LEGIBILITY NOTICE

A major purpose of the Technical Information Center is to provide the broadest dissemination possible of information contained in DOE's Research and Development Reports to business, industry, the academic community, and federal, state and local governments.

Although a small portion of this report is not reproducible, it is being made available to expedite the availability of information on the research discussed herein.

LA-UR-88-2454

LA-UR--88-2454

DE88 014397

TITLE: A SCIENTIFIC VISUALIZATION WORKBENCH

AUTHOR(S): Richard L. Phillips, C-6

SUBMITTED TO: Supercomputing '88, November 14 - 18, 1988, Orlando Florida

/

A Scientific Visualization Workbench

Richard L. Phillips[1]
Los Alamos National Laboratory
P. O. Box 1663, MS B-272
Los Alamos, NM 87545
505-665-1343
rlp@lanl.gov

ABSTRACT

A system for visualization of data from supercomputer simulations has been developed for use by scientists and engineers at Los Alamos National Laboratory. The *scientific visualization workbench*, as the system is called, is based on an industry standard workstation (a Sun 3/160C), the NeWS window system, and a video/graphics add-in card, which is supported by NeWS. Also involved is a frame buffer attached to a 48-Mbit/s Cray channel and a video link from the frame buffer to the Sun workstation. Computation and graphics preparation are performed on the Cray under the control of a NeWS client program. Images are rasterized and sent to the workstation at the rate of 25 frames/s. The workbench permits the scientist to converse with applications running on a Cray but still retain the convenience and flexibility of a personal workstation environment. Moreover, the video link allows animated graphics images to be displayed on the workstation at a much higher rate than would be possible with a conventional network connection, i.e., Ethernet. Also, NeWS affords the convenience of a modern window system. Live video can be displayed in a window anywhere on the screen, while other windows may contain conventional text or graphics.

Keywords: workstation, supercomputer, video, window system, visualization

INTRODUCTION

There has long been an interest in providing researchers with a method for viewing data derived from large-scale scientific computations. The sheer volume of data emitted by today's supercomputers demands a way of representing it compactly; graphics is the solution. Several approaches to visualization exist; they can all be classified as either *interactive* or *non-interactive*.

The *non-interactive* approach involves running a scientific simulation as a batch job and processing the data at a later time. One method of viewing, which has been used by large scientific installations for years, is to produce a film or microfiche on a film recorder. This usually takes the form of a two-dimensional graphical representation of isolines or finite difference meshes. Often an animated film sequence is produced showing the temporal evolution of a simulation. Another approach to offline visualization is the TMDS [1] system, developed at Lawrence Livermore National Laboratory. In TMDS several hundred centralized frame buffers service over 1800 remote video displays located throughout the Laboratory. Graphical information is rasterized on the supercomputers and distributed to a user's display. Under some conditions animation rates as high as 15 frames/s can be achieved, but mostly a series of static images is displayed. TMDS is also used for conversational monitoring and suspension of the progress of a simulation, but this is by no means considered interactive.

More recent approaches to non-interactive visualization involve using modern computer graphics techniques to produce synthesized, realistic images from the simulation data. Here models of lighting, reflectance, and transparency can be applied to render data to an easily perceivable form. Methods such as ray tracing and radiosity, which in themselves consume large amounts of supercomputer time, can produce images of striking reality. Sometimes these tools are used directly by the scientists who produced the original data, [2] but the "service bureau" approach is also widely used. Here experts in graphic design and animation consult with scientists and produce slides, videotape, or film that depict the desired phenomena [3].

The *interactive* approach to visualization actually represents a spectrum of interactivity. In general it means having some degree of capability to view and control a simulation while it is being performed. One extreme would be the ability to periodically interrupt the computation and view intermediate results. The other extreme would be a dedicated supercomputer tightly coupled to a dedicated high-performance display system.[2] Even then, the degree of interactivity would depend on the complexity of the simulation being performed. Moreover, this is not a cost-effective approach to visualization. Recently, a new class of computer has been announced by several companies.[3] Called graphics supercomputers, they have mini-supercomputer (Alliant and Convex class) computing power with tightly coupled, high-performance graphics displays. These machines, however, are in the $100,000 price range and may not be cost-justifiable in situations where many scientists require visualization capability.

A common situation in many large scientific organizations is where hundreds of workstations are connected over a local area network to one another and to one or more supercomputers. This situation not only provides the user with the wealth of software and graphics utilities that today's workstations offer, it also gives him a degree of control over simulations running on a supercomputer from a system conveniently located in his office.

1 This work was performed under the auspices of the U.S. Department of Energy.

2 Apple Computer often uses their CRAY X-MP/48 with attached Ultra frame buffer as a single user system.

3 Ardent Computer, Stellar Computer, and Apollo Computer all offer machines in this class.

Computing tasks can be distributed to machines where the job can be done best and, with modern window systems like X-Window and NeWS, a program running on a supercomputer can control the display in a window on the workstation [4]. Best of all, since workstations are relatively inexpensive, organizations can afford to provide them in quantity for all who need them. The workstation approach to visualization is explored in the remainder of the paper.

WORKSTATION-BASED SYSTEMS

Object Data Approach

The most common local area network used for workstations is Ethernet, a baseband 10-Mbit/s contention-based network. The most often-used protocol is TCP/IP,[4] which because of its overhead, seldom yields an actual throughput greater than a few hundred kilobits per second. This limitation strongly affects how one uses a workstation for visualization of supercomputer data. For example, one can transmit graphical data in object form[5] from supercomputer to workstation and use the local graphics hardware to display it. For complex images, however, or where animation is desired, the limitation of the Ethernet, coupled with the relatively low performance of standard workstation graphics hardware, yields a drawing speed that is just too slow. On standard workstations, even with careful tuning of hardware and software, one can achieve animation rates of only a few frames per second for images containing a few hundred vectors. For proper animation, one would like about 20 frames/s, preferably 30 frames/s.

One can partly overcome the speed limitation by turning to graphics accelerators, specialized boards that significantly speed up graphical operations. Workstation manufacturers provide these as do third-party hardware suppliers. Recently, Raster Technologies and Trancept Systems (now a subsidiary of Sun Microsystems) introduced accelerators that promise to have a strong impact on visualization efforts. Raster Technologies offers the GX4000, a board set that plugs in to a standard Sun Microsystems workstation. This board implements the proposed PHIGS+ standard in microcode and thus provides a three-dimensional graphics capability for the host workstation. This means that full three-dimensional data can be sent from a supercomputer and transformed locally. Thus, an object can be animated just by transforming it and re-displaying it in a new orientation. Depending on the complexity of the object, good to excellent animation rates can be achieved.

The Trancept board is called the TAAC-1, which is an acronym for Trancept Application Accelerator. As the name suggests, this board can be used to speed up all kinds of local computing tasks. When used for graphics applications, it can perform sophisticated rendering operations at very high speed. While the board is an add-on to a Sun Microsystems workstation, it has its own frame buffer and is capable of producing high-resolution 32-bit/pixel images. The board also has the the potential of accepting data directly (bypassing the Sun processor) at 800 Mbyte/s, but further development work is required before this feature can be used.

While the graphics performance of both of these accelerators outstrips conventional workstation hardware, the rate at which data can reach them is, in general, still limited by Ethernet throughput. This may not be a problem for investigations where a relatively small amount of object data is sent from the supercomputer and is expanded locally to produce high-quality images, animated sequences, or both.

There is a large class of problems, however, where this is not an especially useful approach. These are typically simulations involving finite difference calculations where complex meshes are used to describe the spatial aspects of the problem. An example of such a mesh is shown in Figure 1. This represents one instant in time in a sequence of many thousands of such images. As the problem evolves, some part of the mesh changes from one time step to another. Each image comprises several thousand data points and, because of Ethernet limitations, one cannot expect to achieve very good graphics display rates. The accelerator boards are no help here because there is no local algorithmic processing they can usefully perform. An exception would be the case where only those parts of the image that change in a time interval are sent to the workstation; the accelerator would maintain the graphics database and apply the changes locally. This is not a simple task, however, and has not been done in practice.

Image Data Approach

Another approach to remote data visualization is to perform some, if not all, of the graphics processing on the supercomputer and send the complete representation to the workstation as an image, i.e., pixels. But, if the image resolution is say, 1024 x 1024 x 8, a megabyte of data must be sent for each data display. Thus, there is no speed advantage to this approach unless a transmission medium other than Ethernet is used. Providing a second data transmission path, one that has a much higher bandwidth than Ethernet, is the thrust of this paper. Of course, there must be a device in the workstation that can handle that data. This device is another type of accelerator, a board that accepts an image directly and displays it on the

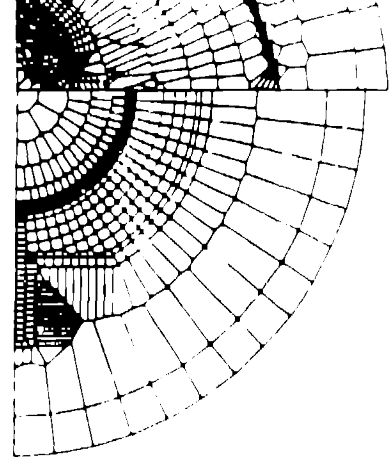

Figure 1.

[4] Transmission Control Program/Internet Protocol.

[5] The term, object data, refers to the numerical representation of the graphical data as calculated by the simulation.

workstation screen. Parallax Graphics, Inc., makes such a board. It can digitize an NTSC[4] video signal in real time, i.e., 30 frames/s. The *scientific visualization workbench*, the subject of this paper, makes use of the Parallax board capabilities, in conjunction with an already-developed supercomputer frame buffer system and the NeWS window system. This image-based approach to visualization overcomes the speed limitations of the object-based approach, but at a price. Image quality is lost by NTSC encoding of a video signal. Thus, one is trading image fidelity for bandwidth. On the other hand, the user sees images as fast as they can be produced by the supercomputer, and temporal fidelity is often more important than image quality. There is another advantage to be gained from processing graphics data on a supercomputer--the flexibility to experiment with new visualization algorithms, rather than rely solely on rendering methods provided by the workstation vendor.

NETWORK WINDOW SYSTEMS AND NeWS

A key element of the *scientific visualization workbench* is a workstation window system that permits distribution of computing tasks in a network environment, in this case Sun Microsystems' NeWS. The acronym NeWS, derived from Network-extensible Window System, suggests the network orientation of this window system. As such, it allows the computing resources in a heterogeneous computing environment to be effectively used. NeWS runs on a machine with one or more bit-mapped displays. It acts as a window server, managing input and output on its host machine. Application programs--called clients--send messages causing NeWS to render images on the display. The clients may reside anywhere on the network. Server-based window systems are often called distributed window systems, because the server and its clients may be distributed over the network. This concept is not unique to NeWS. It appears in Andrew [5], a system developed at Carnegie-Mellon University, and the X-Window System [6] from MIT. Figure 2 depicts this network orientation, where the NeWS server, running on a workstation, serves remote clients running on other machines.

The supercomputer depicted in Figure 2 is the source of data for the *workbench*. As the host for the client program, it provides graphical data, either object or image, to the server running on the user's workstation.

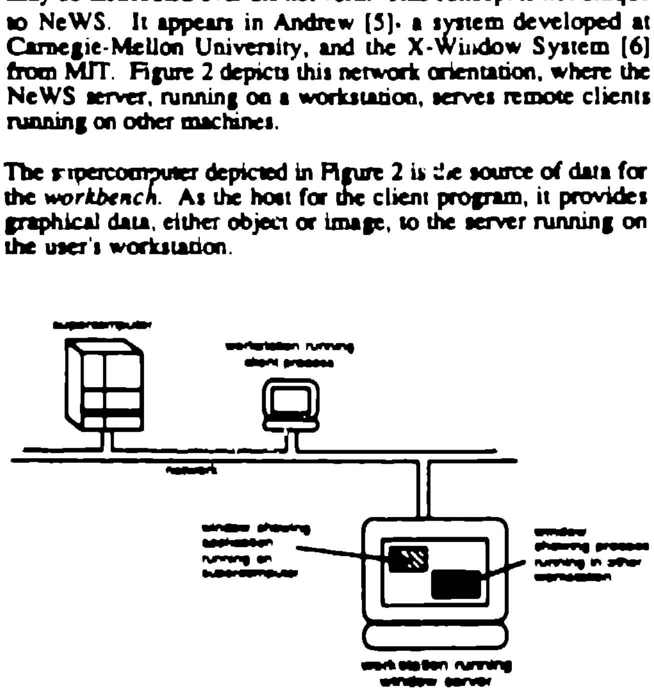

Figure 2.

[4] NTSC stands for National Television Standard Committee. Often referred to as composite video.

One of the key features of NeWS is the use of a programming language for communication between the server and its clients. Instead of sending messages consisting of commands and parameters in a fixed format, clients send programs that the server interprets. The language used by NeWS is PostScript [7], which was originally developed as a language to drive printers. Sun has extended PostScript to include features necessary for operation in an interactive workstation environment. These features are described in detail elsewhere [8]. It is important to note here that the presence of PostScript provides an extremely powerful window system extension facility, making it possible to rapidly develop new tools.

THE SCIENTIFIC VISUALIZATION WORKBENCH

System Overview

The structure of the scientific visualization workbench is depicted in Figure 3.

Each component of the workbench will be described in detail, but briefly, it comprises a Sun Microsystems 3/160C workstation equipped with a Parallax 1280 Series Videographics Processor, a high-resolution color monitor, a Cray Research, Inc., supercomputer equipped with a Solid-state Storage Device (SSD), and a 512h x 512v x 8 bits/pixel frame buffer attached to the Cray by a 48-Mbit/s channel. Communication between the Sun and the Cray is via Ethernet using TCP/IP protocol. The Cray runs the UNICOS 3.0 (System V-based UNIX), and a special Parallax version of the NeWS 1.1 window system runs on the Sun. A NeWS client program on the Cray controls the transmission of rasterized data to the frame buffer and ultimately to the workstation over an NTSC video link. The magnified workstation screen in the lower left of the diagram suggests the presence of several types of windows, in which independent processes are running simultaneously.

The Supercomputer Subsystem

The supercomputer side of the workbench is based upon work previously reported by Fowler and McGowen [9]. They describe the development of a 512h x 512v x 8 bit/pixel frame buffer installed on a 48-Mbit/s I/O channel of a CRAY X-MP/416. Output from the frame buffer is separate red, green, and blue (RGB) non-interlaced video signals with a refresh rate of 60 Hz. The output signal is viewed on a color monitor attached directly to the frame buffer or is transmitted over a fiber optic link to a remote location. Output to the frame buffer is controlled by a highly tuned program that takes advantage of the vectorization capabilities of the Cray. Under certain loading conditions (the Cray is not dedicated to this task) users can achieve animation rates of almost 25 frames/s. The Cray operating system is CTSS, the Cray Timesharing System, and users access the frame buffer control program by logging in on high-speed serial lines. The control process and the display process are disjoint, i.e., they are associated with two different devices.

The workbench uses the supercomputer facilities somewhat differently. The operating system is UNICOS (a UNIX-like operating system developed by Cray Research), and the control and display functions are both incorporated into the workstation. From the workstation end, the user controls Cray activity over an Ethernet link through a NeWS process. This process communicates with a Cray-based NeWS client program, which in turn directs the frame buffer display software described above. There is a separate video link that transmits the frame buffer video signal to the Parallax processor. This is currently an NTSC video signal, derived by encoding the RGB output of

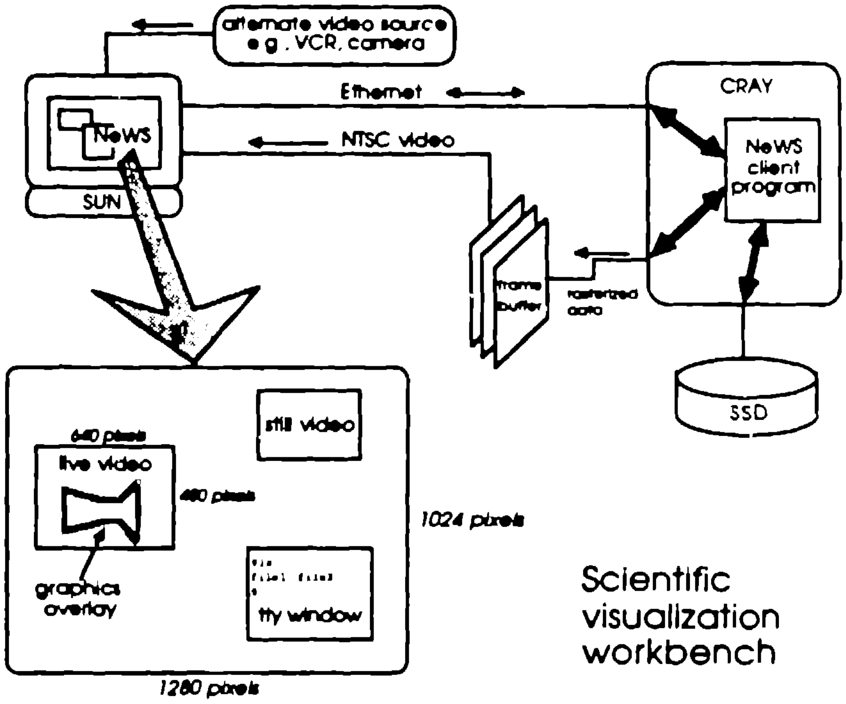

Figure 3.

the frame buffer. The signal is channelized onto a (ubiquitous) cable video network and decoded with a tuner at the user's location. The trade-off of image quality for bandwidth extracted by NTSC encoding was mentioned earlier; a later generation of the system will hopefully involve direct transmission of RGB video to the workstation.

The Workstation Subsystem

The workstation is a Sun 3/160C, but any Sun model with a VME bus interface can be used. This bus accommodates the Parallax 1280-V-8VN videographics card set. Other workstations, supporting different bus protocols, could also be used. For example, Parallax has a Q-Bus version of the videographics card set that can be accommodated by the Digital Equipment Corporation's line of MicroVAX workstations. Naturally, a different set of software would be required. The video display can be either the standard Sun color monitor (1152h x 900v resolution) or any number of third-party monitors. The Parallax board normally produces a 1280h x 1024v signal but it can be programmed to provide a wide range of video formats. The monitor used for the development project is a 19-inch Sony GDM 1950 with 1280 x 1024 resolution.

The Parallax Videographics processor is a complete graphics controller and frame buffer, with a size of 2048 x 2048 x 8. Of this, only 1280 x 1024 is viewable. The graphics hardware of the parent workstation is not used at all; indeed, the Parallax standard graphics performance generally exceeds that provided by the workstation. What makes the board set unique, however, is its ability to digitize and display an NTSC video image anywhere in its graphical address space. Video digitization is performed at 30 frames/s and produces an image size of 640h x 480v, consistent with NTSC bandwidth. Because of a proprietary dithering technique, which uses a 14-bit YUV (luminance and chrominance) color space, the resulting image quality is extraordinary. The board outputs RGB video directly into the color monitor. More details on the Parallax product can be found in the user manual [10].

PNeWS, NeWS Enhanced for Video

Parallax Graphics has ported the NeWS 1.1 system to its videographics card and extended it to support video features. The resulting system, called PNeWS [11], is the foundation for the scientific visualization workbench. PNeWS preserves the purity of the original NeWS product. Only five new operators have been added: four to support video and one to provide general access to all features of the Parallax hardware. One extension, for example, defines a new type of window, a video window. There is a corresponding operator that specifies the source of data to be displayed in the window, i.e., a live video signal. Video windows can contain still as well as live video. Moreover, conventional graphics can be overlaid on any video window. One can selectively digitize live video frames, each in a separate window, and resize them at will. The workbench builds on these capabilities by allowing the user to produce hard copy, film, or printed output of any such video image. PNeWS also supports the Parallax board capability of switching between two separate video signals, a feature that is useful for recording and accessing images on a laser disk recorder. Details of the tools available to the workbench user follow.

System Operation and Capabilities

Launching the workbench is simple. Once in the NeWS environment, the user opens an authorized session on the Cray through a terminal emulator window on the workstation. By pressing a (software) button on the workstation screen, the user begins the client process, which in turn establishes the client/server path on the network. At this point NeWS is in control of the entire distributed environment. The user issues commands via software sliders, buttons, and text fill-in dialogs. Figure 4 shows a typical screen layout for the workbench. In the upper right, for example, there is a partly visible control panel for setting video parameters such as contrast and brightness. The image in the lower right shows a video image annotator and its associated controls. Other controls are used for

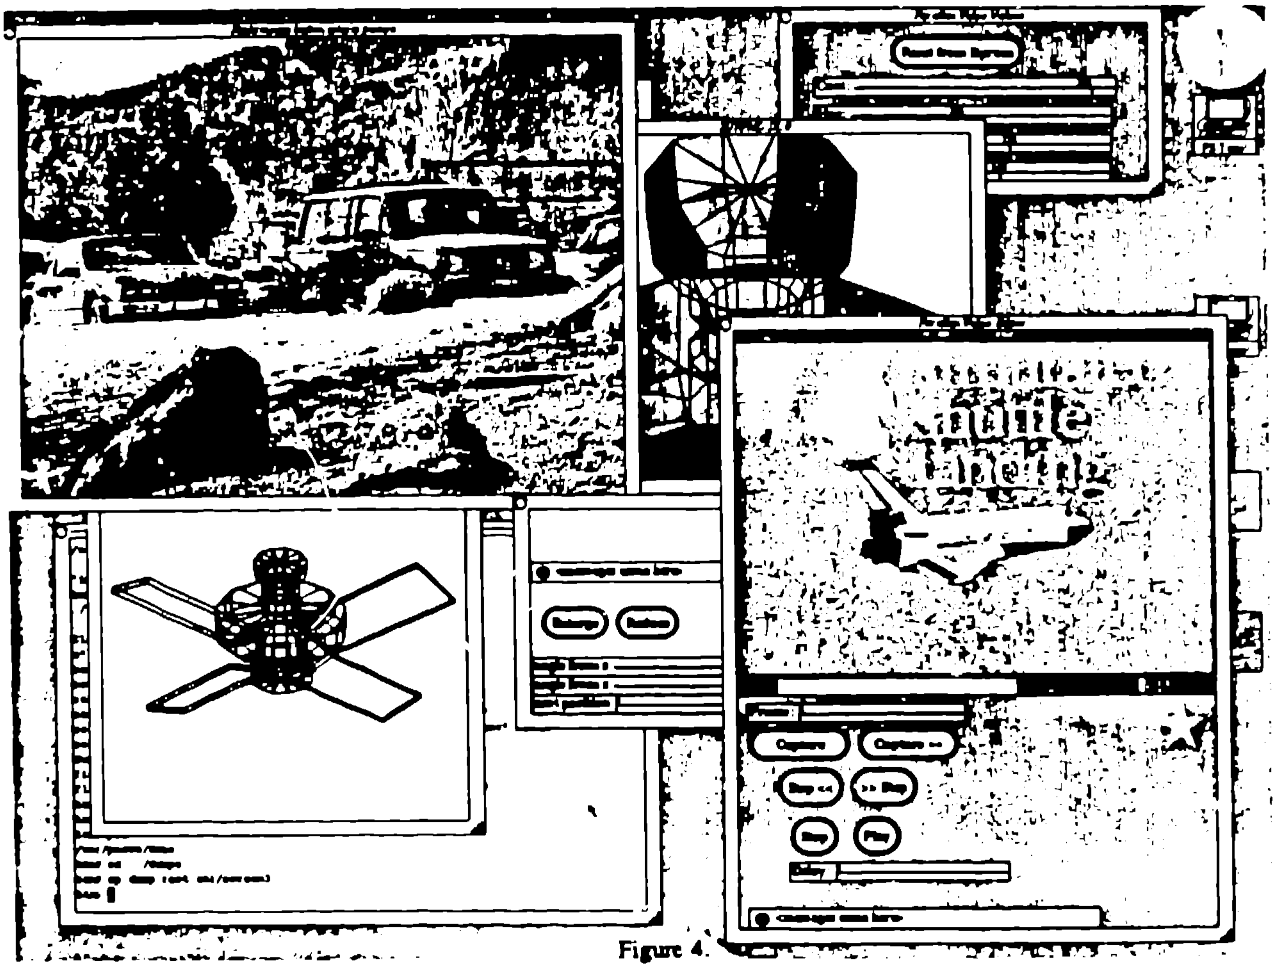

Figure 4.

such operations as video image sequencing and speed, video frame capture, graphics overlay, etc.

The user has control over many aspects of the workbench user interface. Naturally, all windows, video or conventional, can be sized and located as the user wishes. In addition, the default color look-up table can be modified to suit the image being viewed. For example, there is often more insight to be gained by viewing an image in shades of gray rather than color; the workbench provides that flexibility. Also, an alternate source of video images can be selected; Figure 5 shows a situation where the image in the lower right was produced by a video camera viewing an outdoor scene while that in the upper left is a frame from a supercomputer simulation of supersonic jet flow [12]. Yet another tool is the capability of capturing many video images in sequence and playing them back in a rapid loop, or one frame at a time. While there can be only one live video window active at a time, many loop sequences can be captured and, if one wishes, played back simultaneously. Finally, one can apply a 2X zoom to a live video window, causing it to fill the entire screen. Hardware pixel magnification is used to accomplish this but the quality of the image suffers little.

Visualization Tools for Scientific Analysis

In addition to the user interface controls described above, there is a set of higher level tools available to the workbench user. One of these tools has already been mentioned and appears in Figure 4. This is a video image annotator, which also has a database function. Images can be captured from any source and can be annotated with a variety of textual and graphical features. Thus, the user can produce sets of images, all cataloged with appropriate annotations, and recall them for later viewing and hard copy production. The ability to easily produce hard copy should be emphasized. It is due in large part to the fact that PostScript is the language underlying NeWS, the window system used for the workbench.

Often the workbench user wishes to view images in a fixed spatial relationship, i.e., in the form of a mosaic or montage. The display could show several images that bear a temporal relationship to one another, or they could be drawn from different parameter spaces. In fact, one might wish to combine video images from different sources into a single mosaic, e.g., from a videodisk and an online database. The workbench provides this tool, and an example of its use is shown in Figure 6, which shows the temporal evolution of a supersonic jet. Time begins in the lower left corner and progresses from bottom to top, left to right. The user can select any order of presentation and any size and number of images. Another capability, not shown here, is automatic textual annotation of each image in the sequence with an appropriate parameter value.

Another tool allows the user to superimpose arbitrary graphical data on a live or still video image. Plotting isolines, e.g., temperature or pressure contours, is a common use for this tool. The isoline data is sent from the supercomputer (over Ethernet) and is transformed and drawn in the specified window. If the

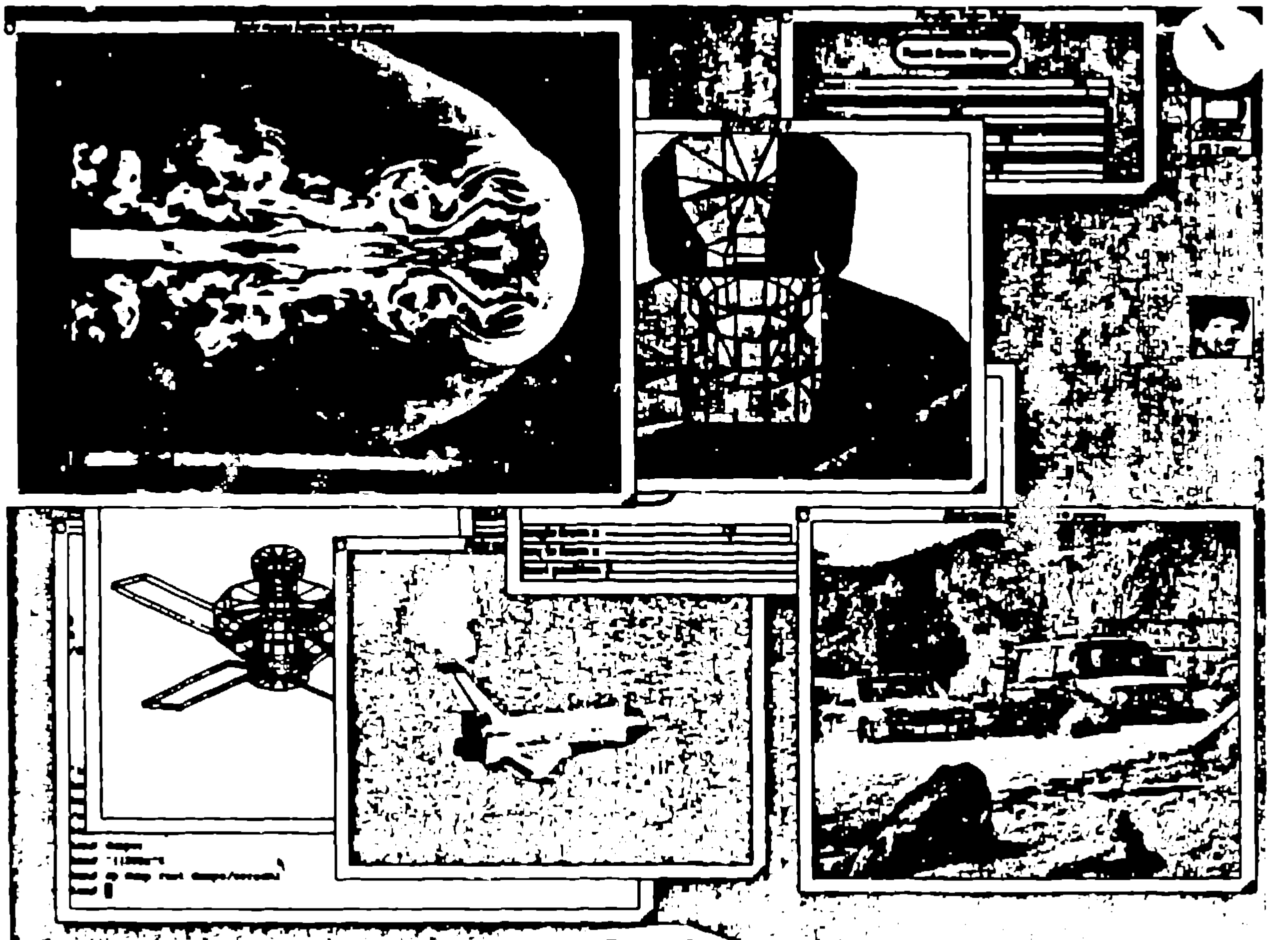

Figure 5.

data in the overlay do not change, the display can be converted to an image for subsequent immediate overlay. This is useful if the window is often closed and re-opened.

USING THE WORKBENCH

Because the workbench is still a development system, it is not widely accessible by potential users. Still, there is some indication how a typical user will interact with the system. When a scientist runs a large simulation on a supercomputer, he generally saves the graphical output in the form of a CGS metafile [13]. The NeWS client program, which is running on the supercomputer, processes those metafiles to produce rasterized output suitable for display on the channel-attached frame buffer. From a control panel displayed on the workstation screen, the user manipulates the display of these rasterized images as live or still video windows. The scientist can control the speed of animation and can browse forward and backward through the graphical data. Interesting, short sequences of data can be captured locally on the workstation and displayed either animated or a frame at a time. Typically, a user looks for features in the data that suggest a further investigation or highlight an error. Selective color table modification is often used in this process.

The video annotation tool described earlier is used as sort of a scientist's logbook. Up to 25 images can be saved in a named folio (file), along with appropriate textual and graphical markup. Any number of folios can be created and later retrieved by a rudimentary keyword search facility. Selected images from the folios can be printed or saved for incorporation into any documentation system that accepts a POSTSCRIPT representation.

SYSTEM PORTABILITY

While the workbench was developed for use in the Los Alamos computing environment, every aspect of the system technology can be transferred to another setting. While a direct connection to a Cray computer is considered an important attribute of the workbench, there are elements of it that will be useful even when this is not available. For those many Cray sites throughout the world, however, all elements should be applicable. All Cray computers are equipped with 48-Mbit/s channels. The frame buffer, which was developed at Los Alamos, is now available as a product.[7] It can be used on either of the Cray DA/DB or DI/DO channels. While the frame buffer produces RGB video signals, there are many NTSC video encoders on the market. The resulting signal can be sent to the workstation via either cable video or a fiber optic link.

As noted earlier, the workstation used in the system need not be a Sun. The Parallax board is available for several bus architectures other than VME, in particular for the Q-Bus, AT Bus, and the bus for the IBM RT/PC. Apollo Computer, for example, uses the AT Bus in their DN-3000 and DN-4000

7 Available from Sivco Engineering, Albuquerque, NM.

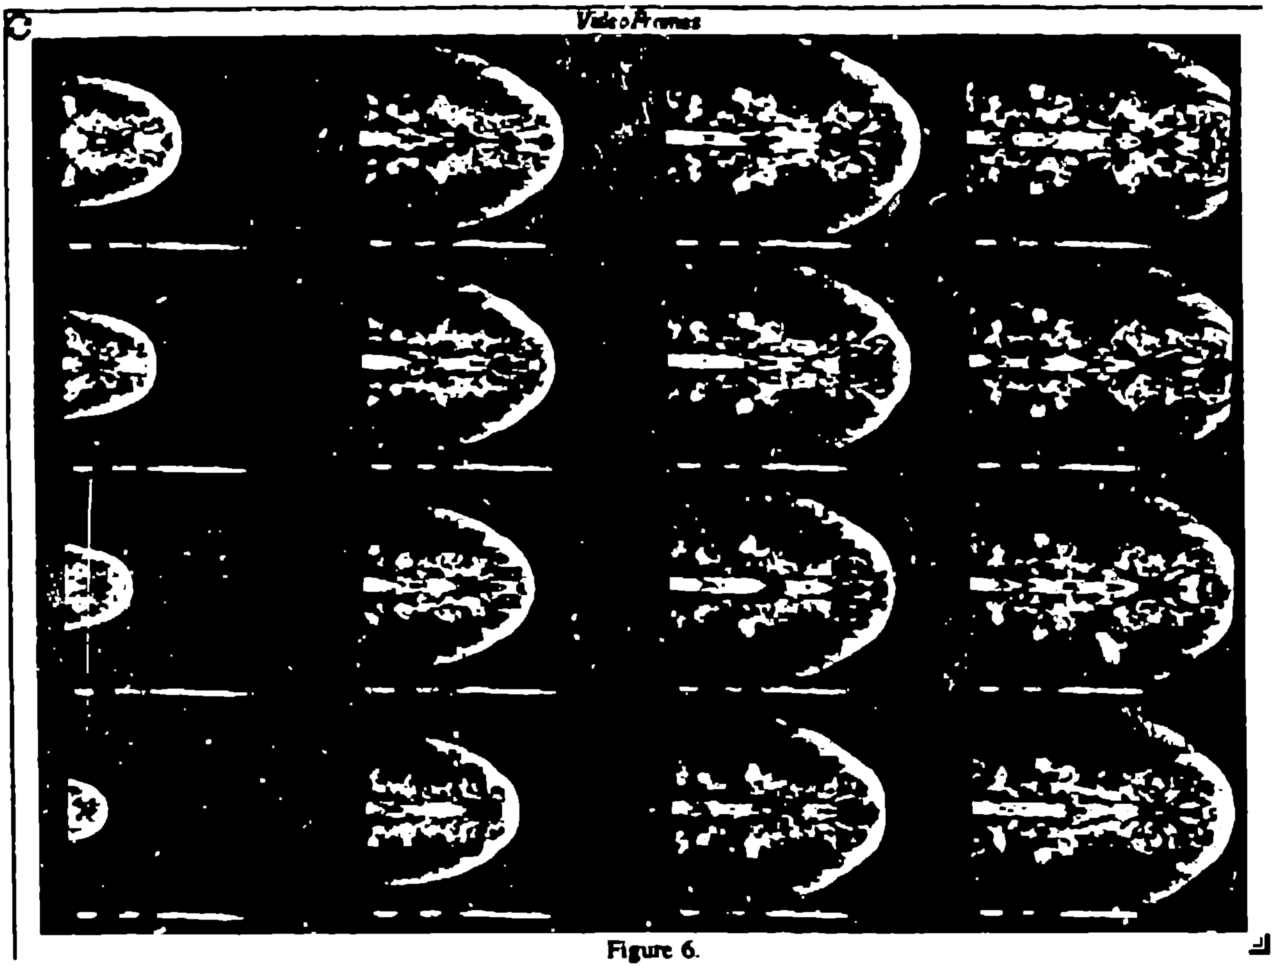

Figure 6.

series workstations and Digital Equipment Corporation uses the Q-Bus in their MicroVAX workstations. UNIX is available for both companies' workstations, and ports of NeWS exist for them as well. Of course, it would be possible to develop the equivalent capabilities of the scientific visualization workbench using the X-Window System; Parallax has software to support it as well as NeWS. But the higher level tools described here are NeWS-specific and would have to be rewritten for X-Window.

FUTURE DIRECTIONS AND ENHANCEMENTS

While the *scientific visualization workbench* is very useful in its current state, there are several features that remain to be added. Among these is a rudimentary page layout and report building facility. This will permit the user to assemble pages of figures and images into a form suitable for direct hard copy production on, for example, the new color POSTSCRIPT printer from QMS.

There are several host-based facilities that must be provided. Most important of these is a programmatic interface to simulation programs that run on the supercomputer. Currently, the workbench client program only processes metafiles that were previously produced by a simulation. Coupling the workbench to the simulation will provide more user control and interactivity. Finally, an image database facility should be provided, complete with run ID and parameter specification. This will allow the workbench to quickly and selectively retrieve images from an archive for display as single images or as part of a mosaic. Other enhancements, not envisioned now, will no doubt be requested by users of the workbench.

REFERENCES

[1] J. Fletcher, *The TMDS Manual*, Lawrence Livermore National Laboratory, UCID 30118, Feb. 9, 1976.

[2] M. L. Prueitt, "A Window on Science," IEEE Computer Graphics & Applications, Sept. 1987

[3] C. Upson, "Issues in Visualization," JPL Workshop on Scientific Visualization, Pasadena, CA, Jan. 1988

[4] R. L. Phillips and D. W. Forslund, "Using the NeWS Window System in a Cray Environment," Cray Users' Group Meeting, New York, NY, May 1987.

[5] J. Morris et al., "Andrew: A Distributed Personal Computing Environment," Communications of the ACM, Vol. 29, No. 3 (March 1986).

[6] R. W. Scheifler and J. Gettys, "The X Window System," ACM Transactions on Graphics, Vol. 5, No. 2, April 1986, pp. 79-109.

[7] Adobe Systems Inc., *PostScript Language Reference Manual*, Addison-Wesley, 1985.

[8] Sun Microsystems, Inc., *NeWS Technical Overview*, March 1987, Part No. 800-1498-05.

[9] John D. Fowler, Jr., and Michael McGowen, "Design and Implementation of a Supercomputer Frame Buffer System," LA-UR-88-890, Los Alamos National Laboratory, March 19^8.

[10] Parallax Graphics, Inc., *User's Manual for 1280 Series Videographics Processor*, January 27, 1987.

[11] M. Picco, "Live Digital Video in a Windowing Environment," National Computer Graphics Association, Anaheim, CA, March 1988.

[12] K. A. Winkler, J. W. Chalmers, S. W. Hodson, P. R. Woodward, N. J. Zabusky, "A Numerical Laboratory," Physics Today, October 1987.

[13] T. N. Reed, "A Metafile for Efficient Sequential and Random Display of Graphics," Computer Graphics, Vol. 16, No. 3, July 1982.