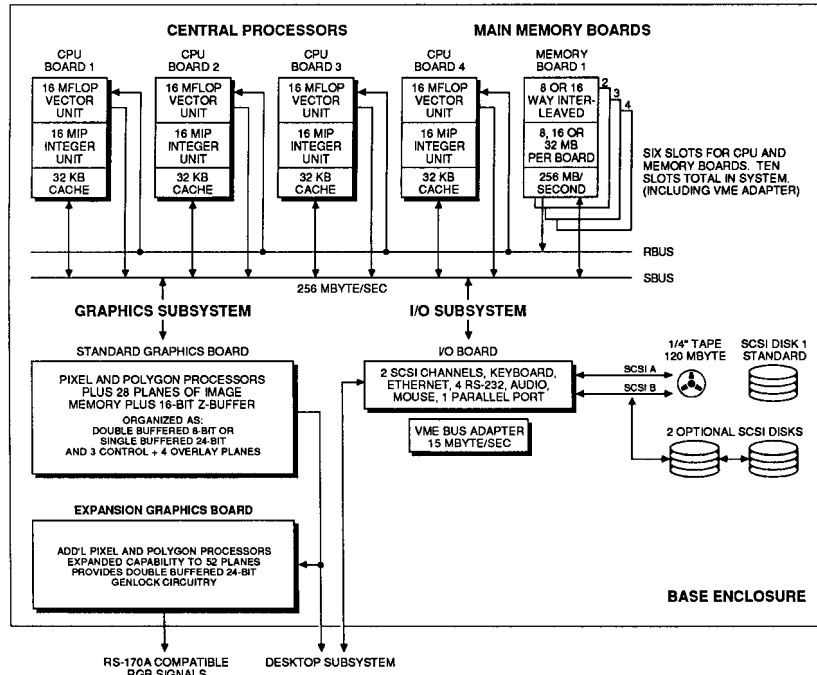


Figure 4

with a set of conventions that adhere to object-oriented programming principles.

Applications Oriented. Doré was designed from an applications viewpoint. No additional libraries are required before the connection is made to application. We have found that an application can be connected to Doré with minimal effort.

Portable. Doré should be easily portable to other graphics systems. These systems must be able to utilize their graphics hardware with the best possible performance.

Future Hardware Orientation. Given the rapid changes in hardware, a library design must be ready for the future, not a reflection of past, hardware limitations.

Future: Graphics Supercomputers, Personal Supercomputers and High Performance Workstations

In order fully to exploit the evolving new class, in preparation for true, distributed personal supercomputers or supercomputing workstations, several areas must evolve more rapidly.

Computation

The basic computer, including its mass storage, is improving at the rate of over 70% per year, providing a doubling of power in less than 18 months. For example, our second generation product is approximately two-and-a-half to three times more powerful than Titan, at nearly the same price. It can be installed as a field upgrade to Titan and provides almost an order of magnitude better performance/price than the Cray YMP. We anticipate that this high rate of improvement will extend into the 90's, whereas we see the supers and minisupers constrained to a 14% per year improvement (at a maximum) in clock speed. While obviously the inherent circuit speed of a less expensive machine can not reach that of a Cray, it can approach it as more of the

processor is put onto a single chip. Furthermore, these chips can continue to have relatively low costs. Several suppliers are in the process of integrating a RISC processor onto ECL and GaAs (the componentry of the Cray 3) chips.

Graphics

The most important graphics and image processing functions will continue to be assimilated into hardware, including renderings approaching the quality of raytracing (e.g. Phong shading), complete anti-aliasing, a variety of primitive solids (e.g. spheres), transparency, texture mapping, constructive solid geometry primitives, automatic depth cueing, etc. In the near future, the fastest graphics workstations are expected to approach one million vectors and one-half million polygons per second.

Titan presently is predicated upon true color, using 24 bits/color. In the future nearly all 3D graphics systems will have this capability. The availability of high definition TV also will exercise a profound effect, by defining higher quality standard video information, and by stimulating development of lower cost components to store, process, and transmit such information.

Networking

Both local and wide area networks lag computation and graphics and significantly hamper distributed computing. While poor networking contributed to the need for graphics supercomputing, it is not good for supercomputing generally. A standard LAN in the form of FDDI operating at 100 megabits per second is essential now, followed by a next generation LAN which fully can exploit gigabit fiber optics and permit communication at memory speeds and video rates. A wide area network providing minimum point-to-

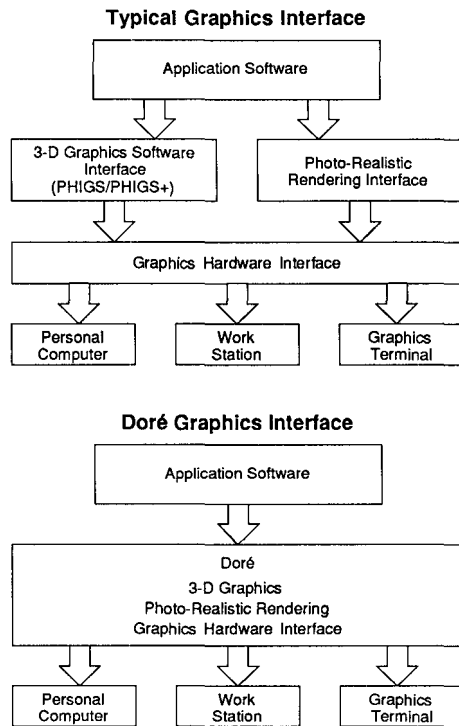


Figure 6

point rates of 32 megabits (ten frames per second for a 1280x1024 true color display) is required for interactive graphics and image processing. Such a network actually will enhance the value and power of distributed graphics supercomputers.

Application Programs which can exploit interactive supercomputing

Many of the application areas where graphics supercomputing should make the greatest impact are those such as mechanical computer aided design which have millions of lines of two-decade-old programs. These programs have several problems. They are neither interactive nor designed for an interactive environment with visualization. Further, they are not oriented to supercomputers which provide substantial parallelism. While current vectorizing and parallelizing compilers can accelerate the execution of such programs to some degree, contemporary programs, designed from the outset with algorithms and data access logistics strategies tuned for the graphics supercomputer, can lead to better performance in many cases.

User Training

Until computer science and engineering curricula reflect the reality of both parallel, vector processor computers and highly interactive graphics, the potential of current graphics supercomputers and the evolving personal supercomputer will be both later in arriving and poorer in quality than we might hope. In addition, new ways of looking at application problems, including new algorithms, need to be encouraged. Such research, as well as the study of fundamental application algorithms, is essential for computer scientists and engineers to gain the insights upon which to base




Figure 7

architectural advances. VLSI design was solved by computer scientists working with electrical engineers. A similar transformation of mechanical design, or of any of the other main application areas, is likely to require a similar collaboration between computer scientists and specialists skilled in the particular application area.

Summary

We have tried to demonstrate that a new class of computer, which can be viewed as a member both of the workstation and of the supercomputer families, has been introduced. Systems within this class will evolve rapidly and will have as profound an effect in creating new applications as have each of the past classes of computers.

Acknowledgements

Titan including the compiler and Doré were designed and implemented by a team of 50 exceptionally competent and dedicated engineers over a period of 2-1/2 years. Paul Ausick performed the layout for the paper.

References

- [1] Borden, B., Visualization Toolkit, *Computer Aided Engineering Review*, November, 1988.
- [2] Savage, M., Andrews, J., and Benzel, M., The Advantages of Vector and Parallel Architecture, *Scientific Computing and Automation*, February 6, 1989.
- [3] Bell, C. G., Miranker, G.S., and Rubinstein, J., Supercomputing for One, *IEEE Spectrum*, Vol. 25, No. 4, April 1988.