A CRAY RESEARCH, INC. PUBLICATION

CRAY CHANNELS

Volume 5, Number 1

Special computer graphics issue

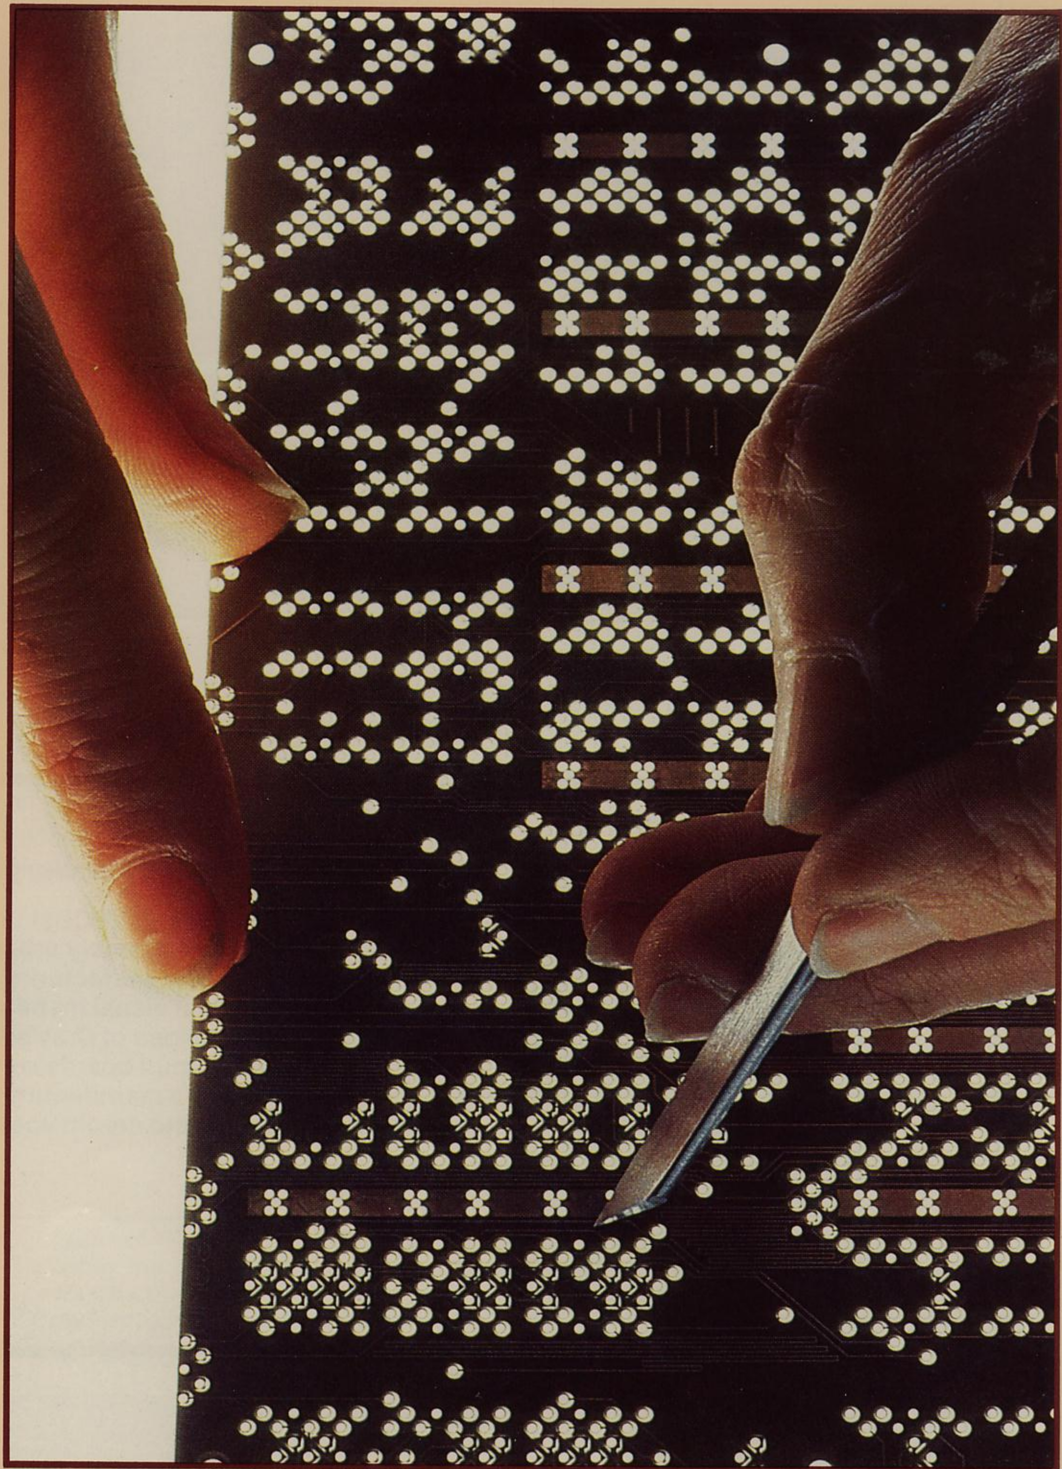

FEATURE ARTICLES:

REGULAR COLUMNS:

In this issue

A picture can indeed show at a glance what may take thousands of words to explain. For this reason, computer graphics has become an important tool for helping researchers in many disciplines keep up with today's onslaught of scientific information. Using computer graphics techniques, researchers can turn a wealth of data and information into concise graphic images for further analysis. Graphics specialists have also applied their techniques successfully to the entertainment medium, as witnessed by the emergence of computer-animated films such as TRON.

Increases in computing power over the past decade have helped make computer graphics a reality, and the advent of the CRAY-1 has played an important role in this evolution in two ways. First, graphic images may be so complex that a CRAY system may be required to generate them. Second, CRAY systems execute application codes so complex that accurate interpretation often requires graphics output as opposed to numeric. In these cases, graphics may be generated directly from the CRAY or downloaded to a graphics system. Recent progress in these two areas has been so exciting that we thought it appropriate to devote an entire issue of **CRAY CHANNELS** to the topic.

First, we present a tutorial article as a refresher on major graphics hardware and software ideas and technologies making computer graphics an indispensable tool. Then, we focus on graphics applications in the petroleum industry and filmmaking, and we take a look at how two other CRAY users are applying graphics techniques to their particular research efforts. Cray Research surveys the Los Alamos graphics network and finally, we present an interview with graphics expert Nelson Max of Lawrence Livermore Labs.

Although our coverage of computer graphics in this issue is by no means exhaustive, we hope the articles stimulate your interest and curiosity, because from petroleum exploration to the electronic artist's palette, computer graphics is becoming more a part of your life every day!

About the cover

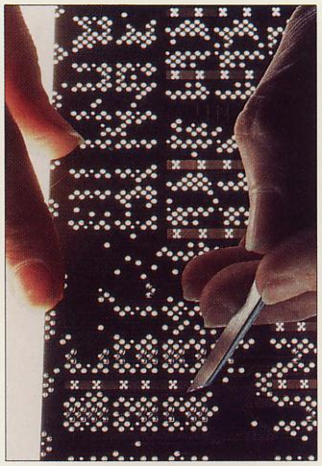

Printed Circuit Boards undergo stringent testing

CRAY printed circuit boards (PCBs) undergo several 100% inspections during the assembly process. Forty percent of the direct labor involved in manufacturing CRAY PCBs is committed to quality inspection. Final visual inspection shown here is done under magnification by one of Cray's 14 final-check inspectors. These same printed circuit boards will again be checked for quality at the company's manufacturing facility prior to the system assembly process. (Photo credit: Joe Gianetti)

CRAY CHANNELS
A Cray Research, Inc. publication Volume 5, Number 1

CONTENTS

FEATURE ARTICLES

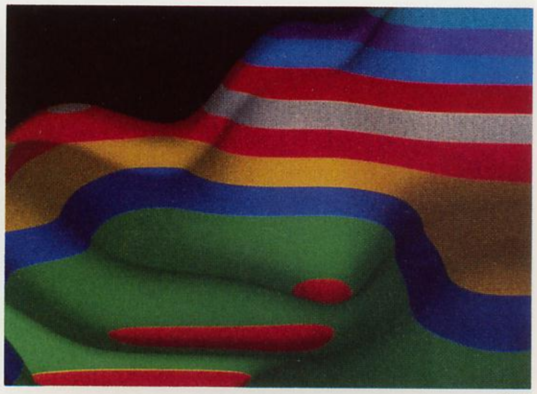

22

REGULAR COLUMNS

Staff

Editor Tina M. Bonetti
Assistant editor Carol A. Smith
Art/Design Chad Jewett
Kathleen Van Kirk
Jeanne M. Jacobowski

CRAY CHANNELS is a quarterly publication of Corporate Communications, Cray Research, Inc., P.O. Box 154, Minneapolis, Minnesota 55440. CRAY CHANNELS is intended for users of Cray Research computer systems and others interested in the company and its products.

Requests for copies of CRAY CHANNELS may be addressed to the Corporate Communications Department of Cray Research at the above address. If you have a feature article or news item to submit for publication, please contact the editor of CRAY CHANNELS either by mail or by telephone at (612) 333-5889.

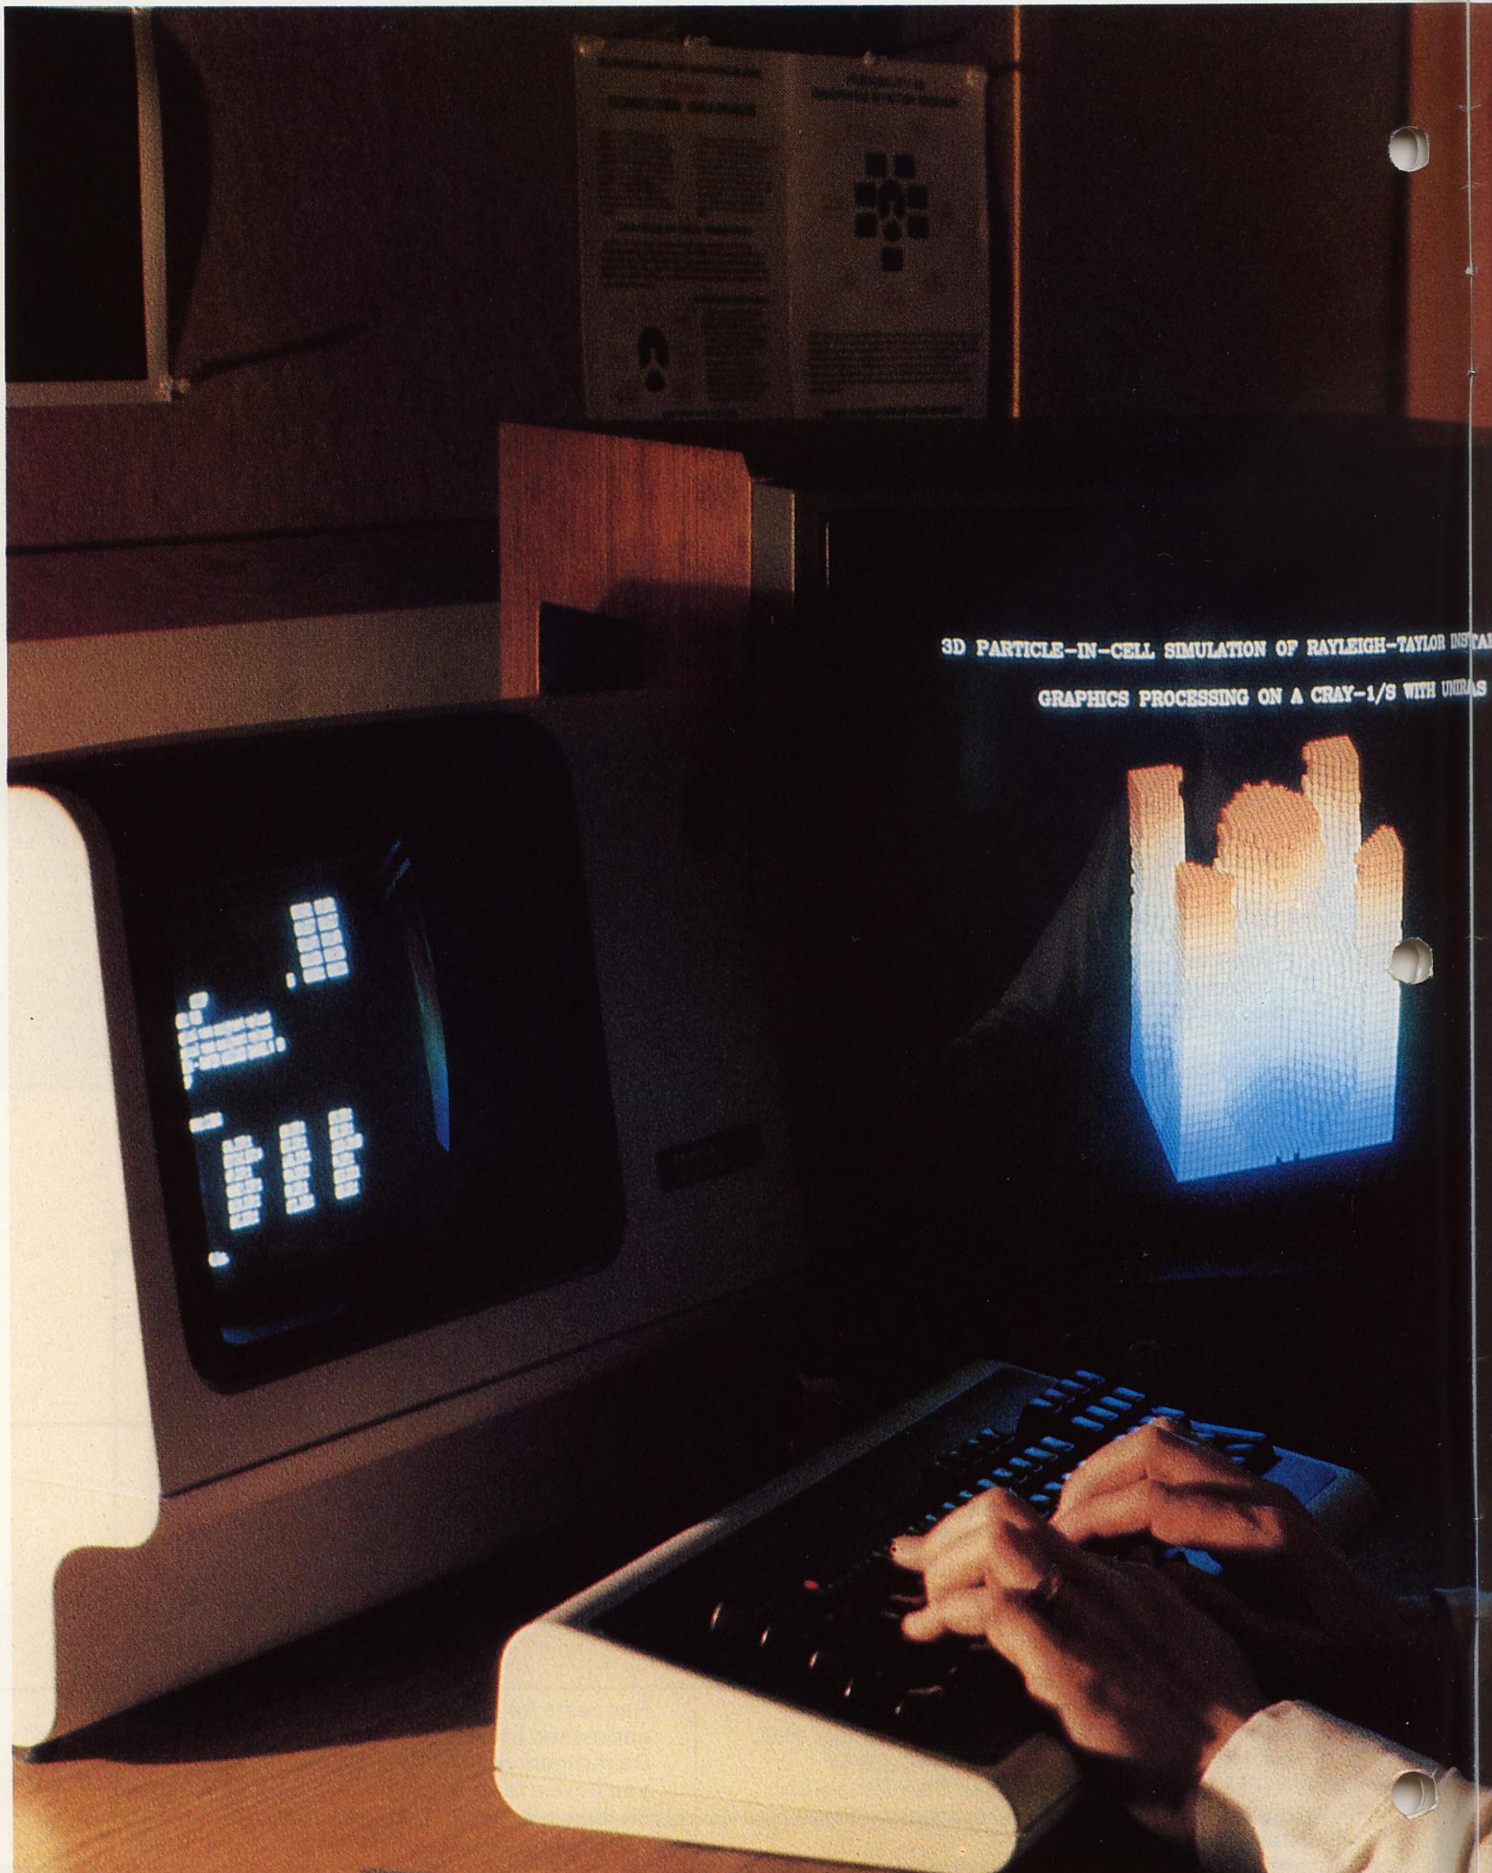
3D PARTICLE-IN-CELL SIMULATION OF RAYLEIGH-TAYLOR
GRAPHICS PROCESSING ON A CRAY-1/S WITH

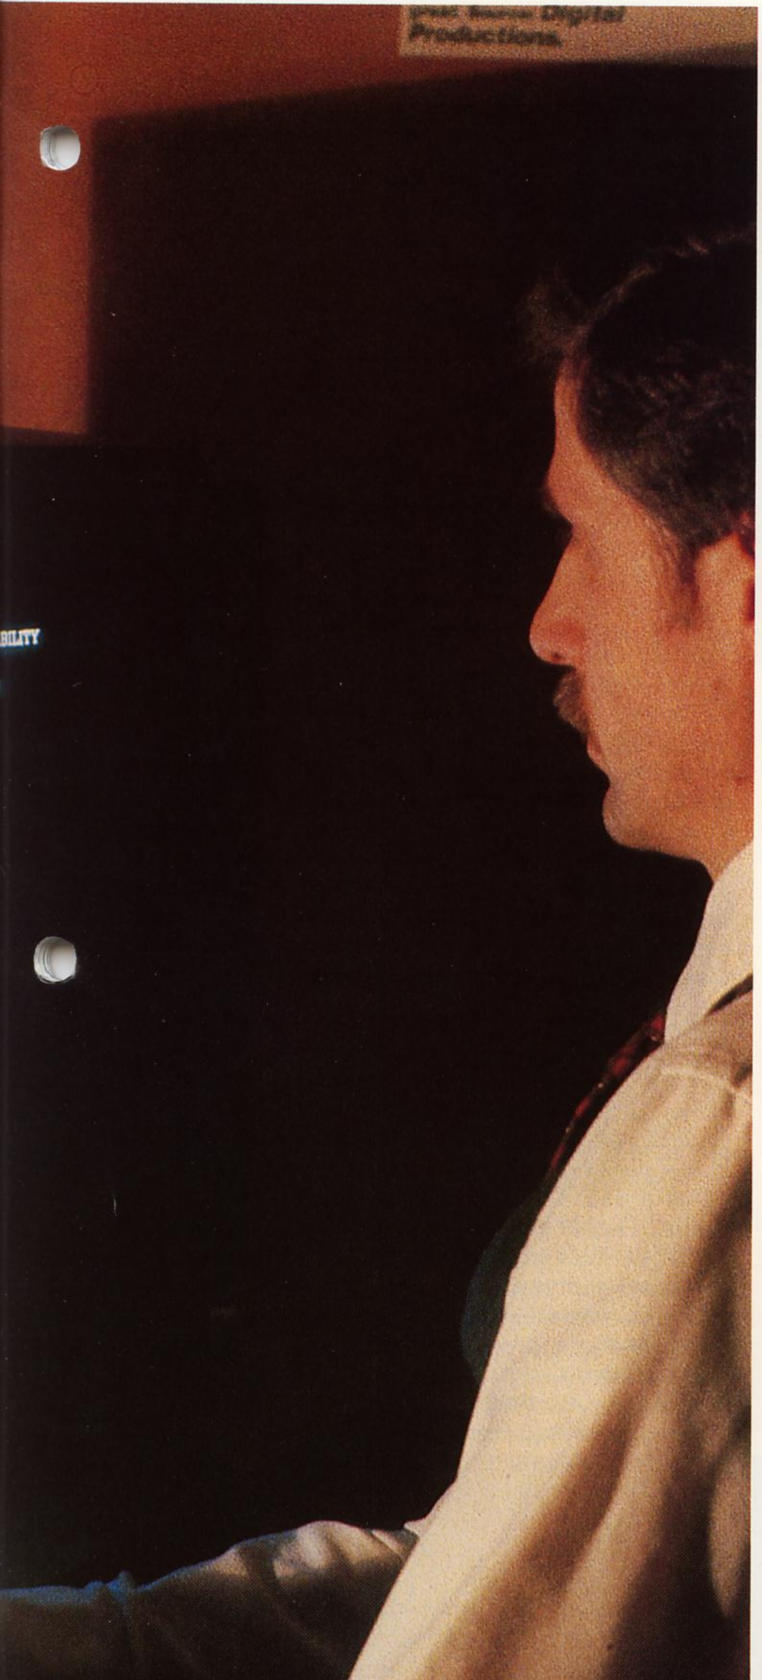

SKETCHPAD realized: an introduction to computer graphics

Computer graphics impacts all facets of our lives today, in everything from our most trivial entertainment to the most important research projects. Actually it comes as no surprise that practical computer graphics has carved its niche so rapidly and pervasively among computer and non-computer users alike. Graphics enables us to take advantage of our well developed eye-brain pattern recognition capabilities that allow us to process many different types of data rapidly when presented pictorially. Unfortunately, creating and reproducing a meaningful picture is not as easy as it may sound. The concept of graphic output from computers is not much newer than computers themselves, but it was not until the 1970's that graphics began to be practical. This tutorial article serves as a refresher of the major hardware and software ideas and technologies that make computer graphics an indispensable tool of research, business and daily life alike.

Historical background

The rudiments of computer graphics can be found in early plotting hard-copy devices such as teletypes and line printers dating from the 1940's. In fact, the Whirlwind I computer at Massachusetts Institute of Technology (MIT) had computer-driven cathode ray tubes (CRTs) displaying output in

1950. But it was in 1962 that the single most important development in computer graphics came when Ivan E. Sutherland, a young Ph.D. from MIT, described a system that would allow one to draw using a computer. His dissertation entitled, "SKETCHPAD: A Man-Machine Graphical Communication System" first called attention to the possiblity of practical computer graphics.

Real progress in computer graphics and interactive systems was less than dramatic in the 1960's. Implementation of interactive graphics remained beyond the resources of all but the most advanced technological organizations. There were many reasons for the sluggish development. The high cost of graphics hardware was only part of it. Computing resources required to support large databases, as well as interactive picture manipulation and postprocessing applications programs whose input came from the graphics design phase were very costly. In addition, writing interactive graphics programs was entirely new for programmers accustomed to writing FORTRAN programs for batch environments. And display device-dependent software involved large expenditures that companies could ill afford.

Significant strides in these and other areas during the 1970's brought computer graphics into the realm of feasibility. Among the most important strides made were those in output device technology. The availability of displays played two important roles in advancing computer graphics: (1) Previously generated images could now be previewed and modified by the user prior to generating hardcopy output, and (2) the user could interactively create and modify images. The two types of displays that dominate the market today are vector and raster video displays.

Output devices

Vector displays

The vector refresh, or calligraphic, display, which became popular in the mid-sixties, is the oldest type of display in use today. The vector refresh display's generated image fades quickly, and the screen must continually be rewritten or "refreshed", typically 30 times a second to prevent flickering. Figure 1 shows the major components of this system. The movement of an electron beam across the surface of a cathode ray tube is controlled by signals from a display processor that receives instructions stored by the user in a display list. These signals control the locations of the beam and can draw from one arbitrary location x_1y_1 on the screen to another x_2y_2.

Continual refreshing of the display limits the number of vector instructions that can be accommodated, but allows the user to modify or sequentially change an image. This important capability is used in 3-D designing but, requires large bandwidth communication channels from the CPU and expensive CPU time. Buffer memories alleviate that penalty and intelligent terminals in newer systems store model descriptions locally while the main computer sends only the information necessary to transform the views.

The vector storage display is a variation of the same basic display. However, its major distinction is that the image does not require constant refreshing. The image remains constant until permanently erased.

Raster displays

Computer graphics based on the raster display is becoming increasingly prevalent. The ability to generate full color and fill areas, gives raster displays distinct technical advantages. Not only that, raster displays rely heavily on integrated circuits which offer increasing compute power and declining costs of memory. Figure 2 illustrates a typical raster graphics system. Its major components are a scan conversion processor, a frame buffer, a display processor and a CRT.

In raster displays, an electron beam sweeps across the CRT in a predefined descending horizontal pattern from left to right. Each scan line on the screen is broken up into picture elements (pixels). A typical system has 512 scan lines with roughly 512 pixels per line, an arrangement that may be thought of as a 512 x 512 grid of a quarter of a million points. Higher resolution screens can range as high as 1024 x 1024 or 4096 x 4096. Under the control of a computer program, the intensity and color of the beam is controlled at each pixel during a scan to produce the appropriate coloring and shading. In some systems, up to 2^{24} different colors can be generated. As with vector display systems, the image generated by a raster display must be

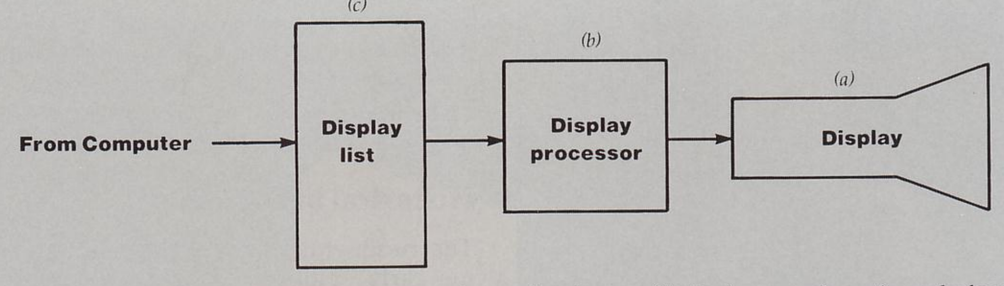

Figure 1 In a vector refresh display, an electron beam excites a phosphor on (a) the inner surface of a cathode ray tube (CRT), producing a glowing image on the face of the display. After a brief period, the image fades and must be "refreshed" by the electron beam. The path of the electron beam is controlled by signals from a display processor (b). By translating instructions stored in a display list (c) by the computer, the display processor switches the beam on and off, controls the generation of the image on the CRT, and controls the refresh of the displayed image.

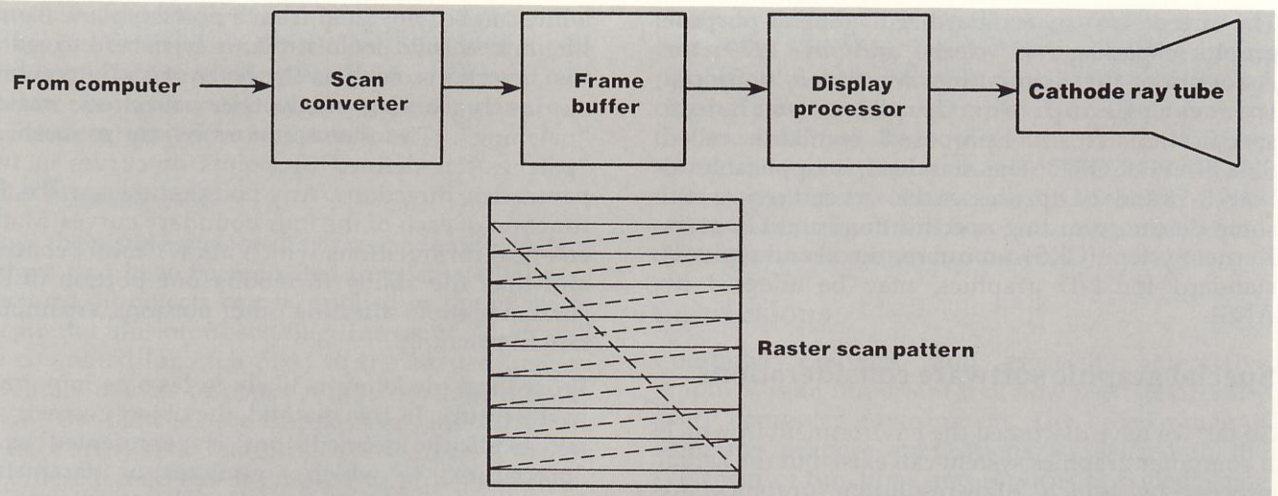

Figure 2 A raster display system. Information about the intensity or color of each picture element (pixel) is stored in the frame buffer, which is roughly a 500 x 500 grid of memory elements. Information in the frame buffer is read 30 times a second by the display processor, which generates appropriate signals to control the electron beam in the CRT. Information needed to draw a vector must be stored in the frame buffer as a set of pixels. Conversion of information from vector to pixel format suitable for display is carried out by a scan converter, located between the computer and frame buffer.

refreshed, and is done so at a typical rate of 30 times per second.

Many types of noninteractive graphics output devices that produce hardcopy output play an important role in computer graphics. Given a wish list, most users of computer graphics would opt for a hardcopy device that is inexpensive and fast, with high resolution, color capability and a high degree of accuracy. The hard facts of life are that such a device does not exist. However, many devices are available that combine some of these features, enabling the user to satisfy individual needs. Pen plotters, electrostatic plotters and computer output on microfilm recorders are all popular devices. Direct image recorders are commonly used to obtain hard copy either from a displayed image or from an electronic signal that produces the image.

Input devices

In many computer graphics applications, an essential element is the on-line interaction of a user with computer-generated images on a display. Many types of graphics input devices have been introduced because of the inherent difficulty of efficiently supplying graphics information to a computer. Of course information can be input through the keyboard, but that can be awkward. A popular specialized device is the light pen, commonly used for pointing at a screen to initiate an action and draw many data points. The pen is actually a photocell device. When pointed at the screen of a CRT, the graphics system identifies the photocell's position. Light pens do not work with storage-type CRTs because the image is painted only once, and no moving beam appears under the pen when the user points it. Some other commonly used devices are joysticks, track balls, and the mouse, all devices that control a cursor or crosshair on the screen. They are used to point to x,y coordinate information that is conveyed to the program.

Software

Software is a crucial, complex and expensive component of graphics applications. The last few years have seen substantial progress in developing graphics software principles that enable applications to be programmed more efficiently. A very important principle is the separation of the graphics software from the application, resulting in modular programs with well-defined interfaces. This allows a commercial graphics application package to be mated with a custom application. Additionally, graphics software ideally is written to maximize device-independence, allowing flexibility in the use of hardware devices. Usually the applications programmer accesses graphics subroutines compatible with FORTRAN or Pascal. Figure 3 illustrates a typical graphics software system.

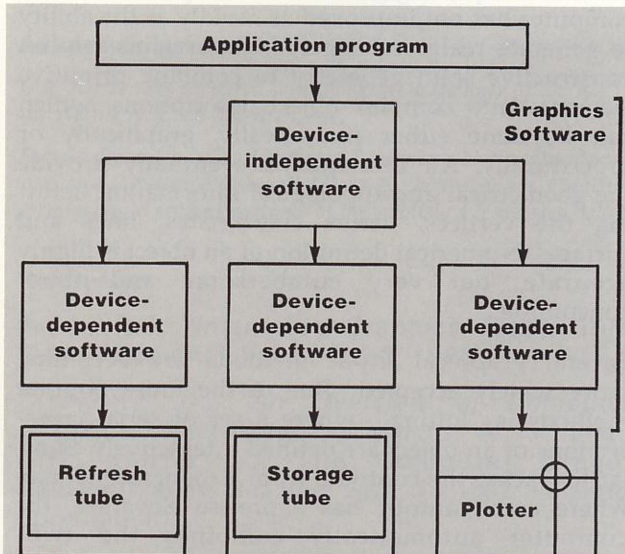

Figure 3 Modular programming techniques permit separation of the application program from more generalized graphics software with device-independent features.

The need for a standardized general-purpose graphics package is clear, and in 1979 the Association for Computing Machinery's Special Interest Group on Graphics published specifications for a proposed standard called SIGGRAPH/CORE. The standard is applicable to both 3-D and 2-D processes. However, there is still some debate regarding specifications and Graphics Kernel System (GKS), an international and national standard for 2-D graphics, may be adopted by ANSI.

Special graphic software considerations

So far, we have discussed the environment in which a computer graphics system can exist, but the actual process by which high-resolution graphics are generated is an entirely different subject. A discussion of specialized software techniques to produce such output follows below. The software that actually produces images is the result of clever programming based on the laws of optical physics.

Five of the major tasks involved in creating a continuous-tone computer image of a simulated object are outlined below. First, a mathematical description of the object must be defined and input into the computer. Second, the three-dimensional description must be transformed into a two-dimensional perspective image. Third, a determination of all visible lines or surfaces must be made. Portions of the scene that are outside of the cone of vision are removed and lines or surfaces not seen by the observer are eliminated. Fourth, an illumination model determines the color or shade of each surface or part of the environment. Fifth, the appropriate red, green and blue intensities must be selected to represent the color specified by the shading model.

Describing the object

The art of inputting geometric descriptions into the computer has not improved as rapidly as the ability to generate realistic images. Most systems rely on constructive solid geometry to combine primitive elements into complex object descriptions, which can be done either numerically, graphically or procedurally. All descriptions eventually provide the geometrical and topological information defining the vertices, vertex coordinates, lines and surfaces. Numerical definition of an object is highly accurate, but very cumbersome and time-consuming.

Several graphical input methods are becoming more widely accepted. One of the more popular methods is "lofting", where a set of serial cross-sections of an object are defined interactively. Similar to tracing the contours from a topographic map where each contour has a precise elevation, the computer automatically combines the two-dimensional definitions to create three-dimensional shapes. A second graphic approach can be thought of as an extrusion method, typically called the "sweep representation method". Its premise is that a line can be generated from a point, a plane from a line and a solid from a plane. Surfaces curved in two directions, such as the body of an automobile, typically employ another method called "patching". The surface, formed by a mesh of "patches", is defined by points or curves in two parametric directions. Any point on the surface is a function of each of the four boundary curves. Mathematical formulations which allow "local" control, and thus the ability to modify one portion of the surface without affecting other portions, are important in patching.

Procedural modeling is likely to become important in the future. In this method, the object geometry is not explicitly defined, but is represented as a "procedure" to which arguments or parametric variables are passed. For example, stairs may be created with only a few input variables, such as the height, width and stair type.

Perspective transformations

Once the task of inputting the object description is completed, the image is generated on the display. With the input methods discussed thus far, each point is located in space by its three coordinates in the object coordinate system. The 3-D description of the environment to be displayed is transformed into a 2-D perspective image by determining the intersections of the view rays emanating from the observer's eyes with an imaginary picture plane. The view ray is defined as the line connecting a point in the environment with the eye of the observer. The picture plane is an imaginary plane, perpendicular to the line of sight, and a fixed distance from the observer. Then, the position of a point on the display is mathematically calculated to correspond to a point on the object by transforming the point from the object coordinate system to the eye coordinate system, where the origin is fixed at the viewpoint. After the perspective transformation is completed, portions of the image that would extend beyond the edges of the display device are clipped.

Hidden line — hidden surface algorithms

After the perspective computations, all points in the environment have been transformed. However, the computer does not know which portions are not visible. Surfaces hidden from view cannot be displayed. The most difficult and computationally expensive portions of the software procedures is the solution of the hidden line or hidden surface problem. In fact, many applications such as computer-aided design do not eliminate lines that should be hidden because of the overhead involved.

There are more than a dozen clever algorithms for hidden surface removal tailored for special applications. For vector displays, the task is relatively easy. It is only necessary to determine the end points of the visible line segments. For raster displays, the majority of the algorithms compute the visibility at each pixel of the image plane. For a polygonally defined environment, the basic task entails a comparison of each polygon to every other polygon to determine which one is closest to the observer. The more complex the environment, the greater the number of polygons to be compared and the more extensive the computations.

Shading models

Once the visible surfaces have been identified, algorithms that have been created to calculate light reflections off objects can be applied to the image to create the illusion of shading. Intensity values will be computed for each pixel in the shaded images. Accurate models of object reflectivity take into account the light source composition and direction, surface orientation and surface properties. A variety of different methods are used to approximate shading based on light physics. Algorithms have been developed to account for the directional distribution of reflected light on a wavelength basis as a function of surface roughness, slope, the material properties and reflection geometry.

Other "ray-tracing" schemes combine the hidden surface algorithm and shading model. Since the intensity of each pixel in the image must be calculated, a "tree of rays" is extended from the viewer, through each pixel, to the first surface encountered. From there, for each ray, the reflection directions are calculated, and the ray continues to other surfaces and light sources. The global illumination information allows accurate simulations to be created.

Color

Finally, color is added to the image by specifying the appropriate blue, green and red values to obtain the correct hues. The CRT has three electron guns, for each of the primary colors. Each pixel is composed of triads of red, green and blue phosphor, and color is produced for each pixel by simultaneously exciting each of the three phosphors by controlling the voltages to each of the three electron guns.

As long as the phosphor dots are spaced closely enough, and the observer is far enough away from the screen, the eye integrates the three luminances into a single color sensation. To produce realistic colors for an image, one must determine the magnitude of electron voltages. The resulting color must approximate the color of the same image as if it were to appear in the real world.

Image enhancements

Several methods exist to further improve the realism of computer-generated images. One important enhancement, the removal of aliasing artifacts, is described below.

In computer graphics, aliasing appears in raster displays because the image is generated as a sequence of discrete sample points, rather than as a continuous signal. The most common aliasing artifact is jaggedness in line edges or in polygon boundaries. One way of reducing the aliasing problem is to increase the sampling rate by increasing the image resolution. A better way to antialias is to filter the original function. For example, if an edge intersects the area defined by a pixel, the resultant color should be a weighted average of the colors of the two polygons on either side of the edge. Until recently, filtering of this sort was too computationally expensive to be used commonly.

Conclusions

Computer graphics, and especially interactive graphics, is an important tool now used in virtually every computer environment. The broad range of computer graphic applications is realized in the spectrum of functions and services provided, some of which are listed below:

- Graphics display of computation results as they occur
- Rapid presentation of large quantities of information
- Visualization of objects for rapid study of design options
- Enhancement of interpretation and impact of data
- Process simulation and verification before committing real resources
- Graphic representation and observation of simulated theoretical models
- Creation of artistic designs
- Entertainment.

To address each of these diverse applications, a plethora of specialized hardware and software has been developed. This brief tutorial has addressed the broad fundamentals of those hardware, software and graphic-specific software systems that can be applied to most of the application environments.☐

Acknowledgements

This article was prepared with the kind assistance of Dr. Richard Weinberg, Cray Research, Inc.

Thanks are extended to Lawrence Livermore National Laboratory and Addison-Wesley Publishing Company for granting permission to reprint portions of the published text listed below.

References

Burchi, R.S., *"Interactive Graphics Today"*, IBM Systems Journal, Vol. 19 No. 3, 1980.

Greenberg, D., *"An Overview of Computer Graphics"*, Greenberg, D., Marcus, A., Schmidt, A.H., Gorter, V., The Computer Image, Addison-Wesley, Reading, MA., (© 1982, Polaroid Corporation).

"Rasters, Pixels, and Joysticks: An Introduction to Computer Graphics", Energy and Technology Review, Lawrence Livermore National Laboratory, Livermore, CA., November, 1981.

Computer graphics for oil & gas exploration

John E. Aldag, Cray Research, Inc.

Petroleum exploration is an on-going cyclical process of building on the knowledge gathered to date about an area suspected of harboring a hydrocarbon reservoir (a prospect). It requires the efforts of many people with varied expertise working in many unique jobs. Geologists, geophysicists, geochemists, computer analysts, technicians, operaters and seismic acquisition crews are among the many professionals involved. They use many highly specialized tools in probing, analyzing and picturing the subsurface. Computers and computer graphics have been an important part of this process since the technology first became available. The combination of the professionals' dedicated efforts and powerful tools contribute to the common objective — obtaining a complete understanding of the geophysics, geology and geochemistry of the prospect.

Several factors point to continued and increasing use of computer graphics in the search for hydrocarbons. Exploration is taking the explorationists into increasingly remote and hostile environments such as the North Sea, Alaska, Saudi Arabia, far offshore Texas, and mountainous areas in the U. S. and abroad. As a result, information about the subsurface is more difficult, and thus, more expensive, to acquire. At the same time, the geology of some recent prospects is more complex than that of classic reservoirs. The information gathered about these newer locations is much more difficult to interpret and requires powerful analytical and display tools in order to obtain an accurate interpretation.

The cost of drilling exploratory or production wells in remote areas such as the North Slope of Alaska or Texas demands that the understanding of the prospect area be as complete as possible. The costs of a "dry hole" can run into tens of millions of dollars. Usually the geoscientist has the difficult task of piecing together incomplete information to predict the geology of the area. For instance, information suggesting a favorable geology above the target prospect depth may heighten suspicions. But inconclusive data at the depth of interest may cloud the interpretation. Because there is not a one-to-one relationship between the recorded phenomenon and the actual geology, available data may be misinterpreted.

Tremendous volumes of information are gathered at all stages of exploration. The data must be acquired, processed, analyzed and interpreted at each step. Powerful data processing capabilities complemented by sophisticated high resolution displays provide more and better information about prospects than has ever been available. Through the use of these tools, explorationists are able to make more accurate and timely decisions regarding the potential of prospect areas. In this article, we will trace the role of computer graphics in oil exploration from the evolution of the original "hunch" to the final and only conclusive test of the geoscientist's skill—putting the drill to the ground.

Landsat imagery

Computer graphics can come into the exploration process at a very early stage. In areas where no other data is available, pictures of the land surface from an aircraft or space satellite (Landsat images) may reveal a surface feature that catches the attention of a geologist. Image processing, involving the analysis and enhancement of surfaces from Landsat pictures is often an initial step.

Geologists can enhance and emphasize the features of interest and display them on high resolution color graphics systems. Information about the geology, the soil, vegetation and wetness conditions can be highlighted. Sediment areas of river deltas are particularly interesting as possible target areas. A linear feature in the image (linament) may suggest an abrupt change in the surface composition, indicating a possible fault line. Figure 1 shows such a display. Because digital pictures can be very large for a prospect area, a powerful processing computer is required to work on the data enhancement in a timely manner. Landsat imagery is a timely supportive tool in onshore exploration. Increasing resolution of graphics tools enhances the interpretation and the utility of Landsat imagery.

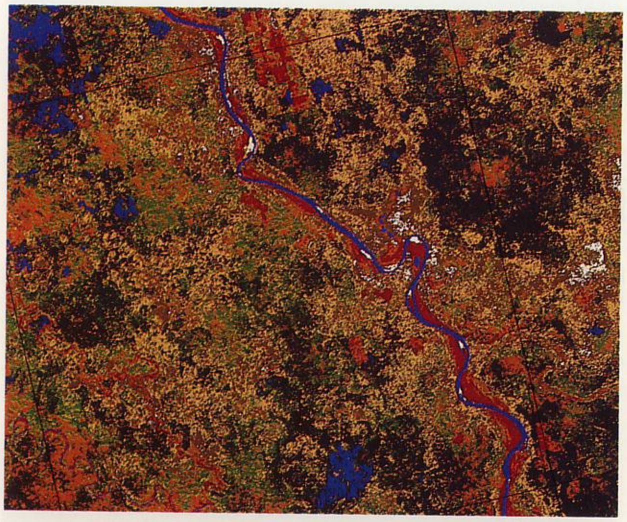

Figure 1 Landsat image identifying vegetation and ground conditions in the Blue Nile area of Sudan.

If the photogeology of the prospect is promising, a geologist may go to the field to study the surface first hand. He or she looks for tilted rock layers or major cracks in the surface crust and attempts to correlate these with features on the photos or images of the area. Samples of the surface rock may be gathered and taken back to a laboratory for analysis. With these added bits of information, the geologist begins to formulate a hypothesis of the region's underlying geology. If that hypothesis includes the possibility of an oil or gas deposit, further data must be gathered to test and modify the idea.

Boreholes and well logs

Other wells in the area can also provide valuable information to the geoscientist. If not, one or more boreholes may be drilled in the area at strategic locations. A borehole is usually only a few inches in diameter and a few thousand feet in depth, but it provides a valuable glimpse of the subsurface. Samples of rock brought to the surface (cuttings) and cylinders of rock cut at strategic depths (cores) give the geologist a record of the substrata at the prospective well location. During or after the drilling, instruments that measure rock properties are lowered into the hole and accumulate information about the rock layers' electrical and acoustic properties. After some processing of the raw data, it is plotted in the form of a line drawing on a display terminal or paper. Substrata information about the rock types, porosity and fluid content can be extracted in this way. These "well log" graphs are valuable records of information and are kept under tight security by the company that owns them. Although a simple line drawing, this record of the subsurface requires an experienced analyst's interpretation.

Graphics for seismic processing, analysis and interpretation

If, after all of these bits of data are collected, the geoscientist's hypothesis still allows for the possibility of a hydrocarbon reservoir, and if management comes to the same conclusion, the most definitive and expensive exploration technique is brought into play — seismic. The decision to acquire and process seismic data for the area is not made lightly because it is extremely costly and time consuming. But if the data is acquired carefully and processed well, the result is the closest thing to a picture of the subsurface available. Large-scale computers such as the CRAY are being used by the petroleum industry to process huge amounts of seismic data. High-level computer graphics capabilities play an important role in portraying and highlighting this data for interpretation.

Seismic exploration involves the acquisition, processing and analysis of acoustic signals reflected from the earth's subsurface. To acquire the signals, a controlled shock or sound wave is induced into the ground or ocean. At each subsurface change from one rock type to another, some of the signal is reflected. The time required for the signal to reach the subsurface interface and return to the surface is related to the depth and shape of the underground structure. Listening devices, systematically spaced in a line, pick up the minute shock signals which have been generated by the source and reflected from below. The returned signals are called seismic reflections, the strongest of which are called events. Using the reflection data geophysicists interpret the rock properties of the prospect area. It is from this information that hypotheses about the prospect area are made and refined.

Acquisition, processing and interpretation of seismic data is a monumental task. Specialists in each of these separate activities work together in obtaining results. Because of the sheer bulk of data to be processed and factored into an interpretation, seismic processing is largely done in batch mode. Even for a small set of 2-D data, tens of millions of data points will require processing. And for 3-D seismic, the task increases more than proportionately. While the 2-D seismic procedures can be thought of as picturing a vertical cut into the prospect area, a 3-D survey results in a spatial coverage of the area. 3-D survey data can be thought of as a block of numbers representing the 3-D structure of area of interest. Large-scale computers such as the CRAY can be used in processing and analysis.

During processing, the evolving digital data may be graphically displayed to evaluate previous processing steps and suggest the next. However, it is after processing that interactive and color graphics play their most important role. Seismic data can be displayed in different graphic formats. Traditionally, seismic data has been output to hardcopy black-on-white raster devices. However, interactive, high resolution color graphics enable efficient, effective analysis that was not previously possible.

Interactive computer programs can allow the interpreter to analyze the data, look for trends and discontinuities and solidify his or her understanding of the geology. In the case of 3-D analysis, the interpreter can look at horizontal and vertical cuts of the block of data representing a volume of the earth. Changes in the subsurface rock types, structure and composition can be inferred from different vertical and horizontal time slices of the subsurface.

Figures 2 and 3 illustrate examples of current graphics output. Figure 2 is a 2-D cross section of seismic data. The strong linear "squiggly" features called events, suggest a change in the rock type and color has been added to indicate reflection strength. Figure 3 illustrates horizontal cuts or "time slices" that suggest the changing subsurface structure at progressively deeper depths. By cutting horizontally through this stack of data and displaying a sequence of these "time-slices" in a movie format, an interpreter can infer subtle changes in the geology from the changing patterns of the seismic data. The strong negative signals are shown in blue, the strong positive in red. Note the movement of the event to the lower right (southeast) as the cut is taken at deeper levels.

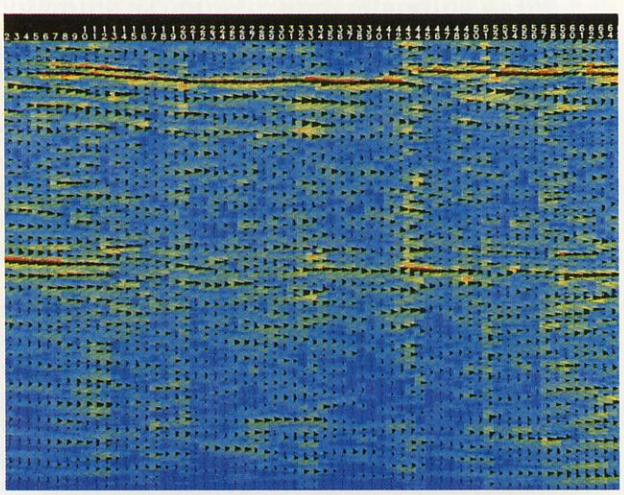

Figure 2 A north-south section of experimental 3-D survey data generated on the CRAY.

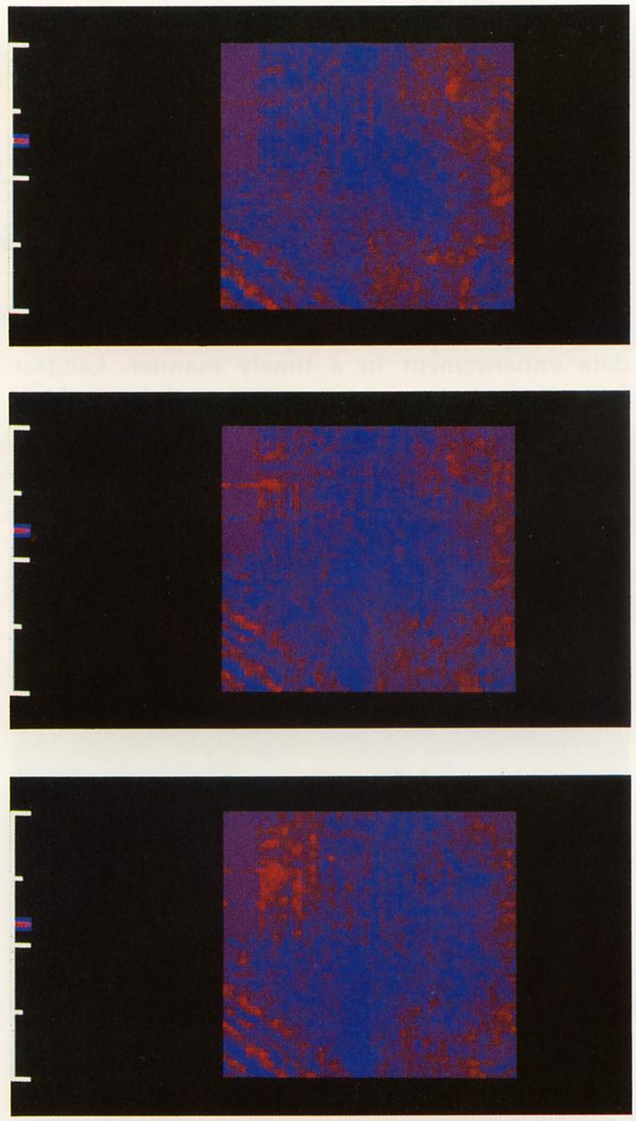

Figure 3 A sequence of horizontal 3-D "time slices" generated on the CRAY.

High level graphics for clearer understanding

For extremely complex geologies, the interpreter may find it useful to test and validate hypotheses about a candidate geology by generating a model of it based on the available seismic data. Because the physics of the seismic process is well understood, it can be simulated with numerical and mathematical approximations. The results of these calculations are graphic displays of: 1) the paths taken by the seismic energy through the model geology and 2) simulated or "synthetic" seismic data. An example of the first display is shown in Figure 4. The ray paths of the resulting reflections from a subsurface structure are shown. The display points out a common pitfall in areas of complex geology. Sideswipe reflections, as they are called, are misleading reflections that can confuse the interpretation of the seismic data. The raypaths of a sideswipe event are the right most traces illustrated in Figure 4. It is the interpreter's job to recognize and understand these 3-D effects.

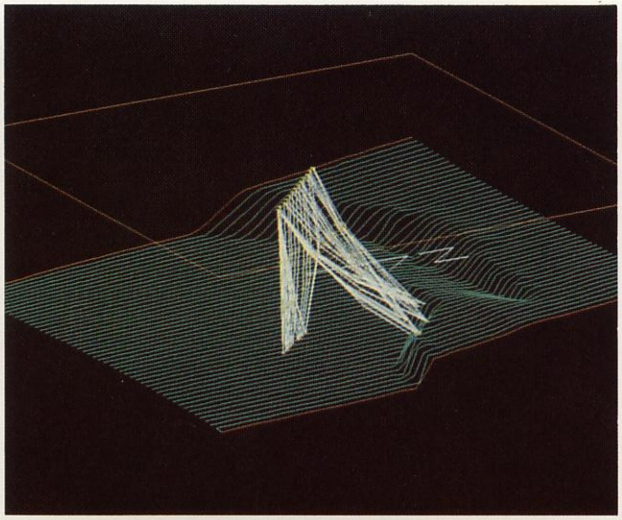

Figure 4 Graphic display of 3-D ray tracing of seismic energy through a model geology generated on the CRAY.

By trial and adjustment and display of the synthetic seismic data to compare to the actual data from the field, the interpreter can eliminate geologies that could not have caused the features observed on the seismic data and refine those that could. This forward modeling approach is often a highly interactive process, and requires a quick estimation of the seismic response to a given trial geology. CRAY computers coupled with powerful graphics can provide the necessary responsiveness required for effective implementation of such 3-D modeling. A CRAY can perform the required calculations over 30 times faster than equipment typically used in the past for the task. For an interpreter, this means wasting seconds rather than minutes for the results of the simulation.

In the final stages of interpretation, the geoscientist will prepare contour maps of the prospect area. A computer often can provide a first and unbiased "opinion" of the shape of the geologic structures based on an incomplete set of data points. A computer contour map is shown in Figure 5. The various colors indicate areas of approximately equal depth. The areas are separated by contour lines of equal depth.

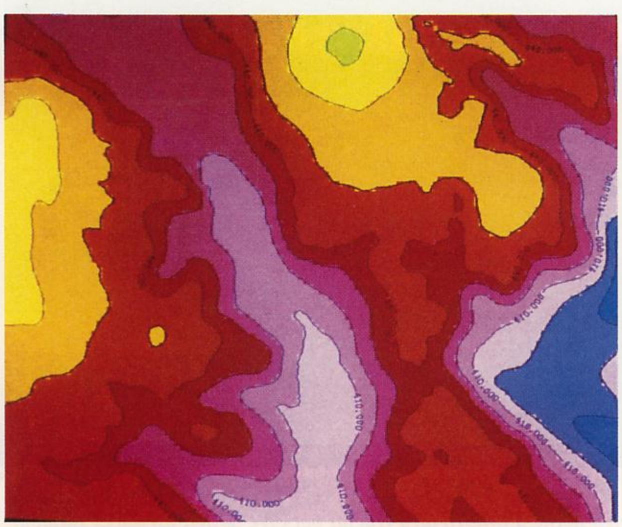

Figure 5 Sample color contour map.

After some fine tuning by the interpreter, the resulting map can be used by the computer to provide unbiased estimates of reserves and other important information for production of the newly found reservoir. The computer's unbiased analysis of reserves is especially important where the reservoir crosses lease or country boundaries. In these circumstances, the companies involved will cooperate to produce a field in the most productive way but will, of course, distribute the production according to the fraction of the field within each lease boundary.

Conclusion

All of this work leads up to the final key decision— to drill or not. After months of effort, an interpreter may decide that the data available is not conclusive or that the prospect is not large enough to make drilling economical at present oil prices.

But the exploration process does not end here. If the prospect is not drilled or is dry, the effort has still provided a wealth of information which will be factored into the next prospect in the area. If the work results in a discovery, all of the exploration data becomes valuable in the production of the reservoir as well as useful for future exploration. Computers and computer graphics are also effective tools in production. But that is an entirely separate topic.

Computers and computer graphics are playing an increasingly important role in the lengthy, complex exploration process. The value of the vast amounts of data acquired about prospect areas could be questionable without powerful tools to help geoscientists understand it. As data processing and display capabilities continue to increase, and as the expense of exploration also increases, computing and graphics, as processing and interpretive tools will play an even more important role in oil and gas exploration.□

About the author

John Aldag has been with Cray Research, Inc. since early 1982, managing graphic applications development. From 1978 to 1982 he worked for the exploration research and development division of CONOCO developing computer applications for the interpretation of geoscience data. John received his Ph.D. in physics from the Univeristy of Missouri, Rolla in 1980.

Acknowledgements

Figure 1 generated by ESL Inc., for Chevron Oil Company of Sudan.

Figures 2 and 3 generated by European Software Contractors, Inc. on the CRAY.

Figure 4 generated by Sierra Geophysics on the CRAY.

Figure 5 generated by Radian Corp.

Seismic data on page 8 courtesy of GeoQuest International, Inc.

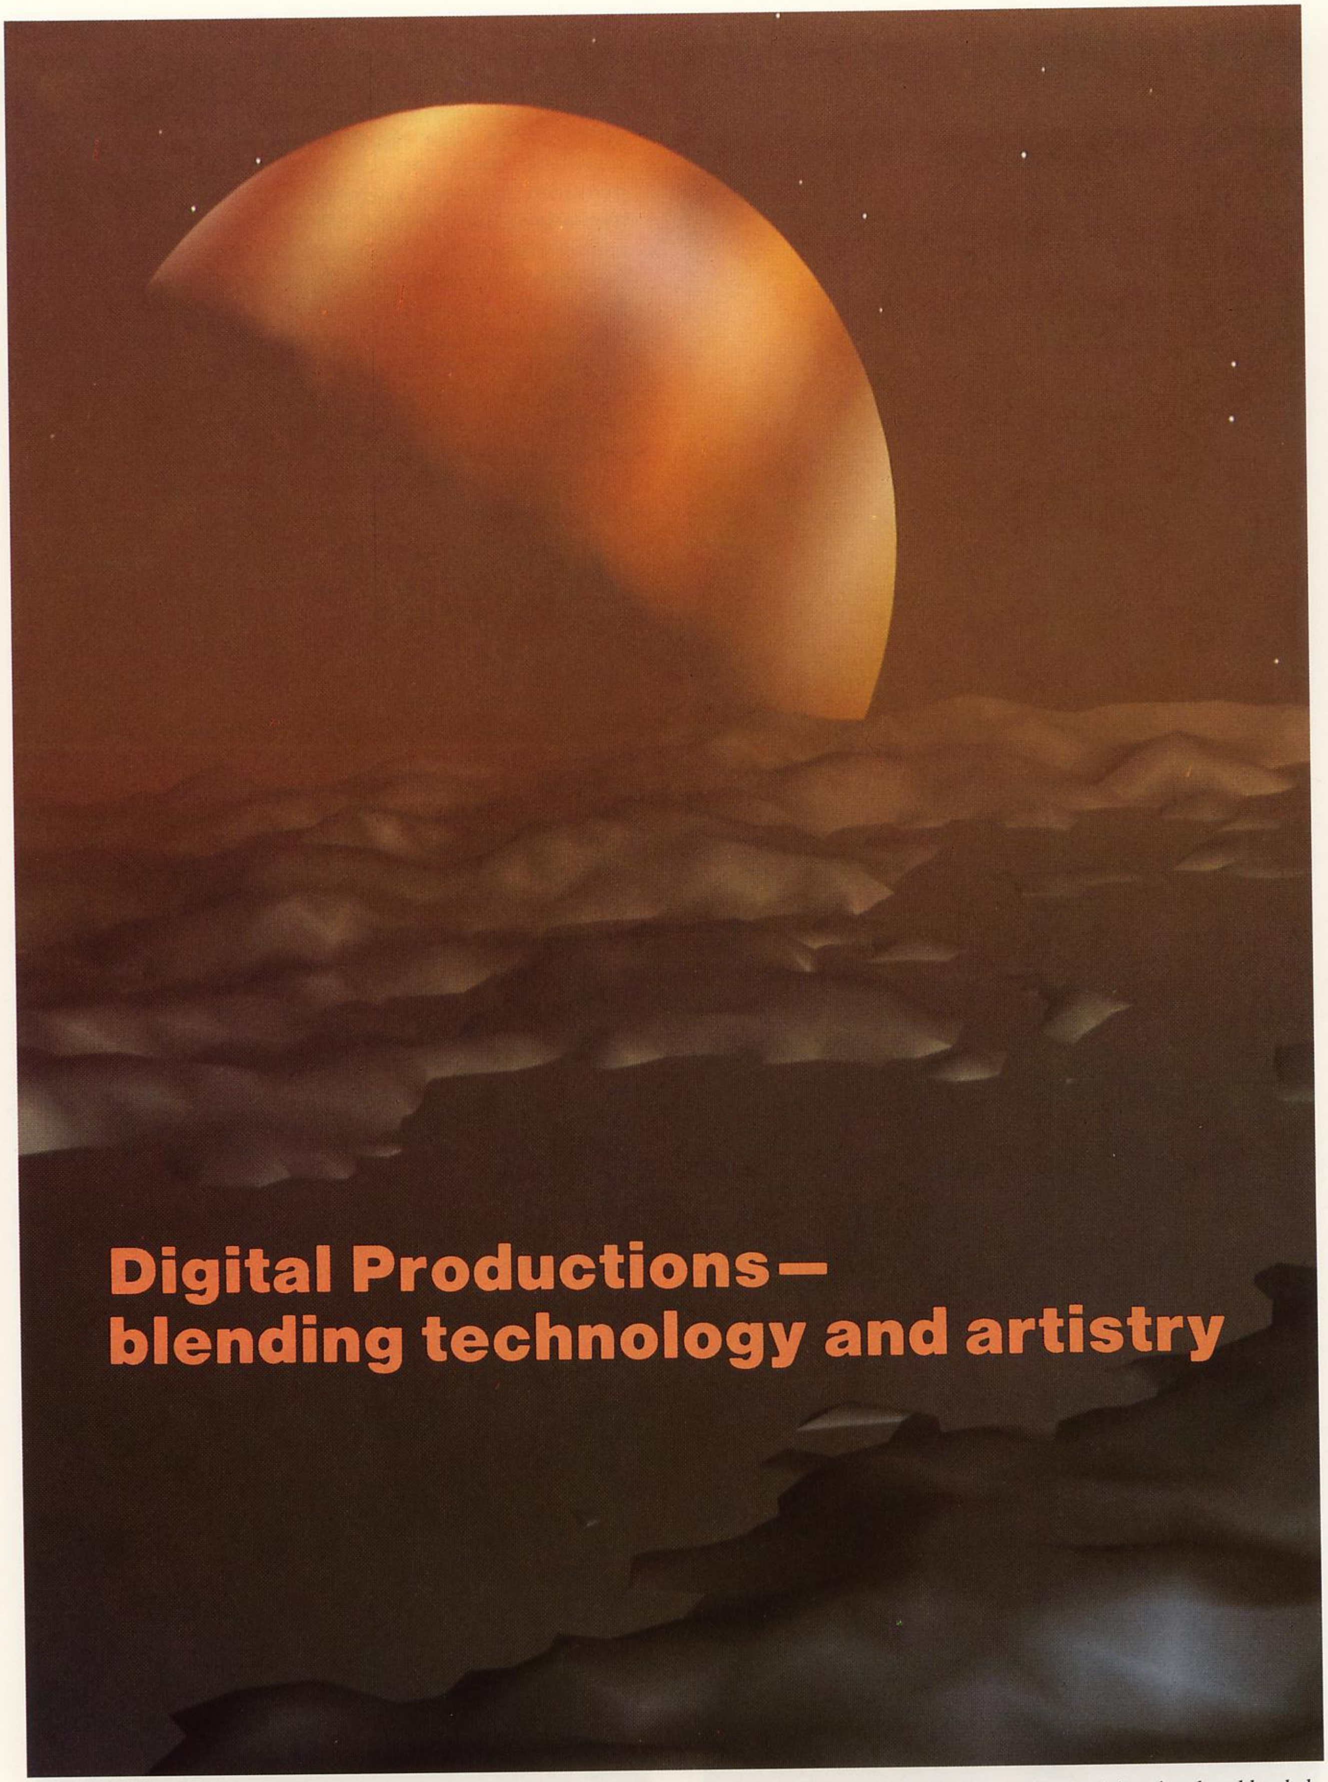

Digital Productions— blending technology and artistry

Digital Productions' CRAY-generated planet terrain. The planet in the background is artistically rendered with soft colors blended together in a rainbow effect to represent clouds. The terrain is modeled using a proprietary fractal package which produces random surface textures. Digital Scene Simulation^{sm} by Digital Productions, Los Angeles, CA © Copyright 1983. All rights reserved.

Digital Scene Simulation^{sm} is a phrase coined several years ago by innovators John Whitney, Jr. and Gary Demos to define their pioneering work in creating complex images from computer-based data. Their efforts have culminated in the formation of Digital Productions, a Hollywood film production studio devoted to the advancement of state-of-the-art computer graphics imaging and animation. The scenes they expect to produce should suspend the film audience's ability to tell the difference between live action intercut with photo-realistic computer simulation.

Digital Productions boasts a unique blend of computer technology and human creativity. It is staffed with bright computer and creative types and is equipped with powerful hardware and software tools. At the heart of the computer network is a CRAY-1/S 1300 system, necessary for this computationally-intensive computer graphics application. VAX 11/780 computers, Ramtek color raster displays and a variety of workstations for encoding and other tasks, round out the necessary tools of the trade.

This CRAY computer is the first to be devoted entirely to computer scene simulation. To meet anticipated expanding work loads, Digital Productions was also the first company to order a CRAY X-MP, which will be installed in late in 1983.

The current system has the capability of calculating complex scenes of up to one million polygons per frame, and can turn out a picture of 4,000 x 6,000 pixels in a reasonable time frame. To illustrate the need for a supercomputer for computer-generated animation, one may assume that it takes ten calculations to determine the value of each pixel. Since commercial motion pictures run at 24 frames per second, it takes 5.76 billion calculations to produce one second of film (24 frames x 10 instructions x 24,000,000 pixels). In actuality, various lighting and rendering requirements cause calculations to range anywhere from 0.1 to 1,000 per pixel. Hence, it takes anywhere from 57.6 million to 576 billion arithmetic instructions to produce one second's worth of motion graphics. The entire image doesn't change in every frame, of course, and computer graphics specialists use some shortcuts to reduce the number of calculations.

John Whitney and Gary Demos have collaborated for many years to define and refine the concept of Digital Scene Simulation^{sm}. Whitney was first introduced to computer generated art by his father, John Whitney, Sr., who is credited with originating the concept. It was during his teen years that John, Jr. began making films with computers. Gary Demos, who graduated from the California Institute of Technology at the age of 20, was intrigued with the principles espoused by John Whitney, Sr. who taught at the institute. In 1973 the two young men joined forces. Since that time, they've worked on the development of three versions of high-powered image-generation software, each more powerful than its predecessor.

Today's software renders images with tremendous clarity. The algorithms being designed accurately model physical attributes such as perspective, highlights, shadows, surface textures, microfaceting, reflections, multiple light sources, volumetric transparency, translucency and refraction. Currently, the company is working on a new action-adventure feature film for Lorimar Productions, entitled STARFIGHTERS. John Whitney says that STARFIGHTERS will achieve a new level of realism in computer-generated special effects for feature films.

An important aspect of Digital Productions' research work actually goes beyond realistic image generation. Demos and Whitney believe that computer power is needed for producing the dynamics of a scene as well as rendering it. Currently, computer graphics in the film industry is limited to producing special effects. While producing realistic images is a major step in computer-generated images, it is a first step, nonetheless.

Demos and Whitney see that the real strength of Digital Scene Simulation^{sm} will be realized when it can carry a story line. The subtleties of motion, expression, gravity, friction and the physical effects of nature, dynamically interacting together, must be handled by intelligent computer programs. These are effects that traditionally have been modeled and analyzed on the CRAY in scientific research, but not intertwined in a dynamically interactive situation. In addition, they foresee that if the subtle but powerful nuances of human behavior and psychology could be modeled and displayed via computer graphics, the the discipline would be elevated to an entirely new level of sophistication. Artists and directors would be freed from the details of execution and could spend the majority of their time and energy on creative work.

Gary Demos explained it this way: "Remember the Disney Productions. His unique contribution was in pioneering animation, taking it from short silly characterizations where figures moved about jerkily into telling full feature-length stories . . . . The ability to express the theatrical was his strength . . . . His work enabled the viewer to see and feel emotions, and this was used in a consummate fashion to manipulate the audience." He went on to say, "If you have the technology to describe such effects, then you can carry the story line. Most of the computer-generated characters to date have seemed cold and lifeless. However, recently we have seen some humor in them."

And when it gets right down to the time involved in producing film, computer-generated animation is really beginning to come into its own. At Digital Productions right now, four minutes of film at a 250,000 polygon complexity level can be generated

Imaginative CRAY tower generated by the CRAY.

per month. "We already have the capacity to produce 45 minutes of animation per year," states Whitney. "When the X-MP is installed, production should increase by 400%. That's not bad when you consider that movie studios like Disney produce on the average of one full-length film every two years."

While Whitney's and Demos' work has primarily been inspired by the needs of the motion picture industry, their software is also relevant to industrial, scientific and military installations. Digital Scene Simulation[sm] software is of interest to organizations that require high-quality, interactive image rendering capabilities. Various applications include seismic research, oil exploration, fluid flow analysis, structural analysis, computer-aided design and architectural design.

The outlook for the future is exciting. "When computer graphics is developed extensively enough," asserts Demos, "we will be able to conjure up images that can express ideas and events better than words." Using what Whitney terms "Interactive Exploration," new areas of simulated visualization will be fully realized, for example, the accurate depiction of the dynamic properties of gases, fluids, and solids in systems. For now, Digital Productions is continuing to use Digital Scene Simulation[sm] to change the way people are looking at motion picture and industrial design. □

Acknowledgement

This article was prepared with the kind assistance of Maxine Brown, Digital Productions.

graphics enables scientists to study the large and small of it

British astronomer and physicist Sir Arthur Eddington once said, "From his central position man can survey the grandest works of Nature ... or the minutest works." Scientists at Los Alamos National Laboratory and Lawrence Livermore Laboratories have embarked on these two paths in researching the Earth's geochemistry and the behavior of atomic nuclei, respectively. In the process, they have taught us more about the world in which we live.

In both instances, the Labs have relied on the processing capabilities of their CRAY-1 supercomputers in combination with sophisticated graphics generation and analysis programs. Both studies found ways to deal effectively with tremendous amounts of data. This article focuses on using graphics to study the large and small of it, from data "sandwiching" at Los Alamos to "filming" the behavior of atomic nucleons colliding at Livermore.

Data integration at Los Alamos National Laboratory

Researchers in the Earth and Space Sciences Division at Los Alamos National Laboratory have developed a rapid, efficient method of resource evaluation to keep up with an onslaught of geologic and geophysical information. The Lab's new data integration system turns reams of data and information from NASA satellites into photographic images that may be analyzed further by scientists.

Thomas Weaver, principal investigator, says "Information is being thrown at us so rapidly that we cannot keep up with the data." The new data integration system, which was developed as part of a U.S. Department of Energy (DOE) uranium resource evaluation program, was given a dry run at Talkeetna, Alaska. The success of this initial test led to development of a full-scale system that uses Los Alamos' five CRAY-1 supercomputers.

The CRAY systems are programmed to accept data from aerial geophysical surveys, geologic maps,

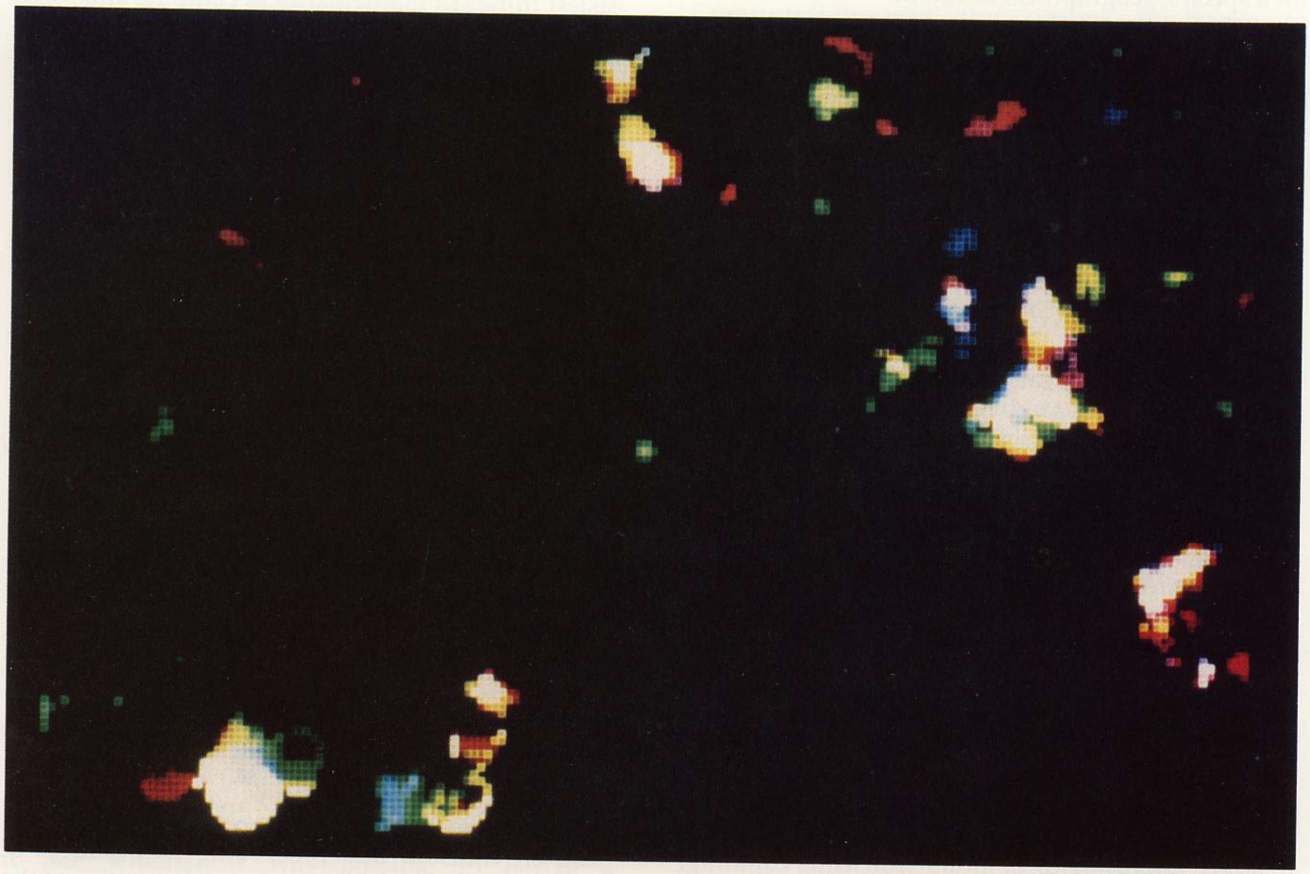

Sample sandwiching output, courtesy Tom Weaver, Los Alamos National Laboratory.

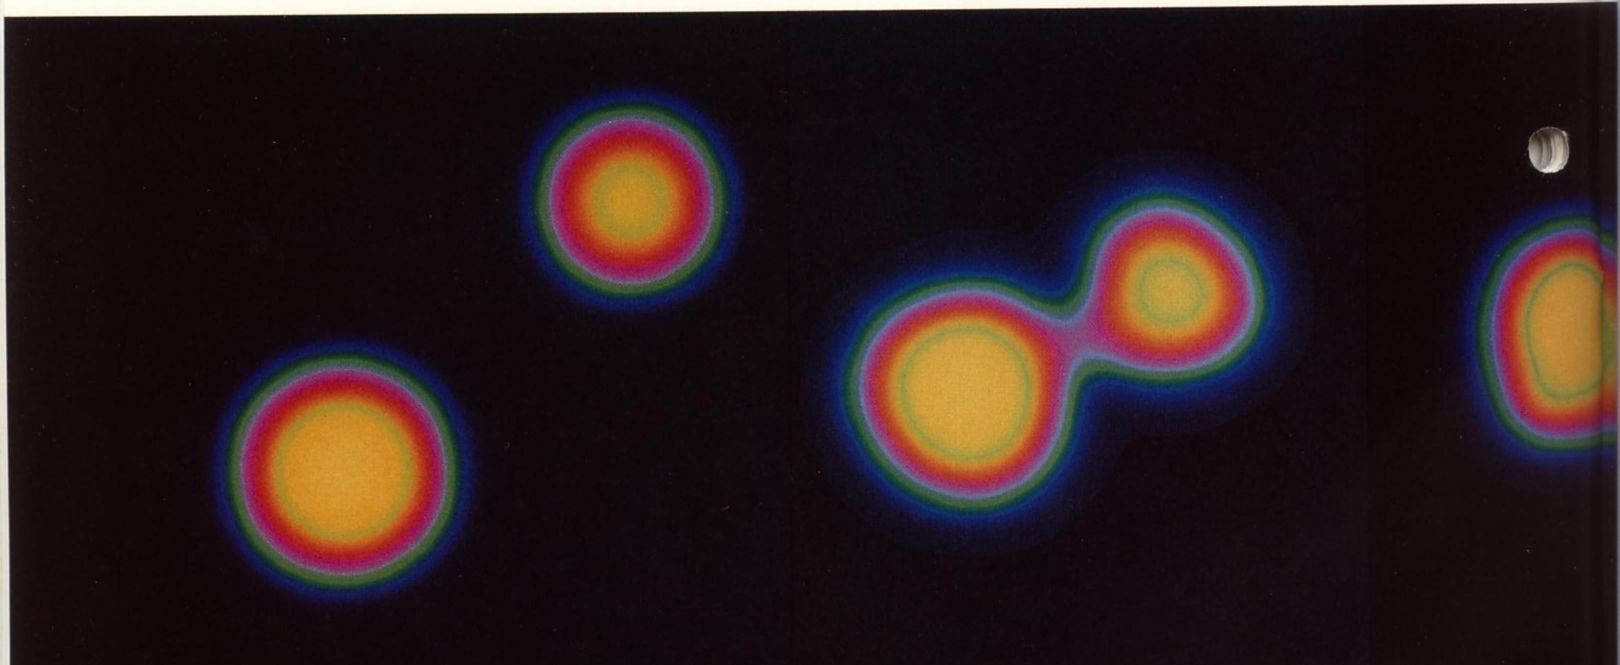

Computer generated images representing krypton-86 and lanthanum-139 colliding nearly head on at the same bombardment energy. The illustrated collision takes 45 x 10^{-22}s, courtesy Dr. Morton S. Weiss, Lawrence Livermore National Laboratory.

geochemical data, and Landsat imagery. The information is digitized, then transformed into photographic images that are overlaid, or "sandwiched" on film. The result is a spatially complete "photograph" of all available information for a specific geographical location.

Following their success at Talkeetna, researchers ran a larger, more complex test on an area in southern Colorado's Montrose Quadrangle, which was also part of the earlier DOE study. The quadrangle was known to contain mineral deposits and several types of mines, and the calibration data needed for the Lab's new data integration system were available. Results of the analysis showed startling evidence of heretofore unknown concentrations of copper, lead, and zinc.

The DOE now has contracted with Los Alamos to produce a geochemical atlas of the entire state of Alaska using the sandwich format. Due to the progress made, much of the guesswork has been taken out of mineral exploration, and scientists say the wildcat methods of searching for mineral deposits may no longer be necessary. Researchers also believe it may be possible to adapt the system to include information on seismicity, gravity, and elevation, enabling the prediction of the probability of floods or earthquakes in various areas.

Quantum mechanics by supercomputer at Lawrence Livermore National Laboratory

At Lawrence Livermore National Laboratory (LLNL), the CRAY-1 has teamed up with the latest in computer graphics techniques to allow scientists to study the behavior of atomic nucleons colliding. By hand, the task is formidable, practically impossible, but the CRAY-1 and the use of graphics provides an opportunity to simulate such reactions and then view them.

The mathematics that describe dynamics in quantum mechanics is the mathematics of wave functions. Wave functions are far from being the easiest things in mathematics to compute and solve. Calculating the behavior of a single atomic nucleus involves dozens or hundreds of wave functions (each proton and neutron has its own) that are related to one another. Thus, they cannot be solved one by one; rather, they must be solved simultaneously.

At Lawrence Livermore National Laboratory, Dr. Morton S. Weiss and co-workers have been able to calculate what happens in various kinds of collisions of a krypton nucleus with a lanthanum nucleus. Then, via computer graphics, Lawrence Livermore National Laboratory scientist Nelson Max has displayed these processes in a film. The nuclear dynamics are displayed by assigning a color to each 10% range of density.

The pictures and the information in Dr. Weiss' study represent the solution of 146 simultaneous wave functions. There is one function for every two nucleons rather than one for each because the mathematics contains one important approximation: Nucleons can come with two directions of spin — spin up or spin down. This theory does not distinguish between the two spin states, so each function actually represents the average of one of each state.

The film shows a dozen instances, starting from a head-on collision and proceeding to instances of greater and greater offset between the centers of the two nuclei until they arrive at a sideswipe in which

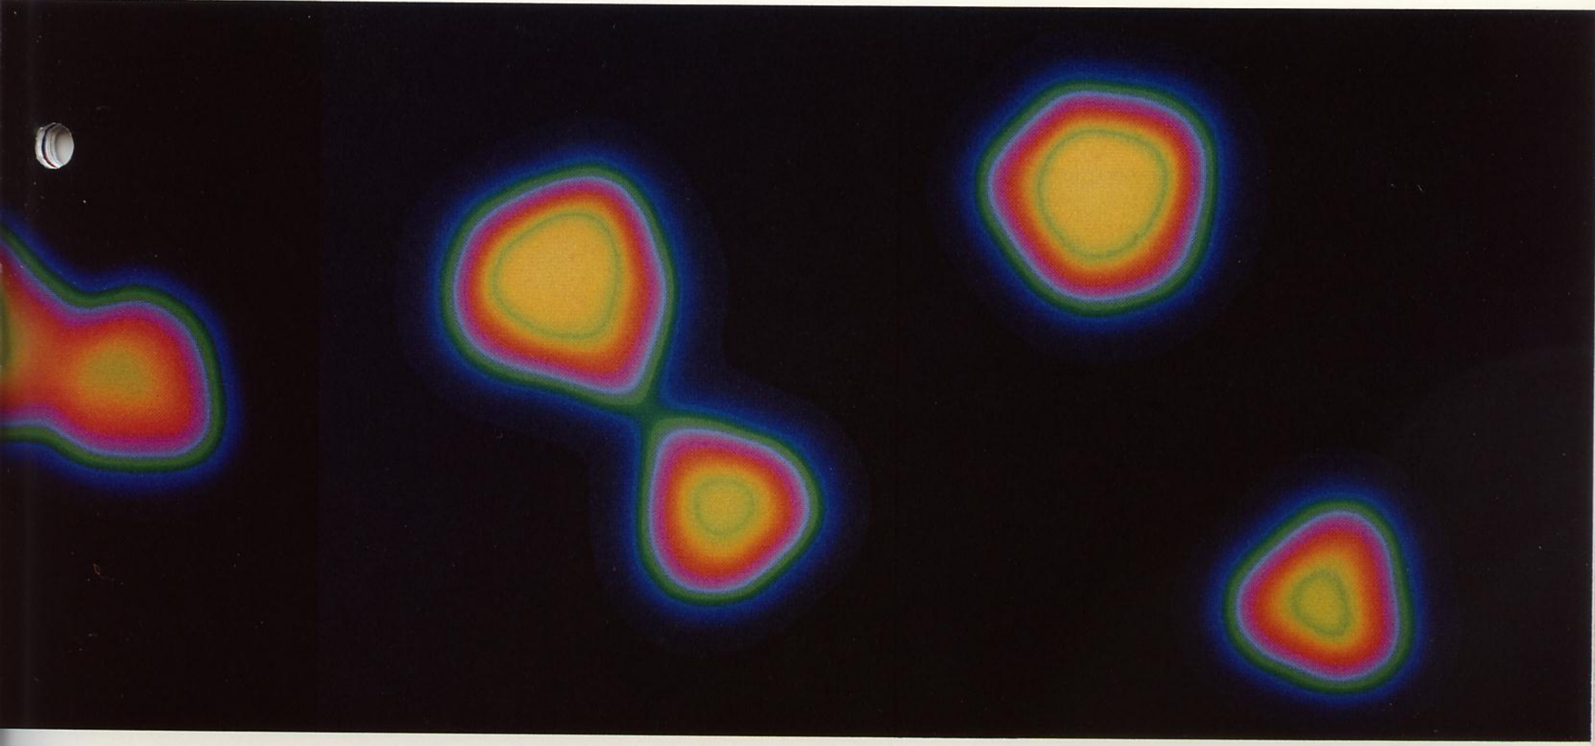

they hardly touch each other. The near head-on collisions result in fusion. A little offset contributes an angular momentum and the fusion product rotates. With more offset, the nuclei collide and lose a lot of energy, but they retain enough to come apart again (known as deep inelastic scattering). Finally, in one instance the two nuclei do not touch, but they become distorted and begin vibrating due to the force produced between the nuclei when they are close.

Dr. Weiss said the film wakes up physicists when he shows it at the end of a talk. He mentions one experimentalist who came up and said how nice it was to be able to actually see the phenomenon. "At least four levels of abstraction," Dr. Weiss remarks, "and he talks about seeing it." Yet viewing the film, one can understand what the experimentalist meant.

The calculations that generated the film images are quite accurate in predicting the main events of the collisions. According to Dr. Weiss, the film accurately predicts the conditions under which fusion is likely and the proper range of deep inelastic scattering. "Nelson Max is a genius," Dr. Weiss comments. "I am glad appropriate credit is being given to those that augment research findings like this in such an important way."

In the pictures the behavior of the nucleons looks like that of a fluid — liquid drops coming together. Dr. Weiss stresses that the effect is actually the sum of discrete motions, those 146 individual wave functions. One of the purposes of the exercise was to see whether it was possible to derive a macroscopic theory that would give a simpler mathematical expression for the overall behavior of the nuclei while yielding the desired predictions. It turns out, Dr. Weiss says, that a macroscopic theory is possible in some particular cases, but the microscopic one, starting from the 146 individual wave functions, remains the general one.

Computation was the serious hindrance to completing this study before now. The theory of quantum mechanics itself was written down years ago by P.A.M. Dirac, a pioneer theorist in the field.

It took Dr. Weiss and co-workers a year to figure out how to feed the computation to the computer. Each of the eleven collisions, which takes about 10^{-21} seconds in real time, took 90 minutes of computation time using the CRAY-1.

The code, which had been developed over a period of about five years, was redesigned and optimized for the CRAY-1. Dr. Weiss says that it simply would not have been possible to execute on a smaller computer. "Because of the CRAY word length, a lot of storage was packed into half-words," he explains. "In order to evolve a collision in real space, a volume of 35 fm x 35 fm x 16 fm (fm: 10^{-13} centimeters) was required. The CRAY is the only system capable of computing the calculations for that volume." □

Acknowledgements

"Data integration at Los Alamos National Laboratory" is based on an article appearing in Los Alamos Science, Vol. 3 No. 3, Fall 1982. *"Quantum mechanics by supercomputer"* is reprinted with permission from SCIENCE NEWS, copyright 1982 by Science Service, Inc. It appears here in modified form.

An interview with Nelson Max

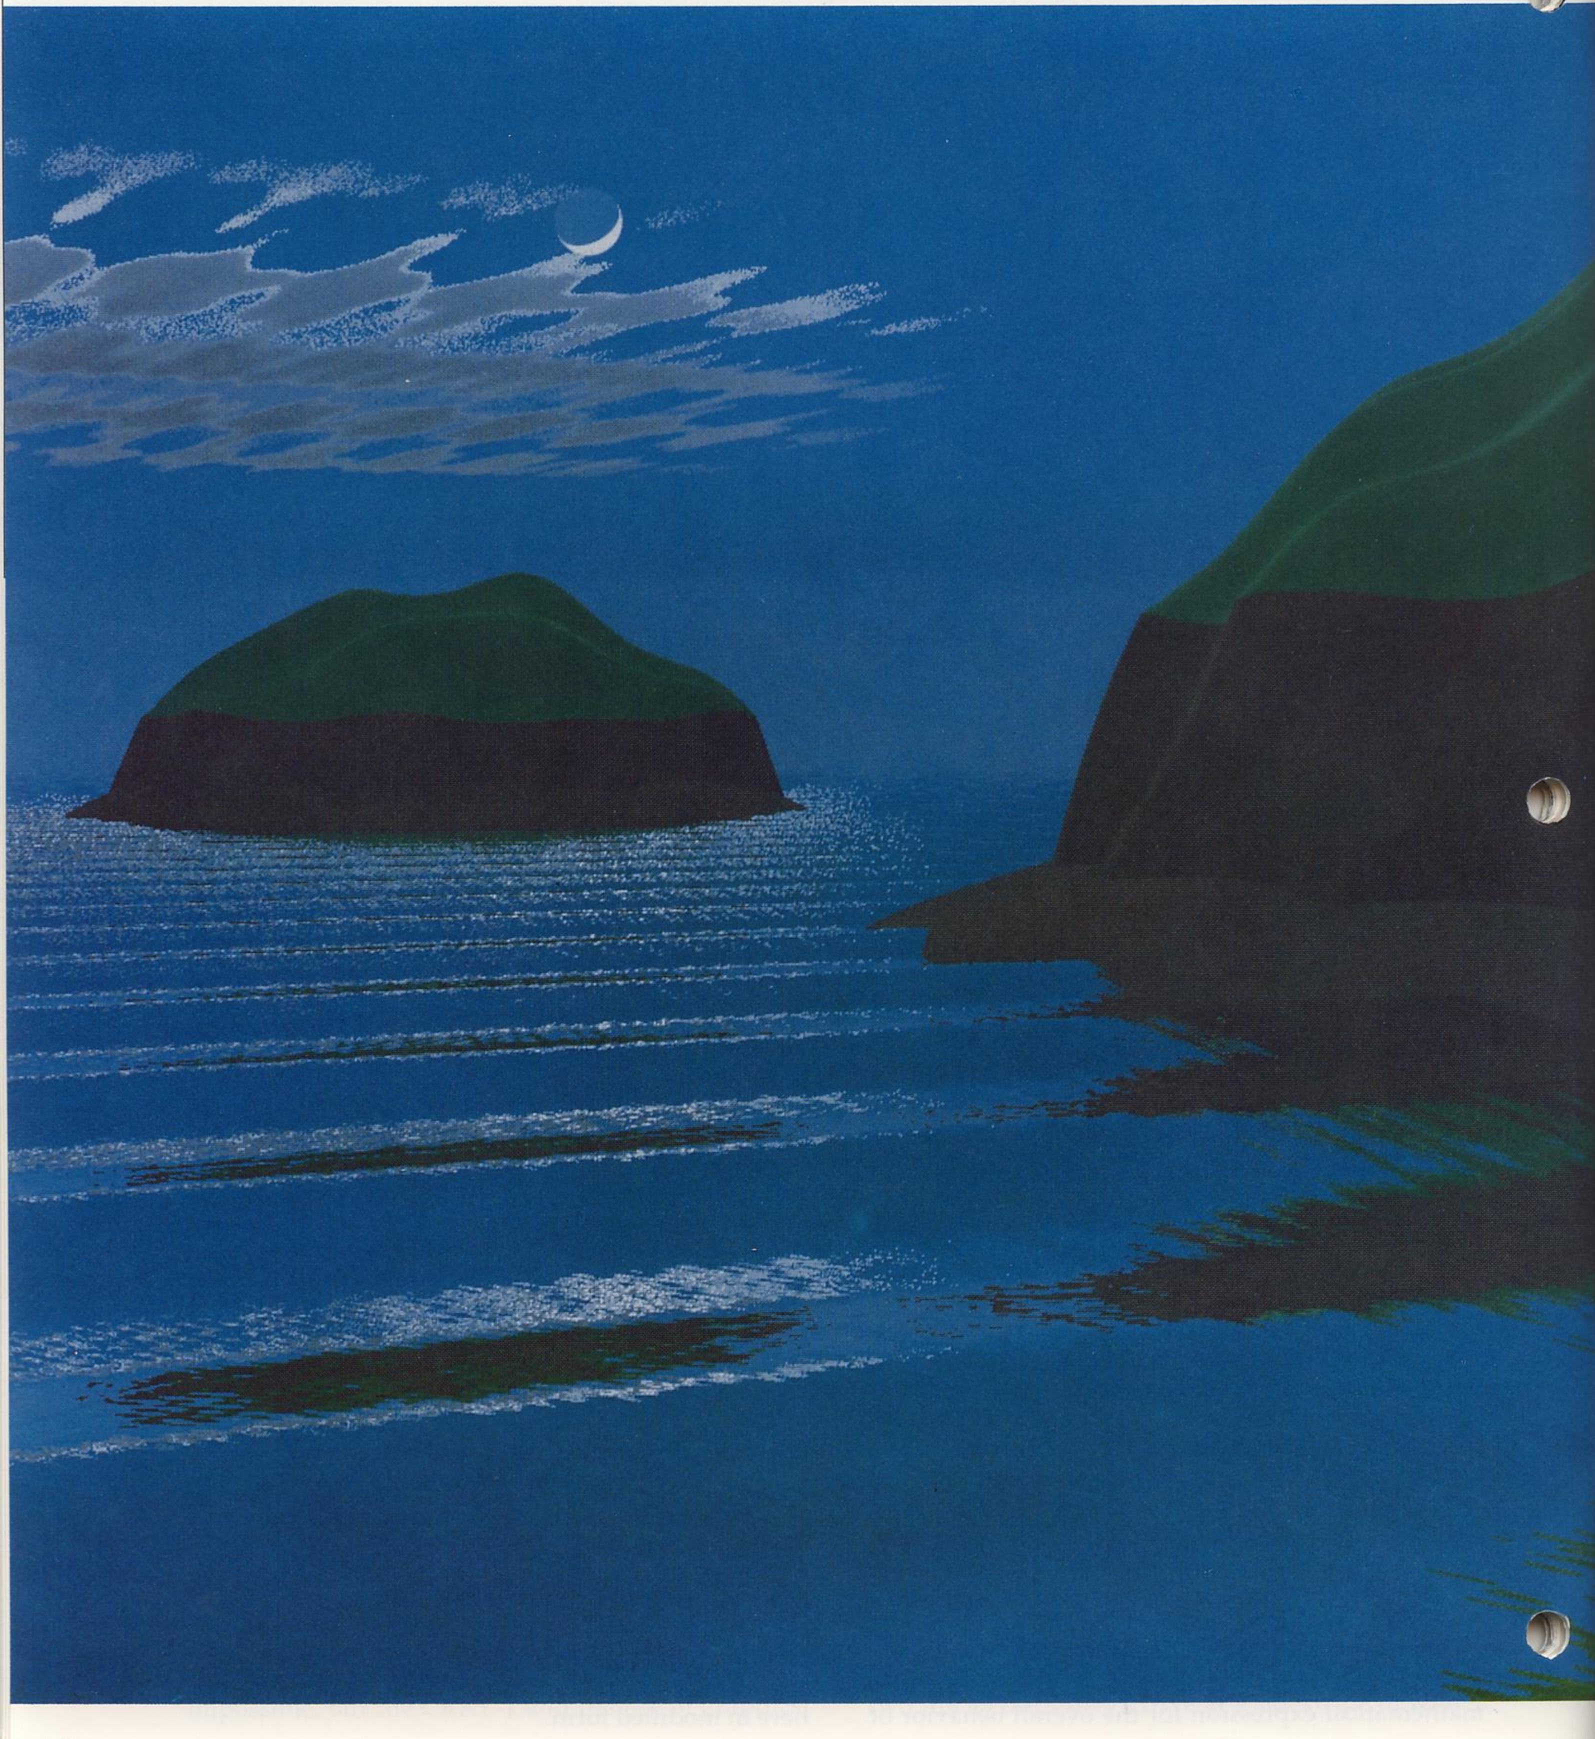

Nelson Max is among the major proponents of computer graphics development today. Back in the late 1960's, shortly after receiving his Ph.D. in mathematics from Harvard, he first developed an interest in illustrating topological phenomenona with computer animated graphics. In the early 1970's, Max produced his first computer-generated films, "Space Filling Curves" and "Turning a Sphere Inside Out". These early works provided Max with the impetus to pursue graphics as more than an avocation. In 1977 he left Case-Western Reserve University, where he had been involved in mathematics research, for Lawrence Livermore National Laboratory (LLNL) to conduct computer graphics research. Since then, he has experimented with, and implemented many innovative computer graphics techniques. At LLNL, Max is heavily involved in molecular and solids modeling. However, one of his more popular films has nothing to do with such things. "Carla's Island", presented at SIGGRAPH '81, is very successful in its realism. The program that produced the four-and-a-half minute film ran on one of LLNL's CRAY systems, and was tailored to take advantage of the computer's vector capabilities. We recently spoke with Nelson Max about graphics techniques he used to synthesize "Carla's Island". The following is an excerpt of that conversation.

Let's talk about some of the techniques you used in creating "Carla's Island".

"Carla's Island" is representative of work that I've had the most fun doing. Since its development in 1981 it's been the basis for a lot of experimentation.

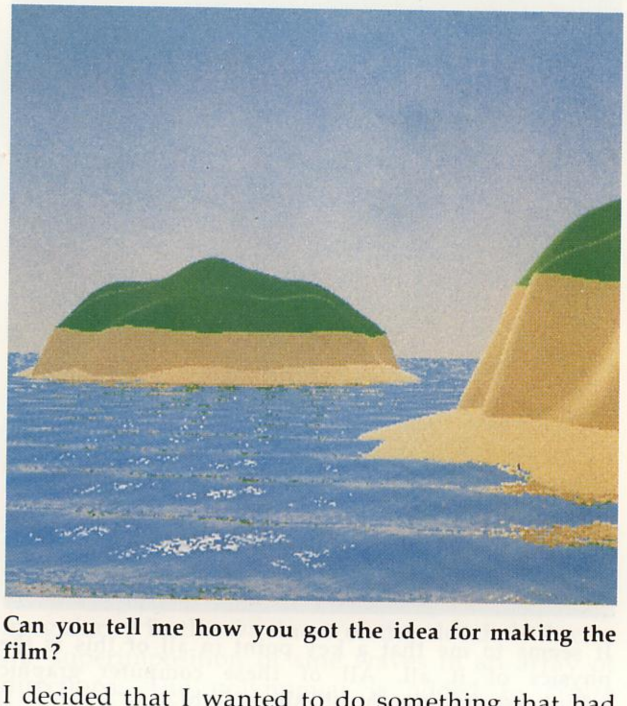

Can you tell me how you got the idea for making the film?

I decided that I wanted to do something that had never been done before at SIGGRAPH '81. In 1980 Lauren Carpenter had done a film with fantastic realism, more effective terrain, mountains and the like, than had ever been done before. One evening after one of the SIGGRAPH '81 sessions, I was on a sunset boat cruise crossing Puget Sound just watching the water. It was then that I thought of creating a film with water reflections. I decided that the water reflections would have to reflect off of something, which is when I decided to include islands in the picture. The terrain in "Carla's Island" is just awful. The island looks like a cupcake. But I just wanted something to reflect in the water. Some of my colleagues at Livermore are currently working to generate more realistic terrain. Later that fall, I began to scheme about how I could create clouds on the computer to make the scene even more realistic.

Where did you get the name "Carla's Island"? Is there a story behind it?

A good friend of mine named Carla went to Hawaii and fell in love with the islands. She's quite artistic and was inspired to paint her fantasy island.

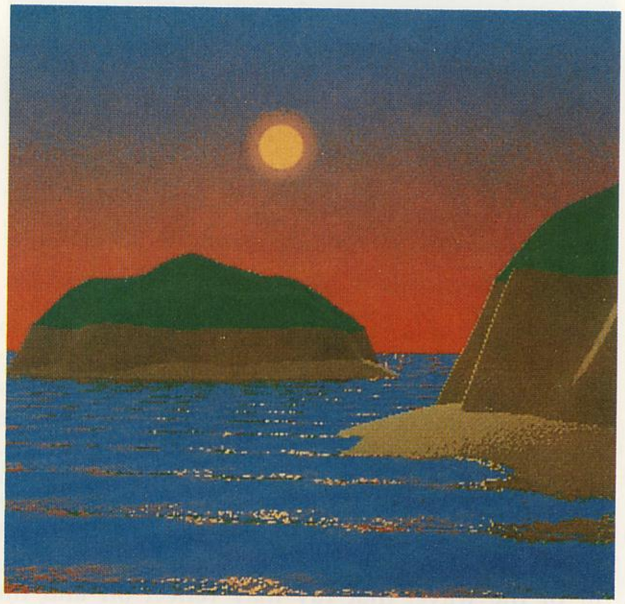

Is the island in the movie the same as the one she painted?

It's close. She wanted a sailboat, palm trees, various things to make it her island. I tried to make it look like a stylized image of a south sea island. The islands obviously need palm trees, so I've sort of been dreaming. When I was in Hawaii recently, I studied the palm trees. They're complex structures that would require complex algorithms to make them seem real. It would be very computationally expensive. Of course it would fit on a CRAY — just keep other users out.

Before we get ahead of ourselves, why don't you explain what the film is all about?

"Carla's Island" attempts to illustrate the changes in skylight and water reflections as the day progresses from morning into the night. Morning light changes to bright midday light and then moves through sunset and moon lighting changes. There's even a thunderstorm and the clouds change shape and grow.

Can you go into detail on how you created reflections?

I used a ray-tracing scheme that takes into account the global illumination information of an image. I'm not the only one who has used this method, it just worked well for "Carla's Island".

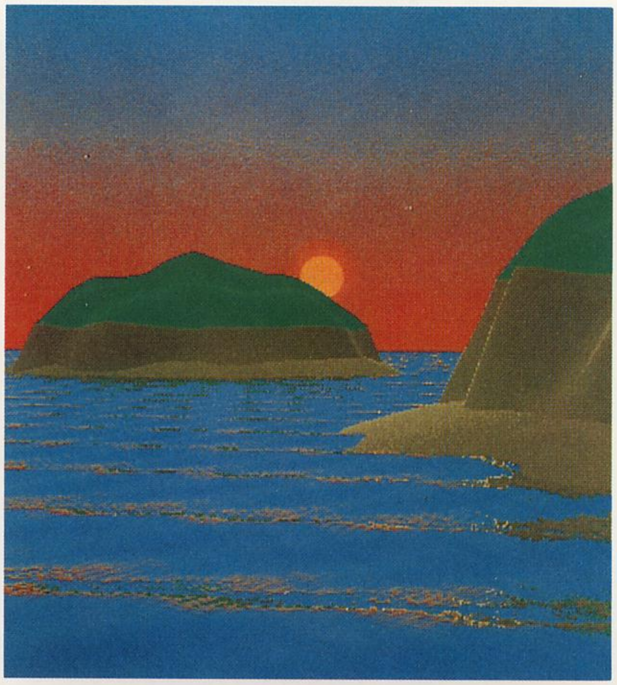

How does it work?

The ray-tracing scheme combines a hidden surface algorithm and shading model. First, I identified the physical structures and the position of the viewer's eye. Then I traced the rays extending from the eye through each pixel, to the first surface encountered. For each ray, and thus for each pixel, reflection directions were calculated as a function of other surfaces and light sources in the image. Let's say a ray hit an object like an ocean wave. Based on the tilt of the wave, that ray may reflect off and hit another wave. From there it may hit the island or reflect back to the sun. To a limited extent I was able to split a ray at a given surface to refract and reflect and then trace both segments. Turner Whitted, another graphics researcher, has done this extensively. The real advantage to ray-tracing algorithms is that all the illumination information is accounted for. The physics of light can be simulated more accurately.

Was color added during this operation?

No, I used a color look-up table for coloring each frame. Basically, different colors appearing in "Carla's Island" were divided into regions in the table. There's a beach region, a cloud region, sky, ocean, sun and so forth. As the sun, the primary light source, moves, colors of the different objects and their reflections change. The color table actually contains normal reflection information in an eight-bit per pixel format. Over 16 million colors (24 bits per pixel) were available, but because of the constraints of the look-up table, only 256 colors could be in any single frame. I couldn't antialias the islands because there just wasn't enough room in the table for all the colors to antialias the sun. But I did save a range of colors between the sun and sky colors.

Sounds like clever use of the color look-up table. How did the setting sun affect the reflections on the waves?

The sky was represented by horizontal bands of colors which turned orange and red as the sun set. A ray reflecting from a wave would be given an index in the sky color region of the look-up table according to the horizontal band that it would eventually hit, which was determined by the ray's angle from the horizontal. However, if a ray reflected into the track through which the sun moved, the corresponding index would be in a new sun reflection region of the look-up table. When the sun was in a given horizontal band, the corresponding entry in the sun reflection region of the table would be sun-colored and the other entries would be the same sky color as the rest of their horizontal bands. Then when the sun moved to the next lower band, the sun reflection entry for the lower band would fade up to sun color and the entry for the higher band would fade back to its sky color. Thus, I was able to reuse a single section of the wave motion as computed by the CRAY, and still give the reflection of the sunset by continuously changing the color look-up table as the pictures were being plotted. The last two minutes of the film took only 20 minutes of CRAY time to compute, but took two days to plot.

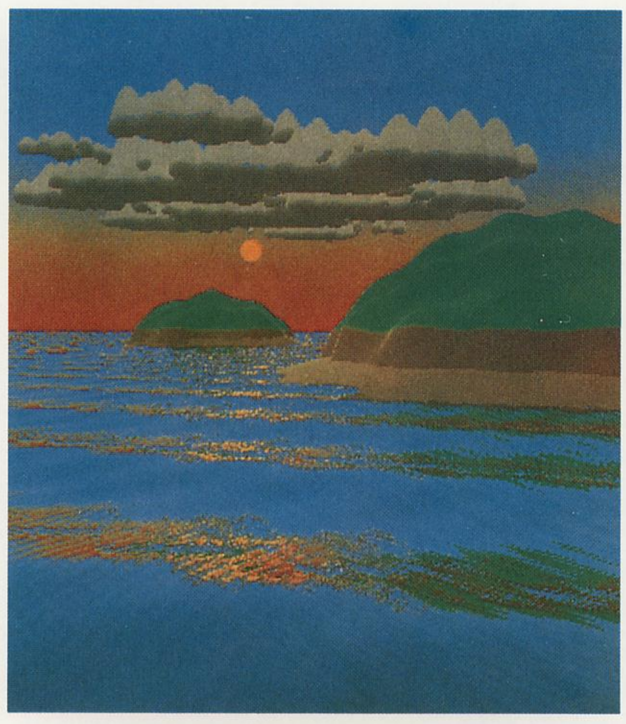

It seems to me that a key point in all of this is the physics of it all. All of these computer graphic techniques replicate the physics of light.

Right. There are other people doing this kind of work. I've learned quite a lot from others about the physics of light diffusion and absorption through clouds.

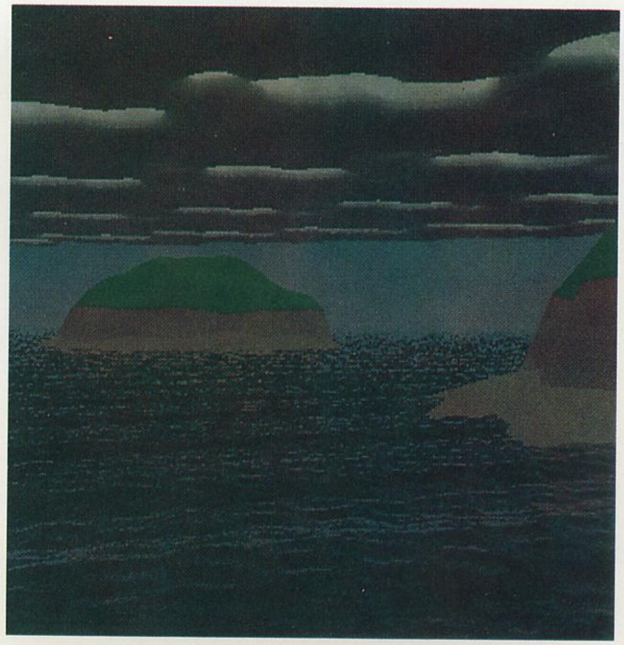

Why don't you talk about that a little bit?

I have some schemes that I haven't implemented yet. The mathematics are written already. I want to make semitransparent clouds that will allow one to see the partial occlusion of the surface behind. I've done a fairly straightforward one, but more elaborate ones in planning will have the sunlight shining from above. The brightness at the bottom of the cloud will be attenuated according to the correct physics. Once I develop the correct calculation for determining how much sunlight should come through each cloud region, I should be able to extend that light below the clouds and create columns of light in the haze. I should also be able to cast shadows onto the terrain below.

It seems to me that in the present form of "Carla's Island", the entire scene darkens a little bit when the sun goes behind the feathered clouds. It's quite realistic. Is that really so? Or am I just seeing things?

Yes, it's true. That's a part of the film I've experimented with. Using the ray-tracing algorithm, I trace rays reflecting from the ocean to clouds. If a ray reflecting toward the track of the sun hits a cloud instead of the sky, the point on the ocean wave will be cloud-colored and never be sun-colored — even when the sun reaches the corresponding band of the sky. So, in fact, you will see the shadow of the cloud in the highlights where the sun reflects on the water. On the other hand, near the end of the film, at sunset, the clouds glow from the red light of the setting sun. In that case, the rays emanating from the clouds should be brighter than the reflection of the sky. Consequently, the reflection of the clouds will be very bright in the water.

When I was in Hawaii recently, flying from one island to another, I realized that the shading on water and land due to the cloud cover would have to be added. So my plan now is to create true shadows. In addition, I'm going to account for rays that are coming up from beneath the surface of the water and refracting to the viewer's eye. What you're going to see is a combination of water color, algae color and turbidity. It's already built into the algorithm but it's not fine-tuned. The thing that I can't do with the color look-up table is correctly add shadows during sunset, because shadows continually move and elongate according to the angle of the sun as it sets even if the clouds are still.

One other thing I'd like to ask about the clouds is how you made them grow and change shape.

They're all defined mathematically. Each cloud's density is defined as a mathematical function, the sum of a polynomial with various sine waves. The irregularity in the waviness of the clouds is due to the superimposition of sine waves in an array of directions. The polynomial is added to create positive and negative regions. Wherever the value is positive, the cloud is visible. Where it is most positive, it's most dense. By changing the coefficients, the clouds grow.

Is there any randomness to it? Or did you physically define the polynomials to suit your representation?

I defined the polynomial. If the resulting cloud didn't look right, I changed the coefficients. If it still didn't look right, I would actually change the formula. For instance, I added extra square roots to round out the edges of the clouds. It was thus, an iterative design technique based on changing the mathematical formulas in addition to changing the coefficients.

Why do you think graphics research is so important?

Computer graphics is utilitarian. Researchers are doing huge computations. How can they be expected to interpret the data? It is important for them to **see** the results. Not only that, if they are able to view motion pictures of the computations, their resultant interpretation might be one that would not have emerged from a sequence of still frames. So film output is very important too. People also like to interactively see how their code is progressing so that if a mesh is crumpling against itself they can stop wasting computer time! Seriously, we have many utility programs that monitor the ongoing computations. Of course slick pictures for presentation purposes are also a popular use of computer graphics. But essentially, many of the research groups really want to generate very realistic views of their computations. By applying some of the principles that come out of research like mine, overall research results are better. □

Further Readings

Max, N., *"Vectorized Procedural Models for Natural Terrain: Waves and Islands in the Sunset"*, Computer Graphics Vol. 15 No. 3, August 1981, pp. 317-324.

Computer graphics at Los Alamos National Laboratory

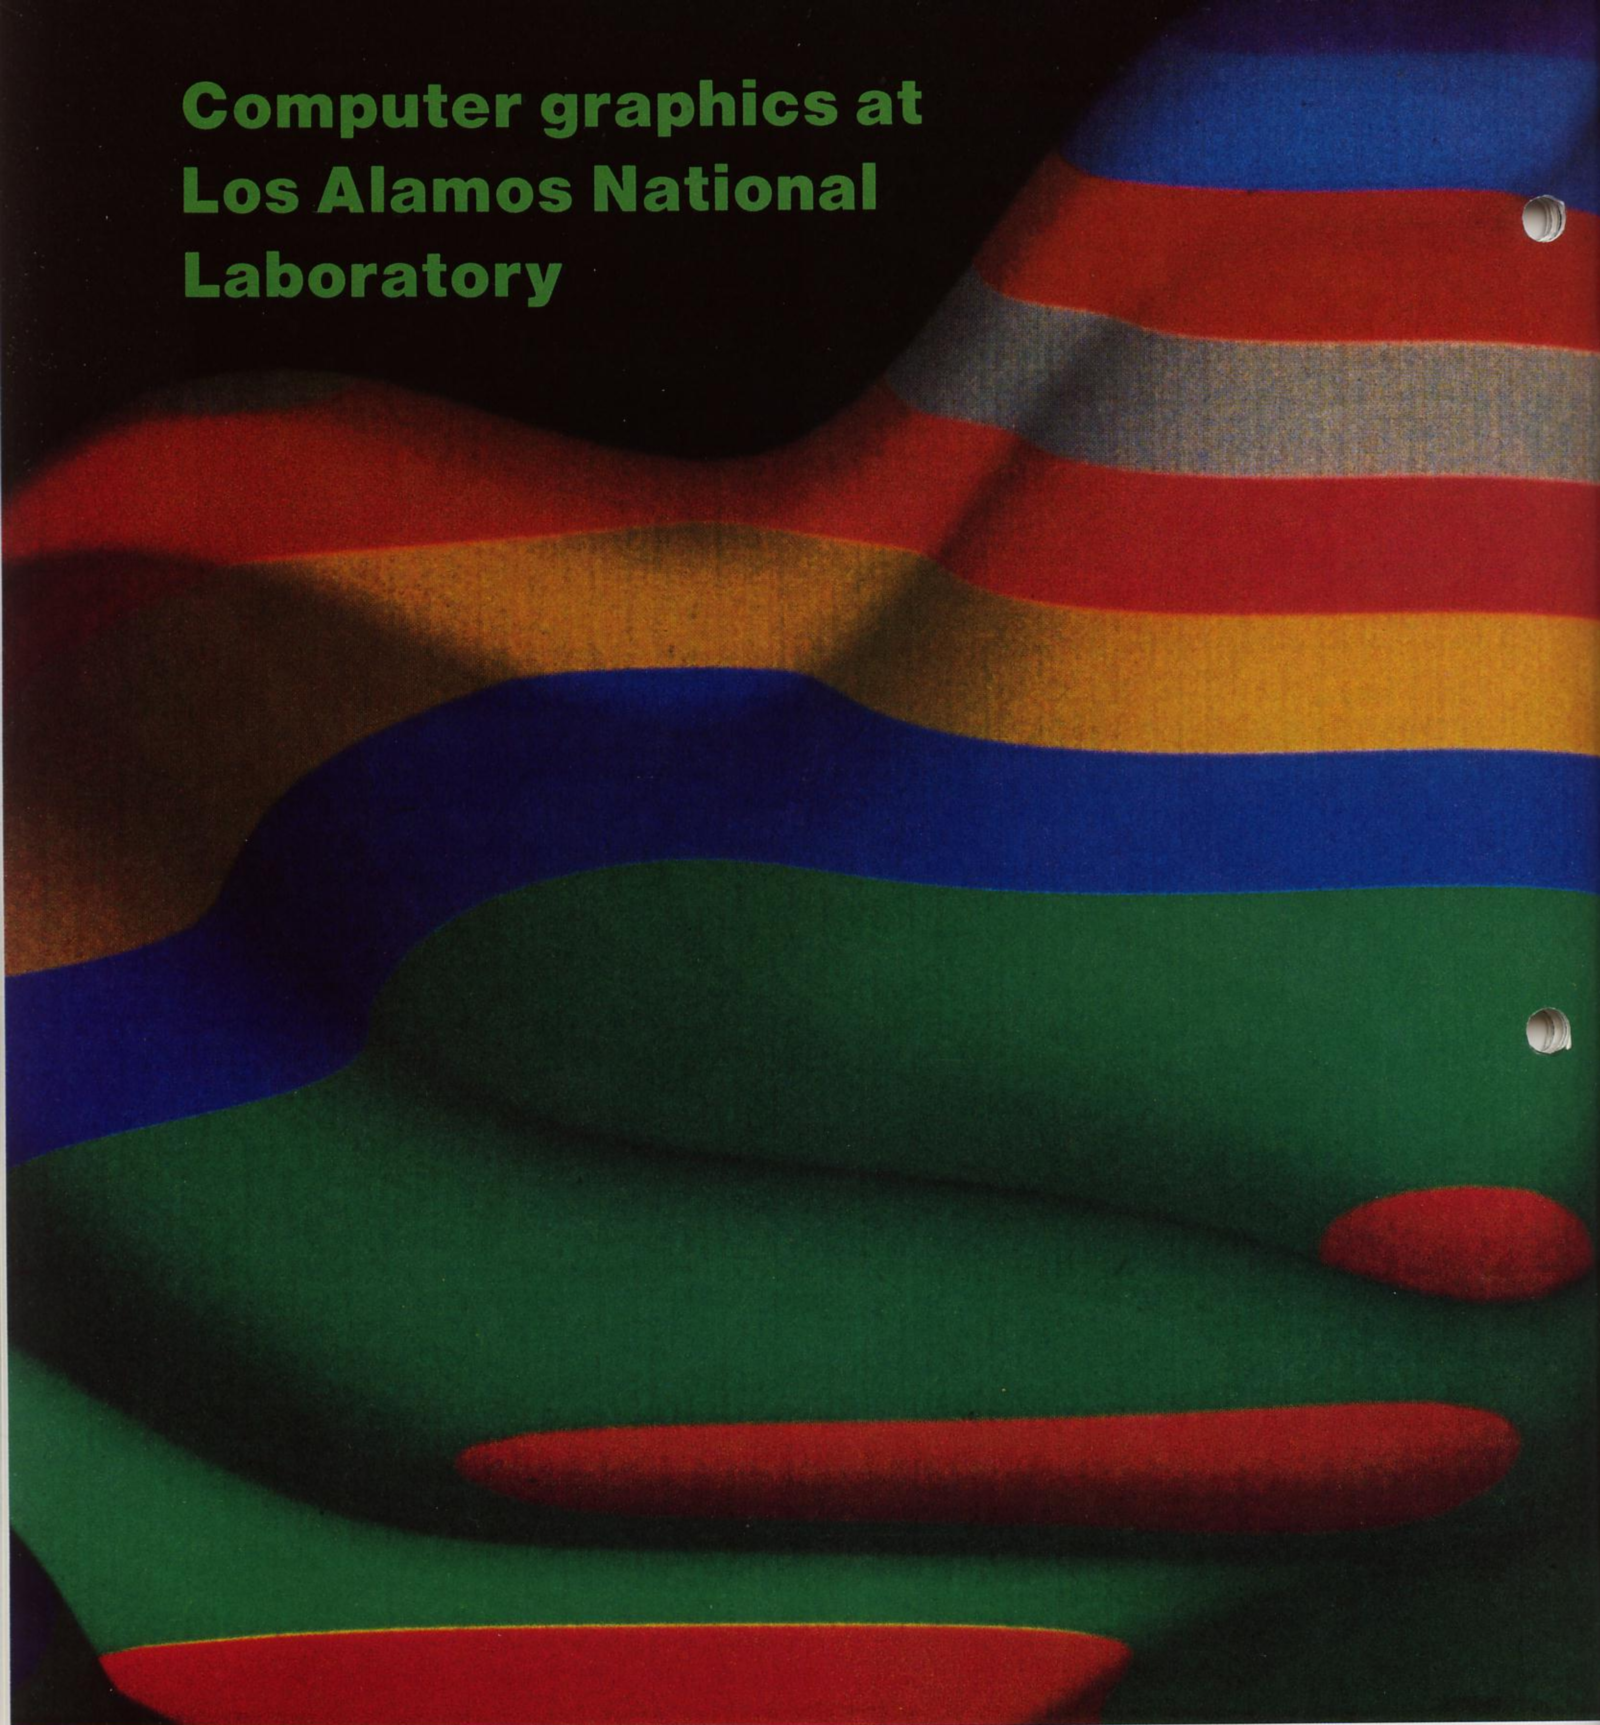

John E. Aldag, Cray Research, Inc.

The computing center at Los Alamos National Laboratory is among the largest and most technically advanced in the world. Indeed, Los Alamos incorporates five CRAY-1/S systems in its computer network. With the complex ongoing research conducted at Los Alamos, it makes sense that dependency on computer graphics would be high — and it is. Therefore, it has been important that computer graphics at the Lab be as functional and streamlined as possible.

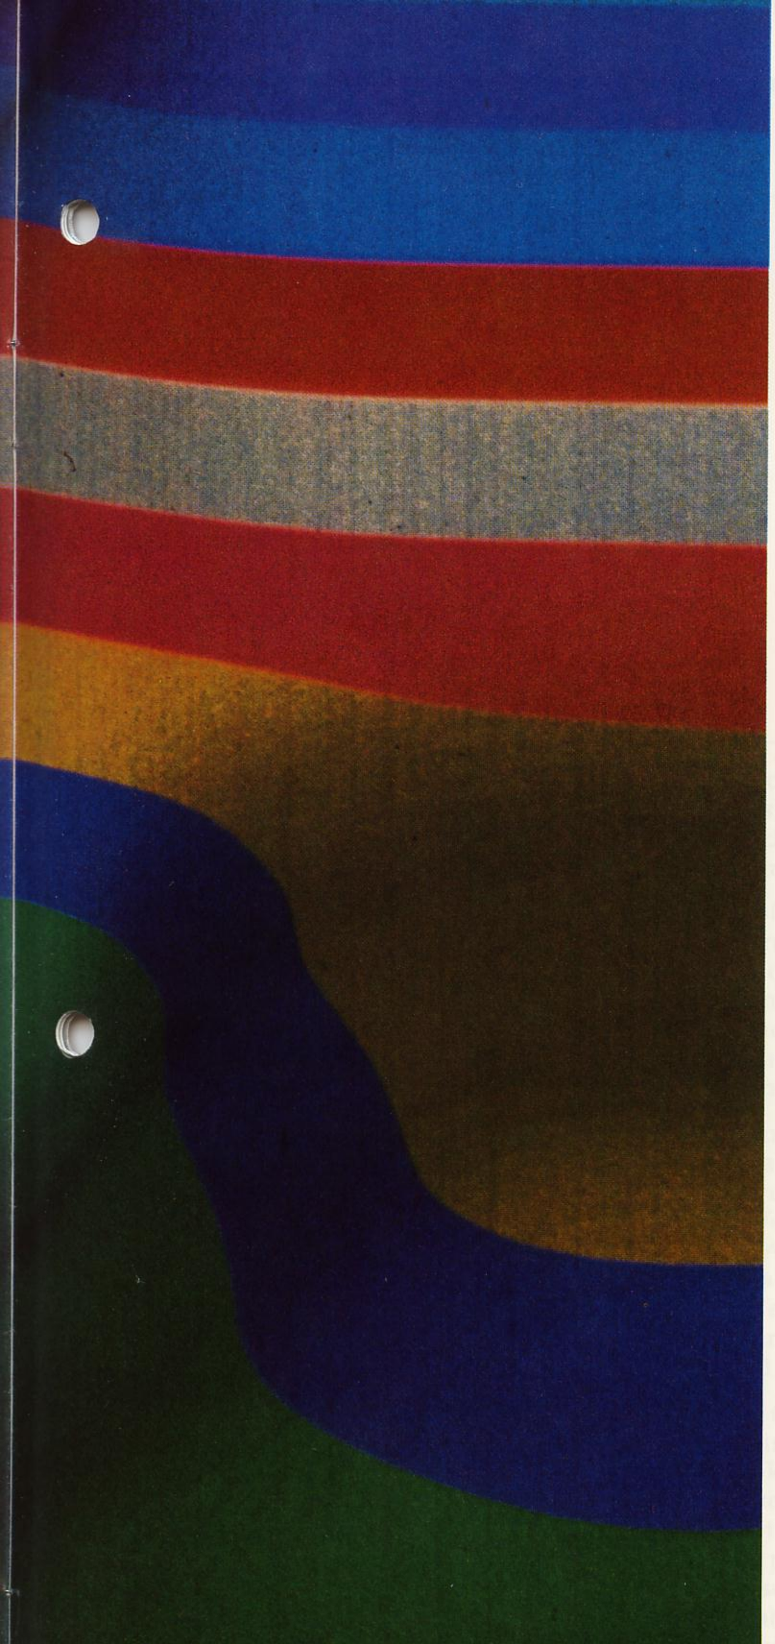

CRAY-generated image created by Mel Prueitt and used in the 1982 SIGGRAPH conference poster.

The Computer Graphics Group at Los Alamos is responsible for integrating a graphics system with the five CRAY computers, four CDC 7600s, three CDC 6600s, a couple dozen or so VAX 11/780s, hundreds of different graphic devices and about 3,500 users. The graphics output needed by the users ranges from simple line drawings and bar charts to high resolution color images of solids with shading and hidden surface removal. The applications vary from presentation graphics and plots of weekly computer activity to sophisticated device designs. The broad spectrum of users is made up of secretaries, clerks, Ph.D. physicists and engineers. Their experience with computers and computer graphics may be a little or a lifetime. Some need interactive graphics on a terminal, some take batch mode output to high resolution film recorders or generate motion pictures. All require fast turnaround.

Sounds like a big job, doesn't it? It is, but it's been made as simple and efficient as possible. In 1976, Los Alamos began implementation of a graphics system that today is a model computer and device independent system. The following article describes the hardware and software capabilities of that system.

Graphics hardware

A wide range of graphics hardware is available to Los Alamos users. At remote locations, direct view storage tubes, such as Tektronix 40XX Series terminals provide the bulk of the interactive local graphics. There are a few color graphics terminals in use at the laboratory, but because of the heavy reliance on computer-generated movies, film recorders are the primary color graphics devices.

The four film recorders and an electrostatic plotter all operate under the control of the Print and Graphics Express Station computer (a VAX 11/780) which receives graphics input from the various worker computers and renders it to the hardcopy devices. The film recorders are in almost constant use. Most of the graphics are rendered on microfiche by the recorders. By changing the recording optics and film mounts, the recorders can be converted for high resolution rendering to 16mm film for graphics movies or to 35mm film for presentation-quality slides. Besides graphics, the recorders are used to produce microfilm listings from the Lab's computers. Users receive their listings on microfiche and view them on their office microfiche viewers. Most of the remaining line printers in the Los Alamos computer center are idle, however, some hard copy is also produced on high-speed page printers by PRINT. An electrostatic plotter, a few color pen plotters, and some VT100 character terminals, enhanced for use as graphics displays, complete the hardware complement at the laboratory.

Computer and device-independent graphics principles

The concept of computer and device-independent graphics is simple enough to describe and its advantages are obvious. A graphics system common to all the available computers allows for efficient, cost-effective use of computer resources.

Scientists and analysts can move their applications programs from computer to computer without changing the graphics calls within the program, thus enabling even load distribution throughout a multi-computer network. With a graphics system common to available graphics devices, scientists can preview computer-generated images on lower resolution or interactive graphics devices before sending them to relatively expensive high-quality color recording systems. In addition, if graphics output is sent to an intermediate disk file, commonly called a metafile, this previewing and final display can be done without rerunning the program that created the graphics. Cost effectiveness and flexibility are what it's all about.

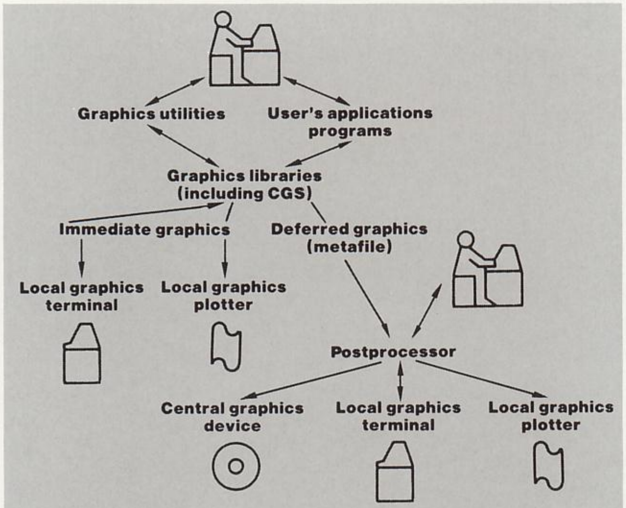

Figure 1 Los Alamos graphics system.

The heart of the system

While the concept of computer and device independence is simple, implementing such a system is quite another matter. The high dependency on graphics and the diversity of user needs made development of a flexible, interactive computer graphics system at Los Alamos a pivotal project. The approach taken by the Los Alamos group in setting up their system is illustrated in Figure 1.

The heart of the system is the Common Graphics System (CGS). CGS was written in conformance with the proposed SIGGRAPH/CORE proposed graphics standard to increase the efficiency of the overall system. It can be used directly or can provide graphics primitives for the more sophisticated programs. As its name implies, it is the one component used by most of the graphics applications at Los Alamos.

Higher level requirements addressed

Users with specific application graphics requirements can take advantage of a broad spectrum of these sophisticated graphic tools at the subroutine call level. This allows them to concentrate on the programming for their research problems rather than on the low level graphics software required to display the results. In order to provide higher level graphics tools to Los Alamos users, the graphics group provides mid-level graphics libraries such as the NCAR graphics library and DISSPLA. These systems provide common but more graphically complex functions and capabilities such as linear, logarithmic, polar axis, projections of 3-D structures, hidden line removal, graphic-art fonts, contours, plots, and so on. All these tools use CGS as their foundation.

Graphics utility programs offering more routine functions that require high quality graphics displays are also available. Programs for producing publication quality 2-D plots, displays of 3-D finite element models, graphic analyses of hydrodynamic data and presentation-quality text slides are categorized in this group. The utility programs run interactively and provide various levels of graphics editing capabilities. Several of them explicitly support a movie output option allowing users to generate multiple frames of graphics for dynamic applications that require graphics interpretation. All of the programs and libraries in this category allow users to view graphics output on a local interactive graphics terminal or alternatively, a pen plotter. Graphics output also can be saved in a device independent metafile for future viewing and analysis. It should be noted that in all these cases, existing software was used wherever possible to avoid massive inhouse development efforts.

Built-in convenience

The metafile and postprocessor combination are significant factors in the system's flexibility. With them, graphics output from programs run in batch mode can be saved for later interactive viewing and analysis, or the same graphics can be sent to any or all graphics devices, as required. A few selected frames of a movie may be previewed before staging the results to a film recorder, or a commonly used display can be stored for future use. The metafile postprocessor also has the ability to overlay or superimpose displays and zoom in on "windows" within a display.

A complete set of tools

The graphics users at the Lab have a complete set of easy-to-use tools at their disposal. Difficult research projects undertaken by the experienced become easier to execute with the available graphics routines, and inexperienced users can use the system with confidence. The system is used heavily. Ray Elliot, Group Leader of the Computer Graphics Group, estimates that 100,000 graphics images are generated on the lab's hardcopy graphics devices daily, some 60,000 of which come from CRAY systems, and that doesn't include the ongoing interactive use of the system. Elliot speaks highly of the system's users and is clearly pleased to relate that individuals with no computer graphics experience can begin to use the graphics system, and within a few weeks, try ways of presenting an idea or the result of a scientific analysis that challenge even the Los Alamos system. □

CORPORATE REGISTER

Canadian Meteorological Centre, SOHIO order CRAY systems

At the close of the year, Cray Research received a number of new orders, including ones from the Canadian Meteorological Centre and Sohio.

The Canadian federal government awarded a contract to Cray Research for the lease and maintenance of a CRAY-1 S/1300 computer system at the Canadian Meteorological Centre in Dorval, Quebec. The new supercomputer will be installed during the fourth quarter of 1983 and will be used to provide weather forecasts for Canada. The contract calls for an upgrade to a CRAY X-MP at a later date.

Additionally, a CRAY-1 S/1300 was installed at the Sohio Petroleum Company in Dallas during the first quarter of 1983. The CRAY system at Sohio will be used in production seismic processing and seismic research.

Shell to install CRAY system in the Netherlands

Cray recently announced that a CRAY-1/S system has been installed for Shell Research, B.V. at the Koninklijke/Shell Exploratie en Produktie Laboratorium (KSEPL). The system was delivered at the end of March and should be fully operational by midyear.

The activities of KSEPL include research into and development of new methods of geophysical data processing and new reservoir models. This Shell Research Laboratory is located at Rijswijk, a suburb of The Hague, in the Netherlands.

Hendrik Vergunst, Computer Services Manager at KSEPL, commented about the recent order for the CRAY-1/S, "We are now using Univac and VAX computing systems in the Laboratory. However, our users have experience in developing new reservoir models on the existing CRAY-1 system in Shell."

"The decision to order a CRAY," Vergunst explained, "represents an enormous breakthrough. The machine has been tested by geophysical engineers at both KSEPL and Shell UK. The CRAY proved to be able to process our present seismic software with considerable speed. At the same time, it was demonstrated that the calculation power of the machine makes presently untrodden paths of research feasible."

NCAR to receive second CRAY-1 system

Cray Research was awarded a contract by the University Corporation for Atmospheric Research recently for the installation of a CRAY-1A computer system at the National Center for Atmospheric Research (NCAR). A one-million word CRAY system will be installed at the NCAR facility in Boulder, Colorado early in 1983. The CRAY system will be NCAR's second; back in 1977, NCAR installed the second CRAY-1 system ever produced by Cray Research.

NCAR, in conjunction with universities around the country, conducts studies in the atmospheric sciences, including international programs in cooperation with scientists and engineers worldwide. The Center is sponsored in these activities by the National Science Foundation, which initiates and supports fundamental and applied research in all scientific disciplines.

"Machines of the capability of the CRAY-1 have made possible many important advances in the atmospheric sciences," Dr. Wilmot N. Hess, director of NCAR, said. "In the areas of climate research, oceanography, severe storms, and sun-earth relationships, both speed and memory are of critical importance. Our first CRAY-1 opened many research doors for us and our university colleagues. The second CRAY will serve many scientsts who are waiting for additional resources."

Chippewa Falls gears up for integrated circuit production

Cray Research's new IC fabrication facility, adjacent to the company's development building in Chippewa Falls, will be completed soon. This new 10,000 square foot building houses a 3,000 square foot Class 100 clean room that will be equipped with state-of-the-art fabrication equipment. The move into the building will begin carefully as major pieces of equipment are installed throughout the year. Del Eberlein, Chief Circuit Designer, said that the operation will start up slowly in 1983.

"This year," he explained, "we expect that only a handful of chips will come out of this facility. We'll be working with wafers that already have the transistor and resistor devices on them. In 1984 we will be making the metal interconnections to complete the chips. By 1985 we should be well into production."

The fabrication facility will provide logic chips necessary for ongoing systems production. In addition, the facility will work with the Development, Engineering and IC packaging groups in the development of chip technologies for future Cray products.

TDC installs a Solid-state Storage Device

Technology Development of California (TDC) maintains an advanced computational facility for NASA Ames Research Center in Mountain View, California. Since late 1981, a CRAY-1/S computer has been a major component of the network. In November 1982, TDC installed a sixteen-million word Solid-state Storage Device (SSD) to augment the CRAY's power and versatility.

Glenn Lewis, Operations Manager for TDC explained, "Certain code written for the Illiac IV had not been transferable to the one-million word CRAY. In addition, data transfer from disk was just too slow. The extra 16 million words and rapid transfer rates of the SSD provide NASA users with power they need." The major classes of problems at NASA that place such demands on the system are derived from computational fluid dynamics and chemistry. Lewis commented, "The performance of the SSD has been very good. In fact it virtually eliminates blocked I/O time. Our users have been very pleased with the results."

Cray Research completes stock offering

In mid-December, Cray Research completed its third successful public stock offering, selling 800,000 shares of common stock. The offering was made through L.F. Rothschild, Unterberg, Towbin; Rothschild Inc.; Piper, Jaffray & Hopwood Incorporated; and Morgan Stanley & Co. Incorporated. The offering generated approximately $30 million, which will be added to the company's working capital for uses such as new product development, system leasing and other corporate purposes. Approximately one-third will be used for further development of the CRAY-2.

CRAY-1/S computer system chosen by U.K. Ministry of Defense

Cray Research recently announced that it installed a CRAY-1 S/1000 computer system for the United Kingdom Ministry of Defense at the Royal Armament Research and Development Establishment (RARDE) in Kent, England. The system was purchased and installed in the first quarter of 1983.

UIS upgrades CRAY

Back in 1978, United Information Services, Inc. (UIS) led the way in offering commercial CRAY-1 processing services on a timesharing basis. Today, the CRAY-1 at UIS is at the heart of a remote computing service. UIS purchased its CRAY-1 S/1000 computer system and completed a system upgrade to two million words of central memory.

"UIS has always kept pace with the batch processing demands of the scientific and engineering community," said Donald S. Bates, President of UIS. "As sophisticated simulation software becomes part of the routine analysis requirements, more and more central memory is needed for resolution." The latest upgrade is the second that UIS has completed in the past three years.

UIS, based in Kansas City, Missouri, maintains major computer services operations in the U.S., Canada, and the United Kingdom. It is a member of the United Telecom Computer Group, a part of United Telecommunications, Inc.

Correction made

Please note the production error made in Iain Duff's article "The solution of sparse linear equations on the CRAY-1" which appeared in our last issue. On page 8, Table 4 should read:

Table 4
MFLOP rates for different implementations of direct SAXPY

Vector length	FORTRAN CFT version 1.09	Cray SCILIB	Boeing CRAY-PACK	Alistair Mills' CAL Code
20	17.7	11.3	16.7	32.5
30	21.2	15.0	21.4	43.2
50	25.4	20.4	27.8	55.2
100	29.7	28.0	35.7	58.6
150	31.5	31.9	39.5	59.5
200	32.4	34.3	41.7	59.9
300	33.4	37.1	44.4	62.8
Asymptotic	35.7	44.4	50.0	66.4

The error is the CAL code column of the table. We extend our apologies to both Iain Duff and our readers, and hope that the correction will eliminate any confusion that may have arisen.

APPLICATIONS IN DEPTH

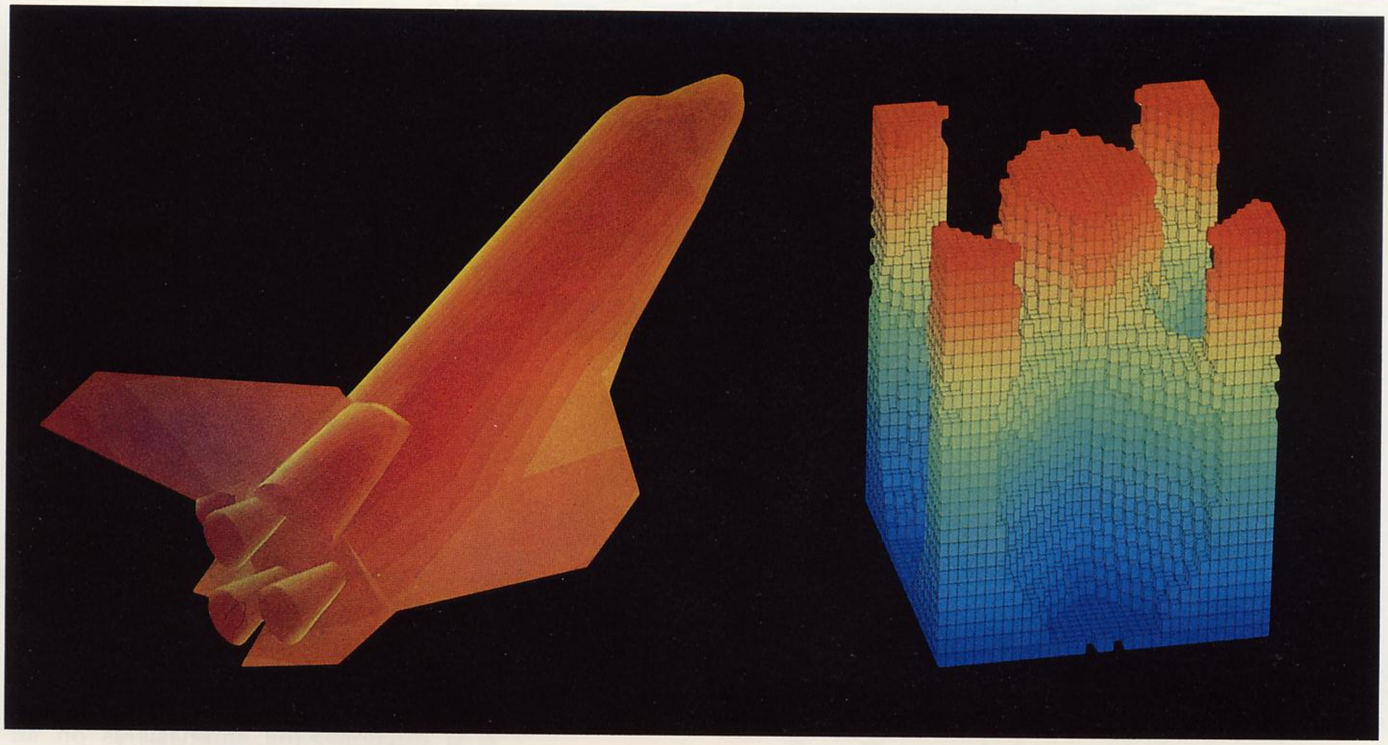

Graphics output, courtesy PDA Engineering and European Software Contractors, Inc.

SIGGRAPH '83

In 1983, the Special Interest Group on Computer Graphics (SIGGRAPH) of the Association for Computing Machinery will hold its Tenth Annual Conference on Computer Graphics and Interactive Techniques. Each year the national conference is held to promote the exchange of information about the latest graphics technologies. It offers a host of educational courses, technical sessions, panel discussions, courses and vendor exhibits, and provides a forum for showing computer-generated film, video and art. Volunteers in the computer graphics community from scientific and industrial entities such as Ford Motor Company, Megatech and the National Research Council of Canada comprise the organizing committee. Dr. Richard Weinberg, a computer scientist at Cray Research and SIGGRAPH '84 co-chairman commented, "SIGGRAPH is the largest and most respected computer graphics organization in the world. Our annual conference is the place where the latest computer graphics research, products, films and art are unveiled."

The SIGGRAPH conference has grown from 600 attendees in 1974 to over 19,000 in 1982. In 1983 it is expected that 20,000 to 30,000 people will participate. The tremendous growth is one more indication of the impact that computer graphics is having on all facets of science and industry.

This year's session will be held in Detroit from July 25-29 and will offer a well-rounded program for experts and newcomers alike. The conference will focus on industrial applications with emphasis on CAD applications robotics. It is being held in cooperation with the Engineering Society of Detroit, the IEEE Technical Committee on Computer Graphics, and Eurographics. Cray Research will host a booth at the exhibit.

For further information about this year's conference contact: SIGGRAPH '83 Conference Office, 111 East Wacker Drive, Chicago, IL 60601, phone (312) 644-6610.

USER NEWS

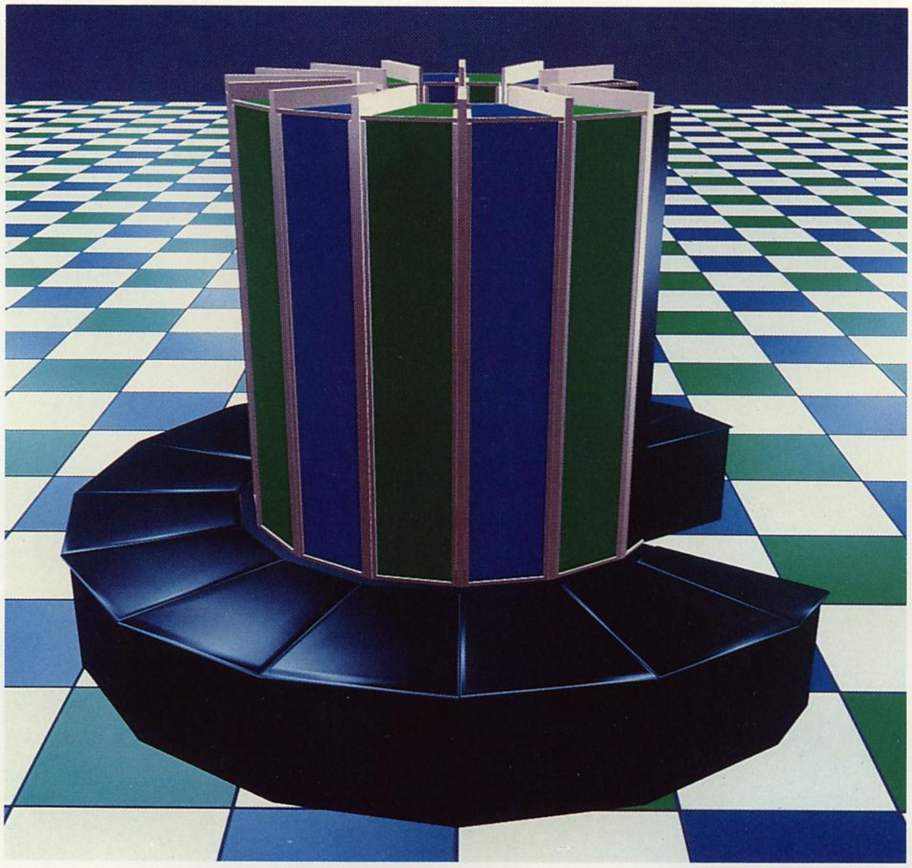

Image of the CRAY-1 generated by the CRAY.

The CRAY-1 — streamlining a self-portrait

Using the CRAY to draw a CRAY... a unique imaging technique conceived by Steve Williams. His involvement began when he volunteered to develop a cover graphic for a computer center brochure at Lawrence Livermore National Laboratory (LLNL) where he was formerly employed. His interest continues, and he still finds time to refine and enhance his creation at Digital Productions, where he currently works. The combined power of the CRAY-1/S computer, Digital Productions' software and the imagination and creativity of the computer artist produced the photorealistic image accompanying this article.

"There's something appealing about having a CRAY generate its own image," Williams states. "If time permits, I'll continue to refine it clear down to the nuts and bolts level." Initially, the CRAY was described as a series of blocks, with minimum color and highlighting characteristics. The first database produced a rough outline of the CRAY, which was used for positioning and testing. It was refined on successive trials, with more information added with each iteration.

The database was processed by a program called HILITE, a general-purpose imaging program for polygon-based objects developed by Frank Crow when he was consulting for LLNL. Williams, who maintained this program at the Laboratory, is now expanding its capabilities at Digital Productions.

The original HILITE was limited in the complexity of the objects it could image. Today, it can handle object descriptions of one million polygons. Also, 16 light sources can be defined instead of the original five, and there are plans to increase this number. To speed up the code, Williams is taking advantage of the CRAY architecture by vectorizing the calculations in some of HILITE's key subroutines. This software, coupled with others under development at Digital Productions, will produce the most sophisticated computer-generated imagery ever achieved. The necessary attributes for achieving the realistic rendering of objects include transparency, shadows, textured surfaces, mirrored surfaces, glowing effects and refraction. As new algorithms are developed, Williams will apply them to the CRAY database.

Williams is enthusiastic about his current work. At Digital Productions, the emphasis is on automating the tedium of image-generation, thereby allowing creative minds to conceive and create new images.

"To make computer graphics a viable film and art production medium," explains Williams, "creative people must be able to use it productively. Programs must be easy to use and extremely flexible in their capabilities. As one of the innovators in this industry, I try to analyze what elements are needed to make an image appear real, and then automate them as much as possible."

Acknowledgement

Digital Scene Simulation℠ by Digital Productions, Los Angeles, CA.

Corporate Addresses

Worldwide Business Centers

Domestic

Central Region Sales Office
5330 Manhattan Circle, Suite F
Boulder, CO 80303
Tel: (303) 499-3055

Albuquerque, NM
Chicago, IL
Minneapolis, MN

Petroleum Region Sales Office
5858 Westheimer, Suite 500
Houston, TX 77057
Tel: (713) 975-8998

Dallas, TX

Eastern Region Sales Office
11120 New Hampshire Avenue
Silver Spring, MD 20904
Tel: (301) 681-9514

Atlanta, GA
Boston, MA
Laurel, MD
Pittsburgh, PA

Western Region Sales Office
5776 Stoneridge Mall Road
The Atrium, Suite 350
Pleasanton, CA 94566
Tel: (415) 463-2800

Los Angeles, CA
Seattle, WA

International

Cray Research France, S.A.
21 Avenue Victor Hugo
75016 Paris, France
Tel: 01-501-5270

Cray Research Japan, Limited
W 1154, Shin Aoyama Building
1-1 Minami-Aoyama 1-chome
Minato-Ku, Tokyo 107, Japan
Tel: 03-403-3471

Cray Research GmbH
Perhamerstrasse 31
8000 Munich 21, West Germany
Tel: 89-56014-0

Cray Research (UK) Ltd.
Seymour House, The Courtyard
Denmark Street, Wokingham
Berkshire, England
Tel: 0734-791180

CRAY
RESEARCH, INC.

Corporate Headquarters
Cray Research, Inc.
608 Second Avenue South
Minneapolis, MN 55402
Tel: (612) 333-5889