Interactive Computer Model Fitting to Observed Space Data: A New Concept in Real-Time Visual Analysis and Display

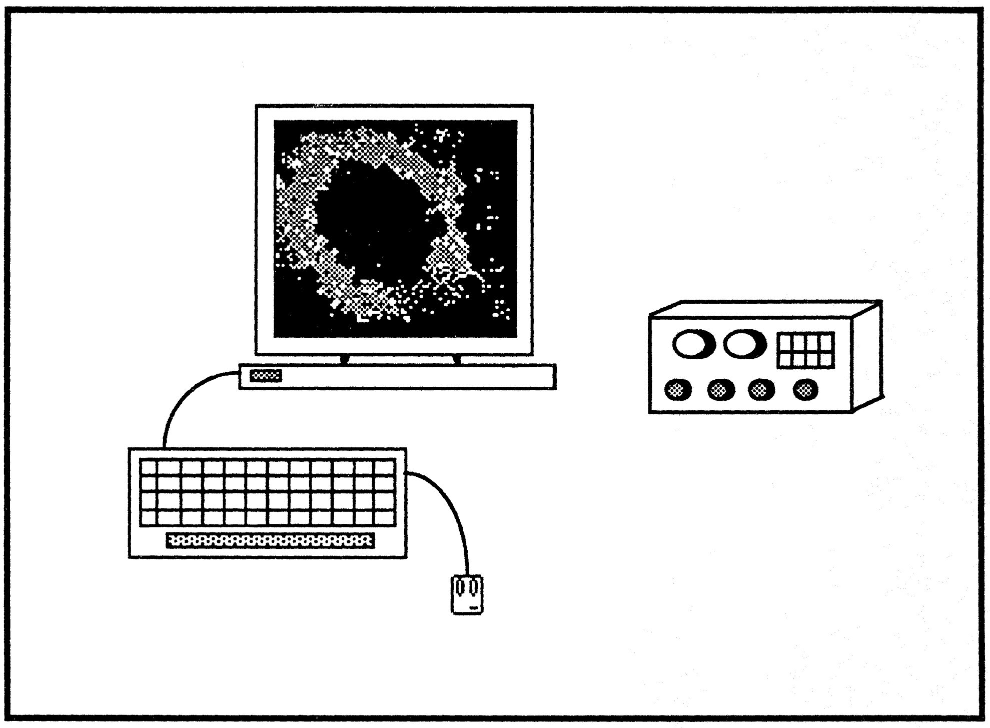

Dr. Jan M. Hollis
Assistant Chief
Space Data and Computing Division
NASA Goddard Space Flight Center
Greenbelt, Maryland

Dr. Robert Burns
Head
Computer Division
National Radio Astronomy Observatory
Charlottesville, Virginia

Goddard Space Flight Center

RESEARCH PROPOSAL

Submitted to NASA Headquarters, Office of Space Science and Applications, Code E

INTERACTIVE COMPUTER MODEL FITTING TO OBSERVED SPACE DATA: A NEW CONCEPT IN REAL TIME VISUAL ANALYSIS AND 3-D DISPLAY

RTOP 656-80-01 Task 01

A Joint Proposal by

NASA Goddard Space Flight Center
Greenbelt, MD 20772

and

National Radio Astronomy Observatory
Charlottesville, VA 22901

PI: Dr. Jan M. Hollis, Code 630
Assistant Chief, Space Data and Computing Division
Tel: (301) 286-7591

PI: Dr. Robert Burns, Head
Computer Division at the National Radio Astronomy Obs.
Tel: (804) 296-0211

Co-Investigators

Dr. R. Pisarski, Code 630
Dr. R. White, Code 680/ARC
Mr. E. Sullivan, Code 684

Dr. D. Wells, NRAO
Dr. E. Greisen, NRAO
Dr. W. Cotton, NRAO

Research Area/Key Words:
Graphics visualization, interactive modelling, mini-supercomputer workstations, science analysis, COBE

Submission Date: June 27, 1988

Cost and Manpower:	FY89	FY90	FY91 (in K$)
Requested R&D Funding	200	200	200
Civil Service Man Years	1	1	1

Approval:

Dr. J. H. Trainor
Director of Space and Earth Sciences
Code 600

INTERACTIVE COMPUTER MODEL FITTING TO OBSERVED SPACE DATA: A NEW CONCEPT IN REAL TIME VISUAL ANALYSIS AND 3-D DISPLAY

TABLE OF CONTENTS

PREFACE

The ultimate goal of computer-aided modelling is to create a system that will enable the researcher to graphically visualize modelling in progress, exercise intuitive judgement by intervening while the model is being computed (e.g., to change model parameters and/or to change rendering attributes), and continue to monitor the emerging results as they are displayed in near real time. Fortunately, from a computational speed viewpoint, such a goal can be achieved with technology presently developing in the mini-supercomputer workstation (MSWS) market (e.g., refs 1 and 2), coupled with innovative implementation of a programmable analog control panel which would afford model and/or visualization parameter changes on-the-fly. In particular the MSWSs are also powerful graphics (e.g., refs 1 and 2) rendering engine workstations with I/O characteristics complementary to their computational speed. What is needed is an implementation of a universal imaging software system to take advantage of the unique hardware system we have envisioned. Further, if the basic tools of that software development can be transported to other, presumably less expensive, hardware hosts and reused, then cost-effective science analysis will result.

The foregoing motivates this Goddard Space Flight Center (GSFC) and National Radio Astronomy Observatory (NRAO) joint proposal whose strategy is to combine the expertise of the science applications programming within the Cosmic Background Explorer (COBE) project, the Astronomical Image Processing System (AIPS) expertise of the NRAO, and the system integration expertise of GSFC/Code 630 in order to demonstrate the concept by focusing on the COBE mission. Joint funding from both NASA and NSF is being sought. Although a demonstration project, the COBE mission's science analysis function has general applicability to any ground-based or space experiment which utilizes two-dimensional format detectors. Moreover, the basic software tools and control panel development provided by this effort will have applicability to any science discipline since the emphasis will be accurate image registration and interactive modelling, both of which will allow the scientist to better visualize and analyze large data sets.

I. BACKGROUND INTRODUCTION

1. Mini-Supercomputer Workstation Hardware Available

A VAX 780 has an instruction execution speed of 1 million instructions per second (MIPS). A typical workstation presently available has the processing power equivalent to several MIPS. The cutting edge of technology available currently is vested in the mini-supercomputer workstation (MSWS) which has 'supercomputer' like performance with regard to central processor unit (CPU) and input/output (I/O) speeds. The MSWS is additionally known as graphics supercomputer workstation since it is also a high performance graphics display unit. These new units are advertised to have about 20 MIPS scalar speed, 30 to 60 million floating point operations per second (MFLOPS) pipeline speed, large memory size, striped disks, 1280 by 1024 video with stereo capability, and graphics (polygon rendering) speeds in excess of 100,000 polygons per second. Such MSWS attributes will soon permit a variety of new approaches to the problem of visualizing large digital images in an interactive manner in near real time. In particular, stereo viewing of polygon-rendered three-dimensional data cubes, with the viewing orientation under the interactive control of the user, appears to have an excellent chance of being the next breakthrough in image visualization technology for science applications.

2. Programmable Analog Control Panel

It is presently common for high-performance graphics displays to use analog knobs, trackballs, or joy-sticks to allow the user to control the TV display of the "object" being viewed. The vendors of the new MSWS also offer optional "control boxes" to support and extend this functionality to as many as 20 programmable knobs. Such a combination of analog-programmable interaction with the MSWS and its powerful graphics capability permits the advent of "interactive modelling" with regard to astronomical image processing. Imagine, for example, that the user has modelled his huge volume of data as processed on the MSWS and wants to view the residuals in real time. Certainly he can do this much faster with the MSWS than before. However, the unique aspect of the analog controls permits interactive model coding in such a way that the user can change the physical parameters of his model (e.g., temperature, density, etc.) with a knob and "visualize" the model change in near real time as he graphically views the results on the MSWS. This concept is new and, if developed, would probably revolutionize the way science is accomplished with respect to large data sets.

3. MSWS Application Software <u>Not</u> Available

The advent of the MSWS and the concept of interactive modelling provides a unique opportunity for NASA to encourage the technological advance of this powerful combination. This means that the interactive modelling concept and the new technology hardware to support it will have to be programmed in some application software system that enjoys widespread scientific use. The foremost such system for astrophysics is the Astronomical Image Processing System (AIPS) which has been developed at the National Radio Astronomy Observatory in Charlottesville, Virginia (cf. ref 3). Once such systems are in general use, they will allow enhanced visualization of data and models and interactive image analysis as previously stated. But they will do more than this. They will almost certainly drive networking to higher rates so that large volume data sets at regional data centers (such as the National Space Science Data Center at NASA/GSFC) will be routinely browsed by scientists from all over the country. Moreover, such MSWSs will undoubtedly become the interface of choice to remote supercomputing resources.

II. OBJECTIVES AND JUSTIFICATION

1. What is Proposed to be Accomplished

We propose a collaborative three-year research and development project to harness the MSWS as a future tool for data analysis. First, we intend to research techniques to apply the benefits of the MSWS (particularly its speed, display capabilities, and potential new interactive capabilities) to astrophysics and we will use the COBE science analysis requirements as a testbed mission. Second, we intend to extend and develop the NRAO AIPS software package to incorporate the MSWS attributes and subsequently make such development efforts available to the science community at no cost.

Specifically in the first year, NRAO (in consultation with Code 630) will competitively purchase two identical MSWSs, one to reside in Charlottesville, one to reside at GSFC. NRAO will port AIPS to the MSWS environment to begin the research and development effort as both a stand-alone system and as a display engine attached to a host. The port will be implemented to allow the new AIPS displays to be as downwardly compatible as possible with existing conventional displays (e.g., a VT240 graphics terminal). The host at GSFC could range from a VAX 8800 to a supercomputer; at NRAO the host would undoubtedly be a CONVEX-class mini-supercomputer. In any case, the host will be used as a file server in the case of medium-performance systems and as a compute server in the case of a supercomputer facility. At GSFC, various science groups will be invited to port their software to the MSWS and to experiment with the new configuration, capabilities, and software. Foremost among these will be members of the COBE science working group (SWG), who will be visiting Goddard frequently as COBE prepares for launch and begins operation. We will thus involve members of the outside community as early as possible. NASA scientists will provide feedback and guidance for software development, based on their experience and initial explorations of the MSWS.

At the end of the first year, a highly flexible and portable new version of AIPS will be available to the entire community on at least two different MSWSs. This will result from NRAO benchmark testing of AIPS on at least two competitive machines. Hence, the second-place finisher in the competitive procurement will at least have a sophisticated AIPS implementation on its machine as a direct result of the procurement process itself.

During the second year, Code 630 and the COBE project will provide personnel to learn AIPS in the MSWS environment as it is being developed at NRAO. Moreover, in this second year NASA expects to use the MSWS for intensive off-line analysis of COBE data and sky models. The results will allow for comparison with those from conventional techniques which the COBE project will be supporting. The MSWS would supplement, but not replace, facilities provided by the COBE project, which will still be used to produce the Project Data Sets and Analyzed Science Data Sets. NASA scientists will participate in writing the requirements for the interactive modelling package, and various NASA groups will be invited to exploit the MSWS outside of AIPS. NASA will provide the principal scientific feedback during the software development. At the end of the second year, the MSWS/AIPS combination will be routinely utilized for COBE data processing at GSFC.

Concurrently in the second and third years, NRAO will expand the MSWS/AIPS display capability in three directions. First, the high transfer rates in the MSWS frame buffers will be exploited to support the movie-loop display of multi-spectral data. Second, polygon-rendering technology will be adapted to astronomical data analysis problems. Although polygon projection has been used extensively in the commercial graphics area, particularly in the areas of animation, modelling, and scene generation, and even in the presentation of astrophysical modelling calculations, it has not been extensively applied for visualization of observational data. Third, NRAO will adapt windowing technology and programmable analog control boxes to provide interactive user control of display modes and model fitting. This concept of the control panel is subsequently discussed in more detail.

The analog control panel concept will allow for truly interactive data model fitting, which is one of the most exciting results of this project. Currently, most scientists fit data through a laborious series of steps. For example, a model might be fit in a batch job overnight. The next day, the scientist examines the data, makes a decision on any model changes or parameter tweaking, and submits another overnight batch job. This can repeat over many days. Generally, the scientist cannot explore as much parameter space as he would like, nor can he compare systematic residuals very easily. With the high speed of the MSWS and its displays, and with possibly twenty parameter values adjustable interactively through the control panel, the scientist could collapse the long iteration process into a single session. The scientist could tell, by displaying residuals, which fits have most of the chi-square in systematic deviations from the model, and which aspects of the data and model are most sensitive to which parameters. For example, the scientist could blank regions of a complicated set of observations such as the bulge and disk of a galaxy, fit the individual components separately, and then put the parts back together. Figure 1 crudely illustrates this concept of interactive modelling in an image context, and a control panel would make the computer a very interactive extension of the scientist's hands, eyes, and brain. The MSWS will be a very advanced version of a scientist at the blackboard or doing a back-of-the-envelope calculation.

At the end of the third year all the unique aspects of MSWS display and interactive enhancements will be available for distribution. Periodic technical reports on the various computer science and analytical technique aspects of the projects will be written, so that the experience gained will be available to the general community and much of the research will not have to be repeated for other NASA-supported systems such as the Hubble Space Telescope's Image Reduction and Analysis Facility (IRAF) and the Land Analysis System (LAS).

2. Justification for the Proposal

The enormous increase in the volume of scientific data, particularly from space missions, and the parallel increase in data quality have put enormous strains on the computer resources of the scientific community who must analyze these data. Increased network and computer speeds are needed to keep up with the data rates, and more sophisticated analysis techniques are needed to extract the new science that the data contain. Such new analysis techniques, in turn, are usually computationally expensive which puts further demands on the hardware. Such a positive-feedback scenario justifies the proposed effort. We subsequently outline some of these new techniques and their application to astrophysics:

a) Interactive display and modelling using color, perspective, rotation and stereo presentation: These tools can help the scientist to examine and analyze large multi-dimensional data sets using the integrative capabilites of the human eye and brain. Examples would be modelling the distribution of galaxies and clusters and comparing the results to sky surveys, or examining an image of a galaxy with a third dimension of velocity and comparing the observations to dynamical and mass-distribution models. Models, observations, and residuals could all be displayed.

b) Fourier analysis and related techniques: These can be used to look for large and small spatial structure scales, anisotropies in the cosmic background radiation (spatial and spectral), for example. Three dimensional FFTs, 2 spatial and one in time, can be used, for example, to filter out the 5 minute oscillations on the Sun, allowing studies of other components of solar variability.

c) Time series analysis, digital filtering and image classification: Studies to characterize the infrared cirrus including fractals, non-symmetric deconvolutions and sharpening of crowded field images, and classification and principal component analysis of multi-spectral band images using Landsat techniques are a few examples.

Knowledge of new techniques is not enough. The scientist must have them conveniently available, in a software system that will run on his own familiar computer. Our proposal to install implementations of these techniques into AIPS would meet these needs.

3. Interactive Modelling: General Applicability to Astrophysics

Since the invention of the telescope, the classic problem in optical astronomy has been removing the effects of the terrestrial atmosphere on astronomical data. During this century, optical astronomers have learned to carefully account for corrections for atmospheric absorption and estimate the effects of clouds. In the process, they have discovered interstellar reddening, interstellar clouds of dust and gas, wavelength variations of extinction (reddening), variations in interstellar dust parameters as a function of position on the sky, interstellar polarization, magnetic fields, and so forth. All of these effects must be carefully accounted for, depending on the observations being conducted.

Astronomy today uses two-dimensional format detectors across the entire electromagnetic spectrum from gamma rays to radio wavelengths. In order to extract the full significance of any two-dimensional map of the sky in a particular wavelength regime, extensive interpretive studies with computer-intensive image processing techniques are required. Today's astrophysical science at the cutting edge attempts to extract very weak signals from the noise in the presence of much stronger signals emanating from sources not under investigation. For example, depending on the weak source to be studied, stronger sources of confusion could be the Earth's atmosphere, the interplanetary medium, the interstellar medium, the galactic signature, or even the intergalactic medium. Hence, even an astronomical observation from space will need unwanted astronomical signatures removed from the data. Moreover, the instrumental signature itself has to be perfectly accounted for in the data processing. Thus, models of the sources of confusion are constructed and literally subtracted from the data so that the noise can be investigated for signals hopefully heralding a new discovery about the Solar system, the Galaxy, or the Universe. Interactive astrophysical modelling would permit such investigations to progress at an unprecedented rate. This increase in rate is necessary to efficiently and quickly comprehend the huge volume of data produced by modern instruments. Astronomers today have more data than they can either comprehend or process and, in that sense, are literally 'data bound.'

a) COBE Project as a Demonstration Project

The Cosmic Background Explorer (COBE) satellite will be launched in approximately May 1989 to map the sky at wavelengths from 1 micron to 1 cm. Its mission is to discover small anisotropies and spectral distortions in the 2.7 K cosmic background radiation, to find the radiation from the first objects or structures to form after the Big Bang, and to study all other sources of diffuse radiation at these wavelengths. These other sources include interplanetary and interstellar dust, hot electrons in the Galaxy, faint stars in the Galaxy, and possibly IR galaxies and hot gas in galaxy clusters. In order to map these primeval and local sources, the COBE will scan the sky repeatedly with its three instruments, building up signal-to-noise until the data are only limited by the astrophysical environment. The three instruments are the DIRBE (Diffuse Infrared Background Experiment) covering 1 to 300 microns with a 10 band filter photometer and a 0.7 degree beamwidth; the FIRAS (Far Infrared Absolute Spectrophotometer) covering 100 microns to 1 cm with an absolutely calibrated polarizing Michelson interferometer with a 5 percent spectral resolution and a 7 degree beamwidth; and the DMRs (Differential Microwave Radiometers) covering 31.4, 53, and 90 GHz with 7 degree angular resolution. All three instruments will make spectacular advances in the sensitivity and accuracy of the data available, and are expected to lead to the discovery of new phenomena in the early universe.

The COBE has two 'atmospheres' with which to contend: (1) the interplanetary atmosphere and (2) the galactic atmosphere. Both atmospheres possess a basic symmetry with significant deviations from that symmetry. The brightness of interplanetary dust is a strong function of solar elongation angle and ecliptic latitude, and at short wavelengths it scatters polarized sunlight. The dust is approximately symmetric about the ecliptic plane and independent of longitude, but the nominal symmetry plane is distorted as it passes near the various planets, there are streams (bands) in the dust due to sources in asteroid family collisions, there are resonances with major planets which spoil the rotational symmetry, and there are traces of the trails of comets. The interplanetary dust composition and size distribution may differ with distance from the Sun as well. Similarly, galactic infrared emissions clearly have the general distribution seen at other wavelengths, with a concentration in a plane, a greater concentration toward the center, and an irregular (cloudy) distribution with variations on all angular scales from fine filaments to huge clouds and loops. The galactic dust temperature is, on the other hand, a very slowly varying function across the sky.

All of these phenomena can be modelled numerically. There are many observations at many wavelengths to be explained simultaneously by a single model, so some complexity is required. On the other hand, there are fewer adjustable parameters than observations. Simulations have shown that it should be possible to model the interplanetary dust emissions and subtract them from the observations with an accuracy of the order of 1 percent. Similar accuracy for the interstellar sources is less likely because of the cloudy distributions.

There are also important concerns about the effects of possible errors in the instruments. The instruments have been designed to avoid all the known errors to an adequate degree, much better than other instruments of their type. Nevertheless, some additional phenomena may have to be modelled with a computer. As an example, the terrestrial and spacecraft magnetic fields affect the DMR slightly, charged particles add noise to the FIRAS and DIRBE detectors, and the expected high solar activity may raise the Earth's atmospheric density to a level where its emissions could be seen with the DIRBE. There is a small effect of the instrument sidelobe responses to the Moon, Earth, and the Sun, particularly on the DMR data, which is amenable to numerical modelling.

The limiting factor for the scientific value of the COBE observations will be our ability to model, explain, and subtract the local background sources and derive a residual of cosmic origin.

The COBE data interpretation is one of the most interesting and difficult parts of the mission, because of the large volume of data, its dependence on multiple parameters (right ascension, declination, solar angle, wavelength, and polarization), and the multiple local sources of diffuse radiation. Displaying these data and adjusting the parameters of models to fit them is a key part of the COBE mission.

The COBE Project has previously selected the AIPS image processing system as the basis for the final display and analysis. Without the benefit of this proposal, the COBE data analysis plan is to generate FORTRAN models on a VAX computer, transfer the results to the AIPS system in FITS format, display them and analyze them, and then return to the FORTRAN modelling program to iterate another model. The procedures described subsequently in this proposal would permit integrated modelling, display, and analysis on a single computer with high speed and user convenience. It would clearly enhance the accuracy and value of the final products, and enable the results to be obtained in much less time. Figure 2 illustrates the kind of new display and visualization techniques we would like to explore with the COBE project, using the MSWS.

b) Why this development effort is necessary to present and future science analysis

The emphasis on computation for the COBE is not unusual. There are many cases where computers are an integral part of a mission. Doppler tracking of satellites and space probes has become accurate enough to test general relativistic corrections. In radio astronomy, the CLEAN algorithms have enhanced the performance of interferometers by many orders of magnitude. Computers model stellar interiors, galactic rotations, collisions, formation, and clustering. In particle physics, sophisticated computer codes make Monte Carlo simulations to remove biases from detector systems, account for complex statistics, and detect patterns hidden in vast amounts of data. Computers have enabled tremendous advances in knowledge, and in many cases they are now limiting further advances.

A recent example for NASA is the IRAS project. Elementary image processing was introduced at a late stage in the project, as people realized the importance of large scale structure across the sky. Huge interstellar clouds and loops were seen, zodiacal dust bands were found, and many correlations have now been made with large scale maps made at other wavelengths. Image processing enhanced the visibility of these features, but was also important in settling the final error correction of the basic data. Early IRAS maps had bright stripes across them, caused by errors in calibration constants of the detectors, some of which were not very stable themselves. Image processing revealed the patterns and enabled physical models and mathematical corrections to be made.

4. Specific Interests Common to NRAO and GSFC

Both institutions have huge three dimensional data sets and attendant visualization problems. Moreover, each institution must make effective use of high-performance computers in relationship to MSWSs; in the case of GSFC, a new large-scale supercomputer will be purchased in the 1991 time-frame; in the case of NRAO, computationally intensive needs are anticipated to be soluble with several mini-supercomputers of the CONVEX class. Further, both GSFC and NRAO need to begin now with the new MSWS technology to meet future computational requirements in the area of graphics visualization. This will expedite science data analysis and provide relatively low-cost visualization tools to researchers to accomplish strategic science goals, both in-house and more importantly to the general science community.

III. TECHNICAL ASPECTS OF THE MSWS HARDWARE CONFIGURATION

At this juncture there are at least two MSWSs that are currently commercially available: The Graphics Supercomputer from Stellar Computer (Newton, MA) and the Titan from Ardent Computer (Sunnyvale, CA). What follows is based, in part, on vendor brochures and several market reviews (e.g., refs 1, 2, 4, 5, 6, 7, 8).

1. Windows

The window concept is the most distinctive feature of all technical workstations. Rather than have an inflexible hardware-based terminal, a completely programmable video frame buffer is made to emulate a variety of imaginary graphics devices. The devices are software abstractions. Indeed, the window system software enables multiple instances of such displays to coexist in the frame buffer, and to overlap each other, with their positions, sizes and overlapping controlled by the user by means of a graphics input device. The input device is almost always a "mouse".

Mouse-driven input from the user is quick and efficient and, even more important, is intuitive. The user quickly learns to control the displays with the mouse. The window software supports a variety of input and display mechanisms, notably menus and "buttons". The menus "pop-up" or "pull-down" on demand, and the buttons emulate the behavior of hardware switches for control of software. In the best software of this type, the user performs most operations by pointing at graphical displays and "picking" options and commands with the mouse; only text input values need be entered on the keyboard.

The frame buffers are arrays of bits, one or more bitplanes deep, from which the video signal is produced by digital-to-analog converters. Text is produced by setting patterns of bits in the bitplanes; thus a variety of fonts is available. Arbitrary graphical information can be drawn in the bitplanes by software; a workstation window can be programmed to emulate any classical graphics terminal. In particular, it is typical for the vendor to supply an emulation of the traditional Tektronix "green-screen" terminal (the AIPS graphics interface standard). A workstation frame buffer with more than one bitplane, especially one with 8 or more planes, is capable of emulating a digital image display which has a similar number of bit planes. Such workstations generally have lookup tables to facilitate grey-scale and pseudo-color image enhancement. It is, therefore, feasible to construct a software emulation of a digital image display. During 1987 several such displays were constructed at AIPS user sites for use with AIPS on Sun Microsystems workstations; early in 1988 the AIPS Group itself picked up this development work. We expect that it will be a straightforward exercise to adapt one or more of these designs for use with the MSWSs.

A window-based image display has both strengths and weaknesses compared to a hardware-based display. Two weaknesses are apparent with the earliest limited number of bitplanes. Such basic display operations as "zoom" and "pan" are implemented in software rather than hardware; performance depends on memory and CPU speeds. Although it is widely agreed that current performance is acceptable for scientific work, large-format images are conspicuously harder to manipulate than for hardware-based displays. The answer to this problem is simply speed in the workstations. For this purpose, the MSWSs are about 10x faster than present conventional workstations, and, therefore, we expect that the speed problem will be almost completely eliminated.

The bitplane problem is that with only 8 or 10 planes, a typical current WS simply cannot emulate the behavior of a hardware-based display which has 28-36 bitplanes. Full-color imagery and fast image blink operations cannot be performed. The answer to this problem is simply MORE BITPLANES in the workstations. We expect to procure MSWSs with about 50 bitplanes.

There are two main strengths of the window-based image displays: they are "virtual" devices of great flexibility and they are servers on local area networks (LANs). The flexibility of the software approach enables these displays to transcend the limitations of their physical frame buffers. For example, the window can be considered to be a "porthole" which looks at a much larger true image. The panning action can be programmed to roam through a 2-D space larger than the physical frame buffer size. Likewise, blink and movie-loop operations can be programmed to roam though a 3-D space larger than the number of image frames contained in the bitplanes of the frame buffer. Finally, the virtual display can be programmed to support image enhancement modes which are difficult to implement in hardware, such as intensity-hue-saturation color rendering.

The notion of the image display as a network service is a critical component of true distributed image processing. AIPS processes in compute servers can display imagery in window-based image displays on workstations. AIPS processes in the workstation are also able to display imagery in the windows and, even more important, to execute interactive image enhancement and cursor-setting operations on the imagery downloaded from the compute servers. If image data files are visible to both the workstation and compute server, an elegant, flexible, natural division of labor is available: (1) the AIPS in the workstation can load and enhance an image, and the AIPS in the compute server can make a cursor setting on the image to specify regions of interest or other parameters for further processing; or (2) the AIPS in the compute server can download an image and the AIPS in the workstation can enhance it and make cursor settings to extract values or establish parameters for analytic code executing in the workstation. Users have the freedom to choose which pattern suits their tastes and requirements.

Conventional workstations, even with their present performance and capacity limitations, are already beginning to replace hardware-based image displays. The new MSWSs have speed and capacity almost equal to even the highest-performance, largest-scale traditional displays. All workstations (both conventional and MSWS) are doubling in speed about every 18 months, and so it is now obvious that, over the next few years, the window-based WS image display technology, with its advantages of flexibility and network-server capability, will replace essentially <u>all</u> of the traditional hardware-based image display technology.

2. Concurrent Pipelined Floating Point Hardware

The new MSWSs deliver exceptionally high computational performance due to their use of three supercomputing architectural features: pipelining, chaining and concurrency. The floating point (FP) hardware of conventional computers (e.g., existing VAXes) performs FP operations in a sequence of steps, one per tick of the system clock; five steps in a sequence is typical. Only one such sequence is in progress at any moment. The pipelining concept is that as the operands of a FP operation proceed to the second stage of the sequence, we can execute the first stage of a new FP operation on a new set of operands. On the third tick, the first set of operands proceeds to the third stage, the second set proceeds to the second stage, and yet another set enters the first stage. After five ticks of the clock, the results of the first operation exit from the pipeline, and, on each subsequent tick, another set of results exits. For an operation to be performed on a long vector of operands (length substantially greater than the pipeline length), pipelined FP hardware can produce results approximately five times faster than non-pipelined hardware executing at the same clock frequency.

Pipelining reduces the time cost to one clock tick per elementary operation. Real applications perform combinations of operations on vectors. Chaining permits such combination operations to be performed in one clock tick per operand element as well. The concept is that pipelined hardware is provided for each type of operation which is to execute in such a combination and that, as the results from the first operation (e.g. an ADD) exit from the first pipeline on the fifth clock tick, they are routed to the first stage of the pipeline for the second operation (e.g., a MULTIPLY). On the tenth tick, the results from the combined operation exit from the second pipeline; operands of the nine intermediate stages are still in the chained pipeline and will exit on subsequent ticks. An alternate name for such concurrent vector operation is "vector triads". For a combined operation to be performed on long vectors of operands, chained pipelined FP hardware can produce results approximately ten times faster than non-pipelined hardware executing at the same clock frequency.

Once one has constructed a computer which uses the pipelining and chaining techniques to obtain maximum FP performance from components executing at a given clock frequency, there is one more architectural technique to use to increase performance: concurrency. We can simply use two, or four, CPUs to double, or quadruple, the aggregate performance of the computer system. It is typical that one can get 1.8x higher performance for 1.3x higher price by adding a second CPU to a system. Thus, dual-CPU computer systems are highly favored. The marginal gains for additional CPUs are smaller, but four-CPU systems make a lot of sense.

In order to realize the full potential of pipelining, chaining and concurrency in a wide range of application code, the hardware technology must be properly supported by compiler technology. Until 1985, only the traditional supercomputer vendors offered "vectorizing" compilers. The first-generation mini-supercomputer vendors, Convex and Alliant, demonstrated that they understood this compiler technology and even implemented automatic concurrency analysis, unlike the traditional supercomputer vendors. The concurrent vector hardware of the new MSWSs is fully supported by their compilers.

The MSWSs deliver high computing performance for comparatively low cost by utilizing a judicious combination of pipelining and concurrency (ref. 1). Stellar's MSWS uses four CPUs implemented in custom VLSI logic. Ardent' MSWS uses one, two or four CPUs which they purchase from MIPS Computers. The Ardent CPUs have nominal performance somewhat higher than those from Stellar.

Because the AIPS software is portable, and because the critical algorithms of AIPS have been adapted for ease of recognition by typical vectorizing compilers, the detailed architectural differences between various systems, such as Ardent or Stellar, are largely irrelevant for purposes of making procurement decisions. Instead, benchmark information for execution of full AIPS tasks on real data is the most reliable indication of the relative success of the efforts of the hardware architects and compiler designers.

3. Disk Capacity, Concurrent I/O

We expect to procure MSWSs with at least 1 gigabyte of disk each. This will allow them to manipulate multiple large-format images and 3-D multi-spectral cubes, and will enable them to properly demonstrate their power as stand-alone image processing computers. The goal for displaying large-format astronomical images is 4096-square, in 32-bit floating point (for dynamic range reasons), which is 64 megabytes per image, and we expect several at a time. A few 8192-square images, at 256 megabytes, will be manipulated. For 3-D cubes, the display target is 1024-square 256 channels deep, which is a total of 256 megabytes in 8-bit pixels, or 512 megabytes in 16-bits. Reserve capacity of 2-4x is appropriate in such systems (to handle cases of multiple images and of images 2-4x larger in any dimension), which leads to the 1 gigabyte specification.

Striping is another supercomputer concept which the mini-supercomputer and MSWS vendors have implemented. Concurrent disk I/O is used to synthesize a disk system with several times higher performance. The idea is that the successive disk blocks of a file can be written to two or more different disk drives concurrently, even blocks to one drive and odd blocks to the other for a two-way stripe. The drives must have independent controllers and sufficient channel/bus capacity must be provided and arranged to avoid contention between them. This striping is a system-installation option and it is transparent to the users in the mini-supercomputers introduced during the last two years. Users see a virtual disk drive which has twice the transfer rate and twice the capacity, or four times higher in both for a four-way stripe.

The goal for the peak disk transfer rate of the MSWS will be 8 megabytes per second. This speed will enable it to display 3-D cubes in "movie" mode. Roaming of large-format imagery can probably be done acceptably with a transfer rate of 2 megabytes per second (about a half-second to fill the frame buffer), but movie display is more demanding. The frame rate target is 8 frames per second at 1024-square for 8-bit pixels, which sets the 8 megabyte per second goal.

Another important motivation for the transfer rate specification is that when the MSWS is doing imaging calculations stand-alone, it will need disk bandwidth in this range in order to keep its high-performance FP pipelines busy. Experimental data on several vector computers can be summarized in a simple rule-of-thumb: the AIPS imaging algorithms run well if the ratio of peak disk rate to peak pipe rate is 0.2 MB/s/MFLOP. For example, a Convex C-1/XP, with a peak pipe rate of 40 MFLOPs, runs well, with CPU/real ratios over 90%, if it has a four-way stripe of Fujitsu "Eagles", a configuration which has a peak transfer rate slightly greater than 8 megabytes/second. A recent measurement on a four-way striped C-1 yielded a sustained transfer rate of 3.6 MB/sec (sustained I/O about half of peak is typical). The peak transfer rate goal which we are specifying will effectively require that the MSWS vendors support disk striping in their operating systems in order to assure full performance in both roles, 3-D cube visualization and stand-alone image processing. Both Ardent and Stellar have implemented disk striping in their operating systems.

4. 3-D Perspective Graphics Rendering Engines

Computer-aided-design (CAD) applications of workstations make heavy use of perspective modelling of solid objects. The objects are modelled as collections of connected polygons, and the process of producing an image from a list of polygon coordinates plus a viewpoint location is often called "rendering". Vendors of CAD workstations now compete on the basis of whose hardware/software combination renders polygons fastest. The new MSWSs have special VLSI support for polygon rendering, and as a result they are claiming the highest polygon rendering speeds available today, about 150,000 polygons/second for 10 pixels/polygon average. This means that a data structure of 10,000 polygons (i.e., 100,000 pixels area, 300-square, or 10% of the frame-buffer area active) can be rendered in 1/15th second. Larger numbers of polygons will take longer, of course, but these numbers immediately indicate the major reason why the rendering engines of the MSWSs are so attractive for our applications: scientific data structures of typical dimensions can be rendered at rates comparable to human perception and reaction times, which will permit the exploitation of exciting new interactive modes for the visualization and analysis of our 3-D data.

For example, the RA-Dec-Velocity Hydrogen-line data on a rotating galaxy can be represented as a surface in the three dimensions and viewed in perspective to visualize structures and relationships in the data. Indeed, a movie of such a rotating, perspective view of a 21-cm velocity field was computed (on a VAX-780) at the VLA several years ago as a demonstration of the technique. The rendering hardware of the MSWS will permit that movie to be computed and viewed almost in real time. Indeed, the user should be able to tie the viewpoint to control panel knobs and continuously compute new renderings of the data, thus permitting the full potential of this visualization technique to be realized.

5. Control Panels

During 1983-1985, NRAO constructed a modular, programmable control panel as a part of its ISU (Image Storage Unit) project. The success of this feature of the ISU (shown in Figure 3) has convinced NRAO that such control panels should be an integral part of all sophisticated image displays, and especially 3-D cube displays. NRAO's 1983-85 panel was not the first of its kind. In 1975, the National Optical Astronomy Observatories (NOAO) fabricated a similar control panel for use with the digital image display of NOAO's IPPS (Interactive Picture Processing System). Many hundreds of astronomers used the IPPS display during ten years of production operation (1976-85). The lessons learned in the IPPS panel were incorporated in NRAO's ISU panel, and lessons learned on this latest panel will be incorporated in the panel interface for the MSWS image displays.

The principal lesson we have learned is that knobs are a superior method for controlling analog variables. The most important ergonomic reason for this superiority appears to be the precision and repeatability of angular motion of the human wrist. Linear motion of the hand has substantially lower repeatability---the geometry simply makes absolute position sensing more difficult in the linear case. An additional ergonomic factor is the familiarity which we all have, from childhood, with adjusting electronic devices using knobs. This makes it easy and natural to close servo-loops through the eye-brain-hand-knob path.

The button and switch controls of the IPPS and ISU displays can now be emulated with the mouse-picked button and menu technology of the window systems of the workstations. The mouse is only a linear analog input device. Boxes with nine analog knobs are available at modest cost from both MSWS vendors as optional equipment and we will procure these to supplement the window panels.

6. Stereo Imaging Technology

Stereoscopic perspective display of translucent data structures with full, free roaming is probably the most effective technique for visualizing complicated 3-D data. A human needs to move the viewpoint at will in order to change perspective to enable the stereo vision "hardware" in the visual cortex to synthesize 3-D structures (sight is in the cortex, not on the retinas). Depth and position are inferred from a complex mixture of factors. Geometric cues include the relative sizes of objects, their parallactic displacements (binocular stereo vision) and the obscuring (and exposure) of one object by another as the viewpoint is changed. Illumination and transparency of solid objects can also provide valuable depth and position cues: if objects in a scene are illuminated by a foreground light source, the distant objects appear darker, and distant objects can be rendered in distorted or attenuated form through transparent foreground objects (translucence, refraction).

Stereo imaging is an option available from both MSWS vendors. The workstation needs to have ample bitplanes, configured as two independent frame buffers. The stereo hardware utilizes a liquid-crystal polarizing screen in front of the video monitor plus passive polarizing glasses. The screen passes either left or right circularly polarized light depending on the polarity of applied voltage, and the glasses have L and R polarized filters in the two lenses. Electrical circuitry switches the polarity on the polarizing screen on successive video refresh cycles, and software switches the display from one frame buffer to the other simultaneously, so that the left eye sees one frame buffer and the right eye sees the other. Refresh frequency is more than 60 Hz and so each eye gets a refresh rate of more than 30 Hz. There are no clumsy cables attached to the glasses, and any number of people can view the display simultaneously if glasses are available. The MSWS polygon rendering engines only need to render the scene twice, once into each buffer, with slightly different viewpoints corresponding to the left and right eyes of the people. The polarizing filters are colorless, and so if enough bitplanes are available, the stereo perspective scene can appear in full color. We intend to procure MSWSs with at least 48 planes (two buffers with three frames -- red, green, blue -- of 8 bits each) to permit full-color stereo operation.

Position and orientation controls (knobs) on the control panel will be used to control the viewpoint for rendering operations. 3-D graphics input devices are available and will be examined for suitability in this application, although knobs will surely work. The question of whether panel knobs should control both image enhancement functions and perspective viewpoint coordinates (control multiplexing to save hardware) is a design trade-off which will be examined.

Should structures in scientific data always be rendered as opaque objects? Scientists have very little experience with this technology, and so we can only speculate here. Opaque surfaces, such as signal level contours, can be defined in data and can be given surface texture and reflectiveness, illuminated by light sources, in order to utilize the software and hardware rendering technology of CAD displays. The signal level of the contours can be controlled with knobs to enable the viewer to explore this parameter of the space. This approach will surely work, and will be a major advance, but there is an alternate approach, translucent-object rendering, which may be more suitable for some scientific data. The idea is a natural one for astronomers, who are often observing sources which have optical depth, either in emission or absorption, or both. It may be that such renderings will be too complicated to properly perceive; it may also be that they will convey a combination of spatial and signal level relationships that will be harder to perceive in the opaque object rendering. Probably translucent rendering will be slower to execute (distant objects which would be obscured now must be fully rendered) and existing hardware and software may need to be extended to support the mode. This may not be practical to do, but the question will be explored in the course of this project.

7. Networking

A major strategic advantage of workstations as image displays is that the WS-to-server interface is via "protocols", not hardware. Of course physical hardware moves the bits, but programmers do not interface applications like AIPS to this hardware --- they interface to the software protocols. This means that new communications hardware can be substituted and the applications interfaces will remain invariant. Also, the protocols do not distinguish between local and distant access. At first we will only use relatively low performance LANs (e.g., Ethernet) but, as communications technology evolves in the direction of increasing available bandwidth and decreasing costs, WS-based AIPS implementations will gradually operate at higher and higher speeds and will be removed further and further from the compute servers; this decentralization is sure to be a convenience for everyone. It is unlikely to happen with direct-connected peripherals because I/O busses always have severe distance limitations. A strategic goal for both GSFC and NRAO is to use supercomputers as compute servers to process worst-case imaging problems, and workstations are the ultimate mechanism for making this server power available where the users are located as communications capability evolves. The protocol-based WS strategy is the optimum way to track the evolution of communications technology.

FDDI (Fiber Distributed Data Interface) fiber optic LANs with approximately 100M b/s peak performance (10x faster than Ethernet) are expected to become available by sometime in 1989. This technology is badly needed for the image processing operations of high-performance workstations in association with even higher performance compute servers (i.e., supercomputers). An important technical point is that the FDDI will be well-matched to the performance of present-day striped disks: both have peak rates of 8 MB/s, and FDDI will probably sustain about half that in practice, about like the disks. This means that display techniques for roaming large format and 3-D imagery can be developed with disks and are likely to work almost as well later with fibers, in a distributed environment. NRAO expects that Convex Computer will offer early support for FDDI, and therefore it will be possible to integrate and test the high-performance distributed image processing environment using the Convex C-1 in the Charlottesville laboratory.

A fact which might become important in the future is that the new routers for the NSFnet are being engineered to support 45 Mb/s trunk lines, a data rate comparable to that of FDDI LANs. If this technology works out, and if the economics and funding levels permit the installation of adequate trunk capacity, MSWS-plus-supercomputer distributed image processing technology may become relevant in a wide-area context.

8. Canonical Architecture and Open Systems

The great dream of designers of distributed systems is that the data would flow transparently, without explicit user intervention, from one environment to the other. The simplest concept which leads to this result is that the disk drives of each AIPS system might be "mounted" on the other through a network. This concept only makes sense if the operating systems, compilers and hardware architectures of the two machines are nearly identical; diversity of design in today's computer market often inhibits the free sharing of data between computer systems. Most CPUs designed in the 80's are byte-addressing, non-byte-swapped, twos-complement machines using the ANSI-standard ASCII character set and the IEEE-standard floating point format. We will refer to this combination as the "canonical architecture". The Fortran standard specifies how variables in COMMON blocks may be EQUIVALENCEd; if the modulo-addressing problem is avoided, as AIPS does by convention, this leads to predictable, machine-independent data structures on machines with the canonical architecture. Because of this fact, an AIPS on such a canonical machine can generally read AIPS binary file structures produced by another such machine with no format conversion.

Many of these machines also support Network File System (NFS), a shared-disk standard which has been promoted by Sun Microsystems. NFS allows data on the disks of either the workstation or the compute server to be automatically accessible to the AIPS executing on the other machine simply by declaring the disks of each machine to also be virtual disks on the other machine. There is a performance penalty due to accessing disks across a LAN, but the gain in convenience and simplicity offsets this. The full benefit of shared disks cannot be achieved until the fiber optic LANs become available in 1989-90.

The fact that file migration in a multi-vendor environment will work best, with minimal coding effort, and most efficiently, for systems with canonical architecture and NFS, means that it may soon become common to write NFS and binary format requirements into specifications for "requests for quotes." In this regard, it is important to note that the traditional supercomputer vendors all have proprietary hardware architectures, although all of them are moving to UNIX as their OS plus NFS for file sharing.

IV. TECHNICAL ASPECTS OF MSWS SOFTWARE DEVELOPMENT

1. Portable AIPS Software

The most fundamental strategic fact about this proposal is that it is based on the use of a proven, already existing, software environment. The proven environment acts as a filter which prevents investment in "off-beat" hardware/software environments: if it won't run AIPS it probably will have trouble with other existing engineering-scientific applications, and therefore may not be commercially viable. The AIPS strategy also "filters" software design by supplying a set of proven design guideline principles and by tending to prevent investment in excessively grandiose software plans: almost any practical software architectural concept can be implemented as a sequence of staged incremental enhancements to AIPS, each of which has predictable incremental cost bounds and delivers predictable incremental scientific payoff in a predictable time frame. Furthermore, the AIPS strategy constrains such enhancements to be made with an eye to backwards compatibility, which preserves the investment in AIPS user training. All of these reasons together imply that the AIPS strategy of simultaneously tracking hardware and software evolution tends to have the lowest costs and greatest predictability.

From the user's point of view (ref 9), AIPS is a command language processor plus a set of application programs, all working together with the peripheral equipment, especially the digital image display, to deliver a high-performance digital image processing total system. A command language processor (called POPS, People-Oriented Parsing System) is used, rather than the host operating system command language, in order that AIPS users will see the same command language in all implementations of AIPS and thereby minimize "finger re-training" costs. This simple principle is very important -- it permits AIPS to easily attract users to hosts which would otherwise deter users because of unfamiliar operating environments. A related design feature is that AIPS imposes its own (portable) file naming convention.

AIPS is portable; its command language processor and application tasks are coded in a portable dialect of Fortran, and the interfaces between these components and the host environment and its peripherals are all encapsulated in layered software interface libraries which can be adapted to environmental and device peculiarities. Indeed, much of the basis for this whole proposal is that attractive features of new hardware and software environments (windows, concurrent vector hardware, distributed computing, etc.) can be easily grafted onto AIPS by simply replacing or modifying various layered interfaces, while leaving the vast body of portable AIPS application code invariant.

The intent of the layered interfaces is that peripheral devices should be regarded by application code as being "virtual" devices. For example, code in the command language process and application tasks of AIPS which interact with the digital image display (the "TV" in AIPS parlance) is not coded for any particular TV model. Rather, it is coded for a parameterized virtual TV with a wide range of potential capabilities, the full set of which are implemented in only a few, if any, real devices. When the code "opens" a device, it receives the description of that device from the interface layer, called AIPS "Y-routines" because the subroutine names all begin with Y (ref 10); it adapts to the various features, attempting to exploit them as fully as possible.

Thus, the designer of a window-based virtual TV for use with AIPS has considerable freedom to vary the dimensionality of the device (bitplanes per frame, pixel dimensions, number of image frames, etc.) within the wide latitude set by the AIPS Y-routine model. The model can be, and has been, extended to include new functionalities. The present project may require the addition of stereo imaging and/or polygon rendering as (optional) imaging modes.

The portable code of AIPS already constitutes a proven astronomical image processing capability. This project will only need to add some new application modules to exploit new capabilities which are available in the new MSWS hardware and software.

A large community of astronomers (ref. 11) are already familiar with AIPS and the COBE project has selected this system for its science analysis function based on the pleasant experience of many COBE project members who previously participated in the successful IRAS/AIPS efforts at GSFC (see Fig. 4). This implies that an experienced user community is immediately ready to exploit the new capabilities which the MSWS project will provide. Changes and additions to AIPS for the MSWS project will be done with an eye to maximizing the continuity of the AIPS traditions precisely in order to facilitate the use of the new systems by existing AIPS users.

2. Machine Independence

NRAO has maintained a policy of confining machine- or vendor-dependencies to the layered interfaces, avoiding them in the large body of portable code. Portable code which depends on machine constants, such as bits per word or bytes per disk sector, gets these constants from COMMON blocks which are initialized by the machine dependent interface layer. These policies will be maintained.

Even with such "magical" hardware as a MSWS, a policy of machine independence will be maintained as much as possible. This will be done by defining new additions to the Y-routines and other interfaces and parameterizing them. The motivation for this is that, if possible, we want new imaging and rendering code to run on more than one MSWS -- i.e., we want to create a "second-sourcing" situation in order to promote competitive procurement.

3. Remote TV Operation

A key component of remote-AIPS operation is that we must support remote-TV operation. The required software for this function is called the "Virtual TV" (VTV) in the AIPS Project. The VTV depends on a virtual circuit connecting the remote AIPS to the local AIPS; the current implementation uses the Berkeley socket interface library and the TCP/IP protocol suite (other protocols could be substituted). At the remote AIPS, the VTV is simply the Y-routines. Calls to the Y-routines are encoded as name-plus-arguments in a machine-independent fashion (using canonical data formats) and are transmitted across the virtual circuit to the local AIPS where the arguments are decoded from canonical format to the local (host-dependent) format. On the local

AIPS, the VTV software calls the local Y-routines, thereby connecting the remote AIPS to the local TV. Whenever the remote AIPS attempts to open the TV, the local Y-routines supply the parameterized description of the local TV to the VTV layer, which encodes the description and passes it back across the circuit to the remote AIPS, where it appears in that environment and is used. This technology already exists and works.

Note that if the local machine is a workstation, then the Y-routines create a TV out of the functionality of the window system of the WS. If the local machine has only one TV (the usual case) it would be shared between the remote AIPS and the local AIPS. This sharing is analogous to the usual AIPS situation in which multiple users and multiple processes of a single user all share access to a single TV. The ability to use the verbs of the local AIPS command language to manipulate imagery downloaded by the remote AIPS is a natural consequence of this architectural model: no new AIPS verbs will be needed and users will need no new training.

The AIPS TV design includes the concept of an image catalog associated with the TV. This allows a display frame buffer to contain an irregular mosaic of images and the AIPS command language and application tasks are able to determine at all times which pixel in which disk file corresponds to each pixel of the framebuffer. Full non-linear coordinate descriptions and physical unit transformations are also maintained in the catalog in order to support interactive cursor-pointing operations over the mosaic. The filename, coordinate and units information is passed through the Y-routine interface; therefore, it passes through the VTV to the local workstation TV, where it is available. Thus, the distributed AIPS model not only includes support for simple image data but also for the full physical description of the data! This technology is already available because it is inherent in the basic AIPS architecture plus the VTV model; a workstation AIPS can point at an object in a scene downloaded by a compute server and can immediately read out coordinates and physical units for the pixel values. AIPS can do this even for cases in which the display contains an image whose original data has been deleted from the disk -- subroutine library reads the (less accurate) pixel data from the display and applies the appropriate scaling which is also carried in the image catalog. This feature is particularly important for remote AIPS operations in which the remote data files are not visible by means of NFS because it will imply that precision pointing and readout operations will still be supported.

The window-based TVs must be constructed in such a way that they have as much portability as possible. Use of industry standards is the key here, and that means "X-windows" and "NeWS" (or "Display Postscript"). These are all "emerging" standards, and the further development of workstation-based TVs will need to track the evolution of this technology.

4. Control Panel Design Features

The provision of control panel functionality is not enough; it is also necessary to find a strategy which will permit such advanced capability to be incorporated into AIPS such that it is an optional function, one which co-exists alongside the traditional AIPS display control capabilities (trackball with buttons plus POPS verbs). This specification assures that AIPS users can control the WS-based TVs using the skills which they already know, learning to use the control panels and knobs selectively and progressively. Furthermore, the capability must be machine-independent and, if possible, capable of being retro-fitted into existing mini-computer-based AIPSs.

The strategy which meets all of these objectives is yet another application of traditional AIPS design principles: the code which implements the control panel functions will be yet another AIPS task (or stand-alone program) which will interface to the TV using the standard AIPS Y-routine interface. That is, it will translate knob and mouse-picked display commands into calls to the Y-routines which will accomplish the desired operations. No changes to the AIPS display architecture (the Y-routines) or to the POPS command language (the TV verbs) are required in order to implement this plan. It only assumes that multiple tasks can contend for access to the TV, a capability which AIPS has always had.

The control panel task will be coded to AIPS standards and will use standard WS window interfaces (almost certainly X-windows) to implement the menu and mouse portions of the control panel emulation. It will talk to the control panel through a standard (RS-232, ASCII) terminal port. Thus, this task will be inherently portable to all WS environments (not just MSWSs) that will be marketed in the 1989 time period.

It is even possible that a terminal-based emulation of the control panel could be constructed, as a version of the panel task which interfaced to an ANSI-standard terminal plus the control panel. Presumably, much code could be shared between the window-based and terminal-based versions of the panel task. While the WS, especially MSWS, version is much more important strategically, and is the main goal of this project, we also know that dozens of fine image displays are already installed (many at NASA sites) and improving their functionality is a worthwhile goal too. We will consider this option during the course of the project.

5. Wide-Field Image Visualization

Digital image processing system designers have always had to cope with the problem of how to display imagery which has a larger format than the size of their physical display, and there is no sign at all that this situation will ever change. Two standard techniques are that the user can either view a designated subimage at full resolution or else can order the decimation of the data in order to view a wider area at lower resolution. A third classical technique is that some digital image displays have a "roam" mode in which, for example, a 512-square window roams in a 1024-square image formed by four 512-square image planes. Even in the latter case, only a factor of two is gained, and so there is always imagery which requires the use of the first two techniques.

The widest fields in VLA imagery are commonly 4096-square; a few cases which need 8192 are often mentioned. The disparity with the 512-square displays which are common today is so gross as to be ludicrous. The operational fact is that such imagery simply cannot be viewed effectively at full resolution with present hardware. Displays of 1024-square are beginning to appear, and they make a big difference, but disparities of four in linear dimension and 16 in area still exist. If displays could be rapidly reloaded by the host computer, this problem might be endurable but, in practice, image display interfaces do not attain transfer rates which enable us to surmount the hardware limitations. Radically new solutions are needed if the high resolution capability of the telescope is to be made fully available to astronomers. We believe that the MSWSs will provide these solutions, based on the intimate coupling of the powerful general purpose computer to the image display.

The computer can compute multi-resolution roaming windows into the data, and can maintain these using the window mechanisms provided in the workstation environment. The concept is to provide one or more small "finder" windows showing the location of the high-resolution window in the larger space. The zoom factor of the main viewing window will be instantly changeable using a knob on the control panel, while the coordinates of the view window will be continuously changeable using two knobs on the panel. As an example, consider a 1280 x 1024 frame buffer displaying a 4096-square image. Suppose the software computes a 256-square image block-averaged by 16 from the original and displays it in the corner of the screen. There will still be an unobstructed area 1024-square at full resolution. A box 64-square can be drawn in a graphics overlay on top of the 256-square finder window to show the area being displayed at full resolution. If the user zooms up the main window (block replication) the box would decrease in size appropriately. "Negative" zoom factors would also be implemented (1/2,1/4,1/8th resolution); very few image displays today support this important capability. This combination of techniques (low-resolution finder windows associated with the roaming window, combined with both positive and negative zooming) will enable the user to perceive the high resolution window in the context of the entire field of view.

The computer must be able to reload the display with new data as rapidly as the user wants to roam. The action must be smooth---jerky motions would distract the user from his objective of visualizing the image. The resources of the MSWS will make this easy. First, several times more RAM will be available than the size of the frame buffer, and the RAM-to-frame-buffer copy rate will be exceedingly high, probably capable of reloading the entire buffer in a time comparable to the vertical refresh interval (15-20 msec in modern displays). In fact, for a roam operation, most of the pixels remain the same, but are moved to new locations in the frame buffer. It is standard to have high-speed block move operations in present-day workstations in order to implement this very common operation. It is also standard to have two frame buffers so that the next video refresh can be computed and loaded while the current buffer is being painted on the phosphor. Finally, the RAM will be backed by magnetic disks which will transfer to RAM using the striping technique. It is nearly certain that transfer rates in excess of two megabytes per second will be achieved, which will permit loading the buffer in less than a half second from a random location. If, as would be normal, the user roams slowly, the software can continuously load new data into main RAM ahead of the roaming window. In fact it may be practical to just let the standard virtual memory management mechanisms accomplish this function automatically.

6. 3-D "Cube" Image Visualization

The visualization of spectral-line cubes produced by astronomical telescopes is one of the most challenging problems of image processing. Two approaches have been used in the past: (a) roaming in the Z-axis (frequency), especially using the "movie-loop" technique and (b) producing perspective views of surfaces of signal contours in the cube. It appears that the imaging MSWS will enable several new approaches to be explored easily: (a) multi-resolution roam and movie-loop, especially with negative zoom factors, (b) "translucent" perspective rendering of multiple signal contours, and (c) true stereo imaging of 3-D cubes of data.

The movie-loop technique was pioneered in radio astronomy by the GIPSY (Groningen Image Processing SYstem) project around 1978. An instant-replay disk was used to store analog copies of images displayed on a digital image display. Spectral line cubes could be explored rapidly once they were recorded. The success of this project has had a great influence on all subsequent approaches to the cube problem. In 1981, NRAO initiated the ISU project mentioned earlier. As with the GIPSY approach, the concept was to record images of planes of the cube from the digital image display, but this time the disk was digital rather than analog. Some reliability problems had been experienced with the analog disk technology and it was hoped that digital disks would be more easily maintained. The other reason for use of digital image storage was that it permitted more flexibility: digital images reloaded into the digital display can be enhanced and zoomed in new ways at will, whereas the analog images on the GIPSY disk cannot be altered once they are recorded---they can only be displayed in any order at any speed. In fact, the ISU has achieved its reliability and flexibility goals, but has sacrificed speed to get these advantages. A major advantage of the MSWSs is that they will increase both the speed and flexibility of a digital cube display by using increased computer power to substitute for the raw speed of analog disks in movie looping.

First, we note that digital disks with speed comparable to analog disks are available today, and that they have been interfaced to digital image displays. The current notable example (ref 12) is the display system at Los Alamos National Laboratory (LANL), which has several hundred megabytes of parallel-track digital disks attached directly to a large Gould-DeAnza image display, all under the control of high-performance Gould super-minicomputer which has a high-performance communications link to a Cray I/O channel. This combination offers the speed of the GIPSY display and, simultaneously, the flexibility of the NRAO ISU. If the imaging MSWSs did not and could not exist, then some variation on the LANL theme would surely be the optimum path. The imaging MSWS, with its frame buffer even more tightly coupled to its CPU, and with a CPU several times more powerful than in the LANL system, can exploit and explore compute-intensive image visualization techniques more effectively, while its use of conventional mass-manufactured disks reduces cost considerably, which can permit more copies of the display system to be purchased.

Two techniques will exploit the increased power of the imaging MSWS: (a) use of the large main memory as a cache for movie looping, and (b) computation and display of negative-zoom averaged images of cubes. Both concepts depend on the fact that users really do not need to view huge cubes of data at high frame rates. Rather, they need rapid overviews of the large cubes combined with an ability to zoom in on limited regions at full resolution. The compute performance of the imaging MSWS means that averaged versions of the cube at half, quarter, eighth, etc., Z-axis resolution can be computed and stored, as a cube is being loaded into the display. A Z-axis zoom knob on the control panel (in a window of the WS, of course) will permit the user to select which resolution is being displayed. This flexibility will enable the user to rapidly explore large cubes looking for interesting features, even though the disk transfer rates may permit only 3-4 frames/s of 1280x1024 size (assuming that the imaging MSWS will achieve only the 3.6 MB/s rate that the four-way striped "Eagles" of the Convex C-1 achieve). For 512-square imagery and a negative zoom factor of 4, such rates will permit a movie loop rate of about 50 frames/s referred to the full-resolution cube (the display frame change rate would be 4x lower, about 12 frames/s); a 512-channel cube would pass through the display in about 10 seconds. Note that the averaging technique not only increases the effective movie-loop rate but also improves signal-to-noise ratio for signals that extend over several spectral channels. Ability to view the cube at several different spectral resolutions in rapid succession will give optimum detection sensitivity for spectral features with a variety of widths.

The concept of the movie-loop cache is that users cannot use arbitrarily high movie frame rates and arbitrarily long movie loops effectively. The frame rate limit is set by human response times for detection while the length limit is set by human short term memory. These psychophysical parameters are subject to endless debate and probably to personal differences; we assume 8 frames/s and 4 seconds as working values. For a 1280x1024 pixel display these parameters imply a need for at least 32 MB of RAM; 64 MB would be desirable. For example, a 1024-square cube of 256 spectral channels averaged down by 4x would fit into 64 MB of RAM, enabling arbitrary roaming at arbitrary rates. The RAM would need about 20 seconds for initial loading, but this is not really a problem: it is more likely that the user would instead choose 8x negative zoom, select a region of interest, and then zoom up and loop over a more limited range. In fact, the real purpose of large RAM is not so much to permit very large cubes as it is to permit cubes of multiple resolutions to be simultaneously available for selection by control panel knob. Furthermore, users will probably prefer to roam slowly through cubes, varying resolution up and down at will, when they are exploring.

The concept of the "finder window," which was discussed earlier in the context of the visualization of wide-field imagery, is likely to be very important for spectral-cube visualization as well. The idea is that multiple windows in the frame buffer can simultaneously display the cube in multiple resolutions and orientations; this will allow the available RAM to be used to vastly open up the field of view. For example, instead of maintaining full-resolution 1024-square images in RAM, we can keep 512-square, gaining a factor of four. The full-field overview will still be available in a 256-square finder window, within which we can roam the high resolution window. Two other windows adjacent to the X and Y axes of the 512-square high resolution window can display X-Z and Y-Z views of the cube, using the planes selected by the (X,Y) coordinates of the cursor. Note that the transposed views will be at the Z resolution selected by the Z-zoom control. The combination of X,Y,Z-roaming and Z-movie-looping with the wide-field finder window and transposed views will give users unique new tools for cube visualization. These new capabilities depend on two features of the imaging MSWS: its large RAM, which is effectively image display RAM, and its flexible multi-window capability. Although a conventional display, such as the LANL system mentioned earlier, can contain and exploit a large RAM, it cannot utilize the multi-window technique which gives the finder-window and transposed-view abilities. Only a true workstation can do that!

7. Advanced Image Enhancement Techniques

Displays that use only trackball and a few buttons are forced to multiplex the controls. In contrast to this, the control panel approach, as used in the IPPS panel at NOAO and the ISU panel at NRAO, permits the user to simultaneously manipulate multiple parameters of the most complicated display modes. This flexibility enables users to rapidly synthesize new combinations of the controls which are well-suited to their particular pieces of data. Flexibility in choosing and adjusting enhancement functions is especially important.

Consider pixel operators which are implemented as lookup tables in image displays. The linear "window-and-stretch" contrast enhancement is the bread-and-butter of digital image processing. Since the advent of the IPPS in 1975, this has always been a two-parameter function, because it is almost always implemented using the X and Y coordinates of a trackball pointing device. But, in fact, a three-parameter function is a powerful improvement (curvature, in addition to the basic slope and intercept). NRAO's ISU panel has three knobs per lookup table to easily give this additional function.

The first such technique is "blinking" (also often called "flickering"). Two (or sometimes three) images are repetitively displayed in registration. The eye immediately notices flux differences, alignments, misalignments, etc., in particular features. Closely related to blinking is "split-screen"; the left (or top) half of one image is displayed with the right (or bottom) half of the second image. The dividing line may be moved interactively to position it on features of interest. Split-screen works well in conjunction with flickering by setting up image enhancement functions in split-screen and then switching to blinking mode. The whole process is almost effortless when a control panel is available because the mode changes, split-screen coordinate adjustments and multiple enhancement function settings, and blink rate manipulation are all performed with separate controls. While blinking can use pseudocolor enhancement, it does not depend on it.

Blinking and split-screen are useful only when the two values to be displayed are identical in character (e.g., two related flux measurements). It is when the variables are fundamentally different that we see the requirement for one or more additional independent variables in the display technique, and this means color. Human color vision supports three independent variables. Scientific data presentation works best when only two of the dimensions are used, generally intensity (the single variable of monochrome display) plus the hue of the color. The third variable, the saturation of the hue, can be exploited in special cases where three variables are needed.

Almost all color display hardware uses the RGB system: three frame buffers, generally 8 bits deep each, drive the red, green and blue guns of the color monitor. But there is an alternate color model, called HSI (Hue, Saturation and Intensity). The HSI color model is widely regarded to be a superior description of human color perception, and many image processing specialists would prefer to use it if they could. The barrier has always been that interactive enhancement operations on imagery expressed in HSI form demand real-time computation of the equivalent RGB form to drive the monitor guns. Most current displays are unable to satisfy this need: they lack video processing circuits of the necessary sophistication and they usually do not have enough storage to hold both the HSI and RGB images. The host computers could do the computation, albeit somewhat slowly, but they generally cannot download computed imagery into the displays fast enough. The MSWSs will have the computational capability to do HSI to RGB transformations at high speed and they will have the necessary memory and transfer rates as well. Therefore, this project will explore the potential of HSI enhancement modes for display of multi-parameter scientific data. Polarization imagery is a case which appears to be especially promising.

During the first decade of astronomical image processing, almost all image displays had 8-bit frame buffers. This was generally adequate for the data being analyzed then, but dynamic ranges have steadily increased in much astronomical data. Recently some displays have been configured with 11- or 12-bit frame buffers, which is a considerable convenience for viewing high-dynamic-range data. The concept is not that a monitor can display 10 or more bits of video (1000 grey levels) or that humans can perceive such a range (in fact, human perception is limited to more like 6-bits, or about 64 distinguishable levels). Rather, the concept is that a lookup-table (LUT) maps image data into the frame buffers. The user has control over the LUT(s), and, by rapidly varying the transformation function, the user explores the dynamic range of the data and synthesizes the impression of the data range. This interactive technique has been used by astronomers for the past decade; it is almost never openly discussed, because it is intuitively obvious. The MSWSs, by having vastly more RAM and compute power in and near the physical frame buffer(s) will be able to have logical frame buffers of greater depth. The obvious option is 16-bits, with LUTs of 65,536 entries. LUTs of such size would formerly have been impractical, but now are reasonable. Note that because the "TV" of a MSWS is a software abstraction, built on top of the window system, it will be easy to explore the option of 16-bit logical frame buffers and 65K LUTs. Indeed, one could have the TV depth be a user-selectable parameter, to be specified at the beginning of each session. These possibilities will be explored in this project.

A subset of Y-routines should be certified before procurement and test criteria specified. As performance is likely to be high for all interesting systems, mere time trials of simple operations may not be particularly useful. Probably the most relevant simple speed rating is the transfer rate from main memory to frame buffer memory. If this speed is high enough, almost any imaging/graphics mode can be emulated with software. From this point of view, one only needs to specify a minimum set of capabilities for the physical frame buffers, a set sufficient for the desired emulation, and demand enough transfer rate and CPU power. That is, in principle, the physical frame buffer and associated window software do not have to provide a "real" TV, only a primitive display capability which is "necessary and sufficient". Then the certification issue becomes the problem of deciding the necessary and sufficient set of primitive functions and the corresponding set of Y-routine functions which will demonstrate that the capability exists and will measure relevant performance parameters.

8. Real-time Display of Calculated Imagery

The imaging MSWS is able to compute pixels as fast as the electron beam can write them on the phosphor, and it can load the pixels into the frame buffer as fast as it can compute them. This is an unprecedented situation for an image processing system designer, one which opens up whole new vistas of fascinating and probably scientifically productive applications. For simple algorithms requiring only a few clock cycles per pixel to compute, the entire frame buffer can be continuously updated at nearly the refresh frequency. For more complex algorithms, the designer can trade off the number of pixels and the update frequency. If we take the human servo loop frequency as about 5 Hz, we gain about a factor of eight. If we further accept a size of 512-square, we gain another factor of four. This would allow execution of 30-50 instructions/pixel/update, and these instructions could include arbitrary floating point operations at full precision. This is fast enough for one or maybe two transcendental function executions. A target of 2 Hz and 256-square would allow 500-1000 instructions, permitting sophisticated astrophysical calculations, such as interactive model-fitting operations with parameters under control of control panel knobs. Probably most astronomers would be so astonished at the ability to display such calculations at all that they would cheerfully accept an update frequency an order of magnitude lower (0.2 Hz, 5 seconds).

Consider the problem of real-time interpolation of imagery. Suppose we accept the traditional bi-linear interpolation. Then we will perform 4 multiplies and 4 adds per pixel (for arbitrary coefficient values), and with a suitable arrangement of the code, the operations will fully occupy the pipelines. If we postulate 40 MFLOPS delivered, then we will be able to interpolate at a rate of 5 million pixels per second, assuming that we do not require the recomputation of coefficients. Interpolation of the 1280x1024 frame buffer should be feasible in under 0.3 second. Coefficient computation can be avoided by allowing only a finite number of scaling ratios and zero-point shifts, specifically a set chosen from the rational numbers. The concept is that the high resolution roaming window can have an almost arbitrary zoom factor; rather than just 1, 2, 4, 8 plus 1/2, 1/4, 1/8, it can also support 1.25 (5/4), 1.33 (4/3), 1.5 (3/2), 0.8 (4/5), 0.75 (3/4), and 0.67 (2/3), plus a host of other rational number combinations, selected by a continuously variable zoom knob on the control panel. This technique has an important implication for the visualization of large-format imagery: image structures, which do not naturally come with sizes which are neat powers of two, can be rapidly interpolated to take full advantage of the format of the display hardware.

Another interesting feature which depends on powerful computing capability intimately associated with the image display is the notion of having a "resolution" knob on the control panel. This can be implemented with a real-time convolution operation. For example, an interesting class of low-pass filters can be synthesized from 7x7 convolving kernels. The simplest implementation is a 49-point gather chained to a weighted sum (the dot-product reduction operator), which a 40 MFLOP pipeline can execute at about 500K pixels/second, permitting an update frequency of about 0.4-0.5 Hz for the 1280x1024 display. Again, most users, for most applications, would cheerfully accept a much lower update frequency, certainly as low as 0.2 Hz, which would permit a larger kernel (e.g., 9x9).

An important application of the display of calculated imagery occurs for 3-D multi-spectral cubes. Often there is a "baseline" signal present in all spectral channels which obscures faint spectral emission signals. Such constant backgrounds are eliminated by subtraction of a reference channel of the cube. Indeed, it can be convenient to view a spectral channel while varying the choice of reference channel which is being subtracted, with the goal of maximizing the spectral contrast and minimizing the background artifacts. An alternate case is spectral absorption seen against an extended continuum background source. In this context, the appropriate function to display is the ratio of two channels. The ability to rapidly view the difference or ratio of any two channels of a multi-spectral cube depends on three properties of the MSWS that we are proposing: random access to enough RAM to hold a large portion of a cube, adequate computational speed and the control panel with appropriate knobs for the channel selections. Of course, panel control of enhancement and viewport into the channel difference image is vital.

During the past decade, many astronomers have asked for a capability to interactively register images. This requires interpolation of one image to match another, and blinking, subtraction or dividing to view residual mis-registration as the interpolation parameters are interactively varied. Expanding nebulae are a possible astrophysical application for this type of display capability: an image at one epoch can be interactively varied to fit an image at another epoch to examine the motion of structures, which are likely to be varying as well as moving (human judgement is required to untangle the situation). The ability to optimize image enhancement plus zoom and roam viewport in order to view the residual patterns is vital! The MSWS will be capable of performing this complicated operation gracefully because of its combination of compute power and control panel knobs for control of offset, scale and rotation for the interpolation, plus the panel control of enhancement and viewport.

9. Interactive Model-fitting

A major new capability which we expect will be of great importance to the COBE project data analysis effort is the interactive fitting of models to imagery. The goal of the COBE mission is to produce parameterized models of the planetary and interstellar media and cosmic background emissions which will fit the COBE multi-spectral data satisfactorily. Display of observed-minus-computed residual images is likely to be the best indicator of adequacy of automatically-fitted models. It is also likely that ability to adjust model parameters while viewing residual images will enable an analyst to improve the model by using human judgement to interpret patterns of residuals. This process will also give COBE analysts an intuitive feel for the sensitivity of the parameter fitting process, and will maximize the likelihood that any new, or inadequately modelled, phenomena will be noticed in the data. This process will be computationally intensive, as well as having demanding display requirements; the idea was nearly inconceivable until the MSWSs became available!

This model fitting will be COBE-specific application code and Figures 1, 2, and 3 help illustrate the concept. NRAO will assist by providing a framework, presumably in the AIPS context, for the code to be created and maintained, and by providing software mechanisms for utilizing the knob panel for parameter control.

All areas of observational science involve fitting models to data, and so the demonstration of interactive capability in a complicated imaging context is likely to stimulate other groups to also exploit the MSWS technology. In astronomy, a likely application will be the fitting of models to polarized radio imagery and Hubble Space Telescope optical imagery of extragalactic plasma jets.

10. NASA/NRAO Networking Interface

NRAO is two point-to-point hops from GSFC via NSI and NSFnet: currently a T1 (1.5 Mb/s) circuit connects the NASA-GSFC LAN to SURAnet Headquarters at U. MD, and a 56 Kb/s circuit connects U. MD to the U. VA LAN; NRAO's Charlottesville laboratory is on the campus of U. VA and has full LAN service to the U. VA SURAnet gateway. Experiments show that effective transfer rates up to about 40 Kb/s are available, when contention is not present, and is adequate to promote the shared development of this technology by the GSFC and NRAO groups.

11. Data Porting Methods

At present, it appears that the canonical-architecture-plus-NFS strategy is the most attractive approach to distributed computing, but we must also address the cases that don't fit this model, especially situations involving computers procured in the past. Astronomers can always use FITS tapes to port data between computers of diverse architecture, but smooth, painless migration of files between two AIPSs is essential if the workstation-plus-compute-server strategy is to make sense in our applications. This capability will evolve in a series of steps.

a) FITS Tactics

The first step will be to utilize tasks FITTP (the AIPS FITS-format writer, see refs 9 and 10) and IMLOD (the FITS reader, see refs 9 and 10) to obtain data transfer between machines of dissimilar architecture, but with the network as the transport medium rather than magnetic tape. Initially we will write FITS files to disk, then order them to be copied from one disk to another, and finally order the second machine to read them from its disk. This capability is already present in AIPS today. We may eventually construct special versions of these tasks which will be able to cooperate: FITTP in one machine sending its output directly to the input of IMLOD in the other machine, in real time. Such tasks will be invoked by the user in the usual AIPS fashion and will run in the background. While this design has the disadvantages that it requires user intervention and consumes significant computer resources, it has the great advantages that it is sure to work with relatively modest effort and will automatically transfer all types of AIPS files which FITTP can convey via magnetic tape.

b) Client/Server Tactics

The COBE database will exist on a VMS machine in a VAX-specific binary format. How can the MSWS gain access to these data, and cope with the data format differences? Two approaches appear to be possible, each with advantages and disadvantages. First, several NFS server implementations exist for VMS already, and other vendors are rumored to be nearly ready for product introductions. This implies that the VMS-resident data files could simply be mounted by the MSWS, with the VAX merely acting as an NFS file server. The VMS-specific file formats would then be "cracked" by ad hoc code on the MSWS, accounting for binary format differences as well as somewhat poorly documented file structures. This may in fact be the highest real-time performance approach, as the vector-concurrent architecture of the MSWS has substantially more computing power available for format conversion than does the VAX.

An alternate approach is that a special server process can be created on the VAX which will accept coded requests from the MSWS across the net, access the VMS-resident database, reformat it into an agreed canonical data format, and send the results to the client process on the MSWS. Because the database access would utilize existing VMS library code, it would avoid the problem of understanding the internal file structures, thereby probably minimizing software development difficulty, although adding some difficulty in the form of the server process itself.

At present, it is not possible to say which of these approaches would be best. The final decision will likely depend on exact details of available manpower, available NFS implementations, exact details of COBE database structures and access patterns, etc.

V. EXPECTED BENEFITS

1. State-of-the-Art Technology Development Driver

This proposal to develop MSWS to provide interactive computer model fitting to observed space data will serve as a state-of-the-art technology development driver for various projects now under the purview of NASA/GSFC Code 630. This will naturally occur due to the proposal's emphasis on interactive modelling, real time visual analysis, and 3-D display. Future space mission projects will demand such data visualization schemes, new techniques for handling large volumes of data and performing CPU intensive computations as a cost-effective alternative to supercomputer resources. Thus, as a result of this proposal, NASA will be able to seed a national effort to develop MSWS for scientific investigation in both astrophysics and earth sciences. This concept is analogous to the national computer networking arena where NASA's Space Physics Analysis Network (SPAN) provided the needed impetus to create an effective national science analysis network to facilitate research.

From a purely commercial viewpoint, this project will probably have a significant impact: the resulting integrated AIPS system will immediately have an extraordinary price-performance-feature advantage for astronomical image processing, and astronomers world-wide are eager to use AIPS processing power! For example, within one year after NRAO procured an advanced vector computer for its Charlottesville Image Processing Laboratory, seven more vector machines were procured by astronomical institutions; AIPS is executing on at least six of the eight machines, and NRAO's first procurement probably caused at least five of the seven to be vector machines rather than scalar. This historical example has two further relevant lessons: (a) NRAO conducted a competitive procurement which fully qualified two vendors using the certification test and published the results; as a result, the first eight procurements were evenly balanced between the two vendors (four and four). (b) Six of the eight procurements were by foreign institutions. We can conclude that "second-sourcing" is good for the market, and probably also good for astronomy, and we can also conclude that national center R & D can have a positive effect on U.S. foreign trade.

The new advanced workstations will have a relatively low price and, therefore, the total number of procurements will surely be larger than for the large-scale vector computers in 1986. The likely outcome of these probable procurements is that during the course of this project, several institutions will be collaborating with NRAO and GSFC on image visualization research using MSWSs and that several of the workstations will be located in the neighborhood of supercomputers, ready to participate in the further development of a distributed AIPS software environment, which is the proper long-term role for these workstations.

2. Astrophysics

Astrophysics research depends not only on state-of-the-art instruments to gather the data but also on advanced tools (both hardware and software) to analyze this data. The NASA-sponsored Data System Workshop held in August 1987 (ref 13) defined the concepts for an astrophysics data system that can be used throughout the community for various astrophysics missions. The MSWS collaborative effort with NRAO will ensure that many of the goals outlined by the Workshop will be addressed. AIPS runs on about 90% of computers used in the astrophysics community presently. In the future, machines such as the MSWS will start to populate this community and, therefore, will need a data system such as AIPS to be ported to them. By setting up this collaborative effort with the original developers of AIPS, NASA will ensure that a truly portable data system will be available to the outside community for this new generation of machines. Code 630 personnel will benefit from this collaboration by understanding the problems inherent in porting data systems to MSWSs and therefore aid other groups in porting such systems as IRAF (Image Reduction and Analysis Facility) to these machines. This NASA-NRAO interaction will also benefit the data system development by promoting the active participation of astrophysicists in developing new analysis techniques that can be used on the MSWS and hence be incorporated in AIPS, IRAF, etc. These ideas are in keeping with the Workshop goals of developing a few transportable operating systems to run on a wide range of machines including the MSWS and involving astrophysicists to help develop new analysis techniques.

Another area identified by the Workshop was algorithmic research. The report stated that NASA should encourage the development of advanced data analysis techniques and encourage the exchange of techniques across project boundaries and from outside of the agency. The MSWS will be a natural catalyst for developing new astronomical software by in-house personnel from various projects such as COBE, ROSAT, etc., as well as outside astronomers. Due to the nature of this proposal, outside personnel are already being involved. The MSWS will also help in exploiting the full information content of various data sets. One problem is removing the instrumental response (spatial and spectral) present in these sets which requires having the right deconvolution techniques and raw CPU power. The MSWS, with its compute power and unique image rendering capabilities (real time model generation), will enable researchers to experiment with new techniques and help them refine previously established ones. The net result will be that better inversion methods will be developed and thus more believable deconvolved data sets will be produced.

The Data Workshop also addressed the problem of accessibility of supercomputers by observational astronomers. In the past, these machines were only used by theoreticians (ref 14) for modelling and not for data reduction. Even though the MSWS doesn't have the CPU power of a large supercomputer, it does have enough compute power to open new vistas for data analysis and can be the logical stepping stone between the current VAX-like machines and the supercomputers. The MSWS uses a UNIX operating system, making it much easier to program than the conventional supercomputers, and already has a built-in image display device. This will allow the porting of software systems such as IRAF and AIPS and enable the user to perform interactive examination of data sets. Also, these MSWSs can be linked to large supercomputers through TCP/IP network protocols for intensive CPU operations.

3. Future GSFC Projects

In the era of NASA's Great Observatories there will be a demand for a low-cost, high-speed graphics engine so that mission guest researchers at small institutions can compete with colleagues at facilities with vast computational resources. The proposed MSWS effort can be used to begin this process which will help develop innovative new ways to visualize results from various space missions and ancillary computations performed by supercomputers. There exist at least three levels of complexity for display of this information (ref 15) The first one involves the least amount of processing and is the standard display most users are familiar with. This includes simple plots, pseudo-coloring of the data, and simple image processing. The second level is associated with higher-resolution data sets that originate either from high-data-rate space instruments or from detailed computations. These data sets are multi-dimensional and have a very high information content. A MSWS could be used to develop new techniques to display these data sets, such as slices through a multi-dimensional surface, etc. The MSWS's liquid crystal display (LCD) could also be used to effect stereo images of data by blinking on and off different polarization states of the LCD. For example, this would be applicable to data obtained from the synthetic aperture radar (SAR) instrumentation which will be aboard the Japanese Earth Resources Satellite (JERS). Viewing would be accomplished through the use of polarized glasses. The third level would encompass data that is almost solely produced by computations. Here, extensive rendering techniques must be used to display the physical systems as envisioned by the researchers. The MSWS could be used to produce "movies" of these systems (fly-bys of planets, geological processes in action, etc.) which would enable the scientists to develop a better understanding of the processes. These "movies" would require large amounts of computations that the MSWS could perform in reasonable times. The experience acquired from developing this software/hardware concept would help GSFC retain in-house developmental capabilities in these areas which are so important to scientific analysis; such an enhanced in-house capability can only help foster good relations between NASA and the outside scientific user community.

It is clear that the current generation of CPUs can not keep up with the huge volumes of data generated by both satellites and computations performed by full scale supercomputers. Data rates from small satellites are approaching tens of megabytes per second and data sets are being measured in gigabytes. In order to deal with these large data volumes, new techniques will have to be developed. A MSWS can serve as an inexpensive alternative to a full scale supercomputer for both data handling and new software development.

The MSWS can attain I/O speeds of approximately 100 Mb/sec which exceed current mini-computer speeds. Coupling the MSWS with a high-performance disk drive (e.g., Ibis 80 Mb/s, etc.) will produce a system that can attain these high data transfer rates. This system could serve as a testbed to study new techniques of dealing with large data sets. So, as NASA enters the space station era, personnel will be available to deal with the large data rates envisioned for the future.

4. Earth Science Impact

It is obvious that the problems of visualization of large-format and 3-dimensional data are not unique to space science. Analogous applications occur in remote-sensing imaging of all kinds (earth resources, medical, seismic, etc.). We believe that development of new visualization technology for the worst-case problems of radio synthesis imaging and COBE data analysis cannot fail to have significant impact on these other applications, even though these potential applications are not a goal of this proposal. For example, multi-disciplinary earth observing programs, such as the Earth Observing System (EOS), require computers which can not only process large amounts of data, but effectively relay their meaning as well. Satellite experiments such as the Landsat Multi-spectral Scanning System, the Airborne Imaging Spectrometer, and the Airborne Visible Infrared Imaging Spectrometer transmit such large quantities of multi-dimensional data that conventional methods of display (2-D, x-y plots; 3-D contours) prove inadequate to extract the information contained within. Therefore, the use of high-speed graphics engines driving high resolution display devices such as the MSWS will become the machine of choice. These graphics workstations will be used for such tasks as volumetric rendering, surface rendering and animation of the data. The volumetric rendering will be used to model a 3-D object, such as a data cube, and, via knobs, change physical parameters in both the model and data. This technique will prove quite useful for multi-spectral data. The surface rendering will enhance techniques for studying topological structures from different perspectives. In the past, such data manipulation required too much CPU time. Finally, animation will help the researcher view data along a dimension that changes in time. For example, geological processes could be viewed in accelerated time.

5. Relationship to GSFC/Code 630 Supercomputer Purchase

A MSWS could be the ideal candidate to be incorporated with the NSESCC "Super Data/Computing" system envisioned for the near future. A high-speed data link (100 Mb/s) would transfer data sets from the supercomputer to the MSWS for further processing, graphics display, etc. This data link would also be used for transferring satellite data sets for special processing on this machine. As a result this MSWS would act as both a graphics-rendering engine for the supercomputer and a display for other high-density data sets.

Further, the MSWS could be used as a compute host for the supercomputer and thus alleviate some of the CPU burden, especially in data-visualization operations. Since the MSWS uses true vector architecture, new algorithms could be developed and tested using this machine. These algorithms then could be transferred to the large supercomputer with minimal impact on supercomputer resources.

VI. TECHNICAL EXECUTION PLAN

During the 3-year course of this experimental development effort, we intend to document our progress to the general science community via presentations (e.g., at the annual American Astronomical Society meetings, special NRAO and NASA workshops, at International Astronomical Union colloquia, etc.). What follows is a strawman scenario of anticipated milestones during the course of this endeavor.

1. First Year

Specifically NRAO, in consultation with Code 630, will:

- Conduct a study of commercial product(s) which can satisfy our technical requirements for interactive computer model fitting and permit real time visual analysis and 3-dimensional display.

- Benchmark test at least two systems from different vendors. NRAO will install AIPS on each candidate workstation, and then will execute the "DDT" benchmarking and certification procedure. The DDT procedure will verify that each workstation can perform radio synthesis imaging calculations correctly; only machines that pass this critical test will be allowed to enter the final procurement process.

- The DDT test verifies only the computational aspects of the machine. The graphics capabilities of the workstation are also critically important. Ideally, a full implementation of the AIPS graphics interface (Y-Routines) would be part of the benchmark. Unfortunately, graphics interface standards are only now at the point where this may be possible. If the Y-Routines can be developed in the limited time available for benchmarking, they will. Otherwise, an evaluation of the attractiveness of the graphics capability will made, based on the available documentation.

- Competitively purchase two identical hardware and software configurations from one vendor. One system will remain at the NRAO and the other at GSFC. Vendor selection will be based on a combination of factors: (1) DDT performance measures, as an indication of both hardware performance and compiler sophistication; (2) Y-routine performance measures or estimates (e.g., how long to load a specified image from disk to virtual display); (3) specific display features which will be conducive to development of the advanced image visualization functions; (4) price; and (5) factors such as maturity of the OS, support for peripherals, compiler quality, and subroutine library quality. The overall aim will be to select the "best" vendor at the moment of the choice while fully qualifying one or more additional vendors.

- Hire a workstation/graphics specialist to work with the the AIPS Group for this project. The person should be chosen as soon as possible to take up residence in Charlottesville, hopefully, at about the time of arrival of the first MSWS.

- Develop or further develop the Y-Routines so that a fully functional AIPS system, including all graphics capability, is present.

- Fine-tune AIPS to the purchased hardware/software system as a stand-alone system, following established AIPS standards.

<u>Specifically Code 630, in consultation with NRAO, will:</u>

- Install the AIPS software, as designed by NRAO, on the hardware/software system located at GSFC.

- Begin the research and development effort as a display engine attached to a host which is anticipated to be a VAX 8800. A critical goal will be to move data efficiently from the cluster environment to the workstation via Ethernet.

- Transform the data received over the network at the MSWS system into a disk-compatible FITS format so that it can be ingested by AIPS.

- Assist various science groups (in particular members of the COBE science team) to accomplish their specific programming needs by (1) designing and implementing graphics algorithms which maximize the data presentation and which fully utilize the graphics rendering hardware; (2) optimizing data transfer between the graphics rendering engine and the VAX 8800 cluster environment; and (3) supplying aid at the systems level to ensure proper disk management, security, and communications facilities for the user.

2. Second Year

Specifically NRAO, in consultation with Code 630, will:

- Extend AIPS, where possible, to take advantage of windowing technology and make the 'virtual TV' concept a reality.

- Investigate the use of the high transfer rates, which the MSWS frame buffers support, and develop new AIPS functions to exploit the high transfer rates to produce, for example, movie-loop displays of multi-spectral data.

- Begin research and development to adapt polygon rendering to astronomical data analysis problems and, if appropriate, develop routines in AIPS to take advantage of this capability.

- Study programmable input control panels and extend AIPS to support these devices.

- Improve interactive user control of display modes and model fitting.

- Extend AIPS to use the more advanced features of the MSWS display. Areas of study could include display of wide field, 3-D images and spectral line cubes, various methods of stereo rendering, etc.

Specifically Code 630, in consultation with NRAO, will:

- Test and validate software developed at NRAO and provide feedback for advanced display code development.

- Adapt the programmable input control panel hardware to COBE analysis model fitting. This task is the heart of the proposal, is very time-intensive, and will be done in collaboration with COBE project personnel, including the Science Working Group.

- Continue helping various other science groups to port their specific software products to the system for research and development purposes. Code 630 will aid this effort at the systems level and provide expertise learned from adapting the system to interactive user control of the model-fitting techniques derived from the COBE experience.

3. Third Year

The third year will be used to complete work begun previously, since the quantity of proposed work, at this juncture, is quite ambitious. For example, we hope that several significant problems from other science disciplines besides this COBE initiative will help demonstrate the feasibility of real-time interactive modeling and 3-D display. We intend to publish the results during this third year in refereed scientific journals and NASA and NRAO technical memoranda. By the end of this project, we envision that the NRAO will be able to integrate most of the proven, experimental package into its public AIPS software distribution.

VII. REFERENCES

1 "First 'Graphics Supercomputers' to Tackle Scientific Visualization," Computers in Physics, **9**, January/February 1988.

2 "Graphics Supercomputers Tackle Real-Time Visualization," Computer Design, page 21, 15 March 1988.

3 "NRAO's Astronomical Image Processing System," by D. Wells, an invited paper in Data Analysis in Astronomy, Proceedings ed. V. Di Gesu et al., Plenum Press, (Meeting held in Erice, Sicily, June 1984).

4 "Stellar Graphics Machine Stresses Interactivity," Electronics, page 97, 17 March 1988.

5 "A Flight Along the Edge of Workstation Technology," Computer Graphics Review, page 28, May/June 1988.

6 "Supercomputing for One," by C. G. Bell, G. Miranker, and J. Rubinstein, IEEE Spectrum, April 1988.

7 "Ardent Launches First Supercomputer on a Desk," Electronics, page 65, 3 March 1988.

8 "Michaels' Goal: Ferrari Performance at the Cost of a Toyota," by Alan Michaels, Electronics, page 68, 3 March 1988.

9 "Astronomical Image Processing System Cookbook," by the NRAO AIPS Group, The National Radio Astronomy Observatory, Edgemont Road, Charlottesville, Virginia 22903, 15 October 1986.

10 "Going AIPS: A Programmer's Guide to the NRAO Astronomical Image Processing System," Volumes 1 and 2, by W. D. Cotton et al., The National Radio Astronomy Observatory, Edgemont Road, Charlottesville, Virginia, 22903, 15 April 1987.

11 "1987 AIPS Site Survey," AIPS Memorandum, Internal Memo Series of the National Radio Astronomy Observatory, 29 April 1988.

12 "Ultraspeeed Graphics for Computational Science," by Karl-Heinz Winkler et al., Cray Channels, **4**, Summer 1987.

13 "NASA Astrophysics Data System Study (Final Report)," March 1988.

14 "Graphics Workstations, Supercomputing, and Remote Communications," by Thomas Lasinski, Cray Channels, **10**, Summer 1987.

15 "The Visualization Challenge in the Physical Sciences," by Robert S. Wolf, Computers in Science, **16**, January/February 1988.

VIII. FIGURES

Figure 1 Caption: Shown on page 38 is a sequence of panels which illustrates the use of interactive modelling in an image context. These data, obtained with the Very Large Array (VLA), represent diffuse, extended 6-cm continuum emission in proximity to discrete, compact 6-cm continuum emission in the elliptical radio galaxy designated M 87 (also known as Virgo A, NGC 4486, or 3C 274). The top left panel shows the raw aperture synthesis VLA image which is convolved with the theoretical instrumental beam shown in the top right panel. The bottom left panel shows the image-beam deconvolution with the CLEAN algorithm (a particular means of beam sidelobe removal). The bottom right panel shows a much improved image-beam deconvolution with the Maximum Entropy algorithm for this particular case. Employing several algorithmic approaches and/or interacting with the modelling process in near-real time permits the user to subjectively judge improvements to the processed data (e.g., increases in dynamic range, subtle data artifact removal, etc.).

Figure 2 Caption: Shown on page 39 is new way of visualizing all-sky data. The sphere represents the cosmic background radiation anisotropy as the DMR will see it. This particular display principally shows the dipole anisotropy (red and blue shift) due to the motion of the earth (observer) with respect to the cosmic background restframe. The three mirrors permit the viewer to see the whole sphere at once without the distortion inherent in planar mapping. The graphics MSWS will allow the scientist to turn the sphere at will, to fit the data in real time, and to display successive models or residuals.

Figure 3 Caption: Shown on page 40 is the programmable control panel associated with NRAO's Image Storage Unit (ISU) project. In conjuction with the commercial IIS image display, the ISU control panel functions include: single image display; interactive split screen display between 2, 3, or 4 images; blinking between 2, 3, or 4 images; intensity-hue display in spectral or cyclic colors; arithmetic image operations (add, subtract, multiply, and divide); 512x512 roam through 1024x1024 images; color overlay of up to 3 images; various modes of image retrieval from storage, including time sequences, movies, and split-screen movies; graphics plane display control; scrolling (registering) of individual images; and transfer functions control.

Figure 4 Caption: Shown on page 41 is a 60 micron IRAF image of the Cygnus Loop supernova remnant displayed with AIPS. The diagonal stripe is due to instrumental noise effects. AIPS was also utilized to intercompare this infrared image with radio imagery obtained with the Very Large Array and with X-ray imagery obtained with the Einstein Observatory satellite.

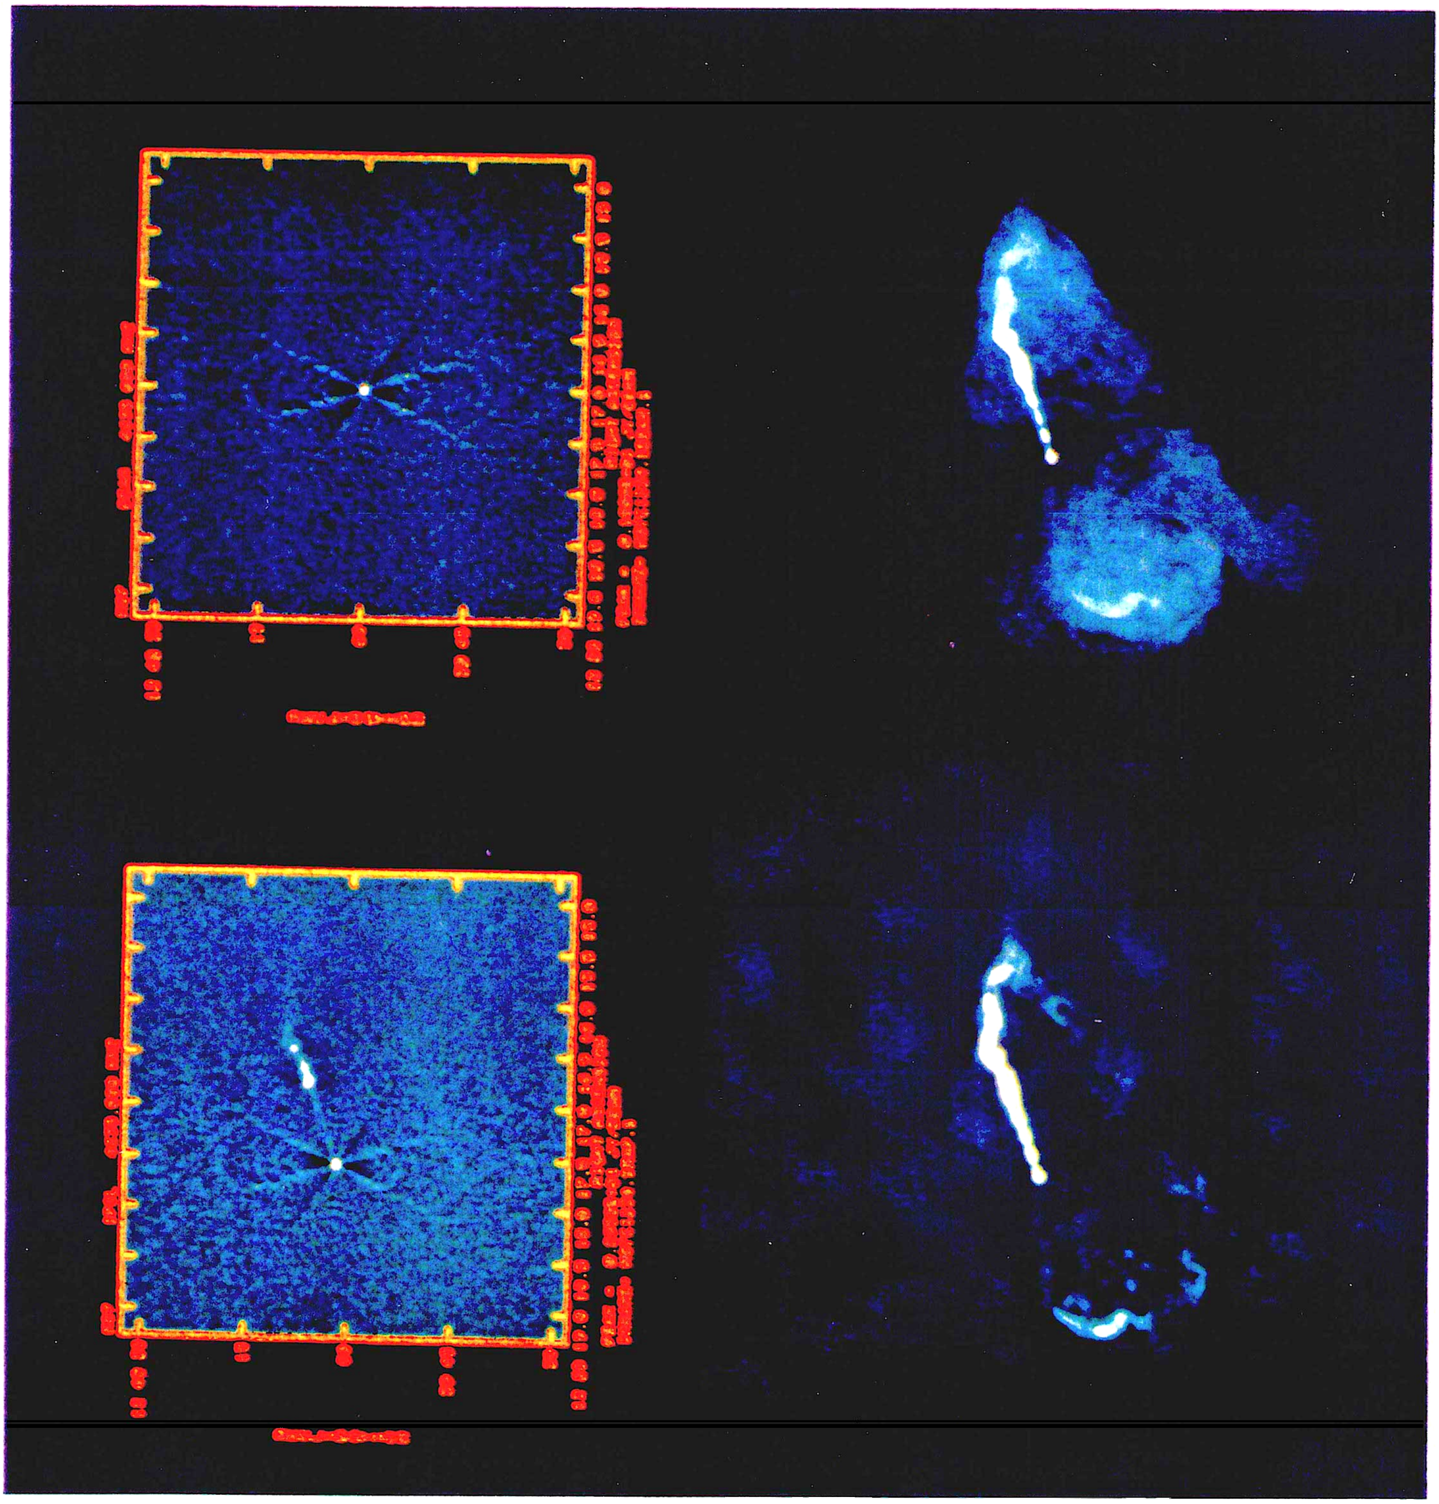

Figure 1

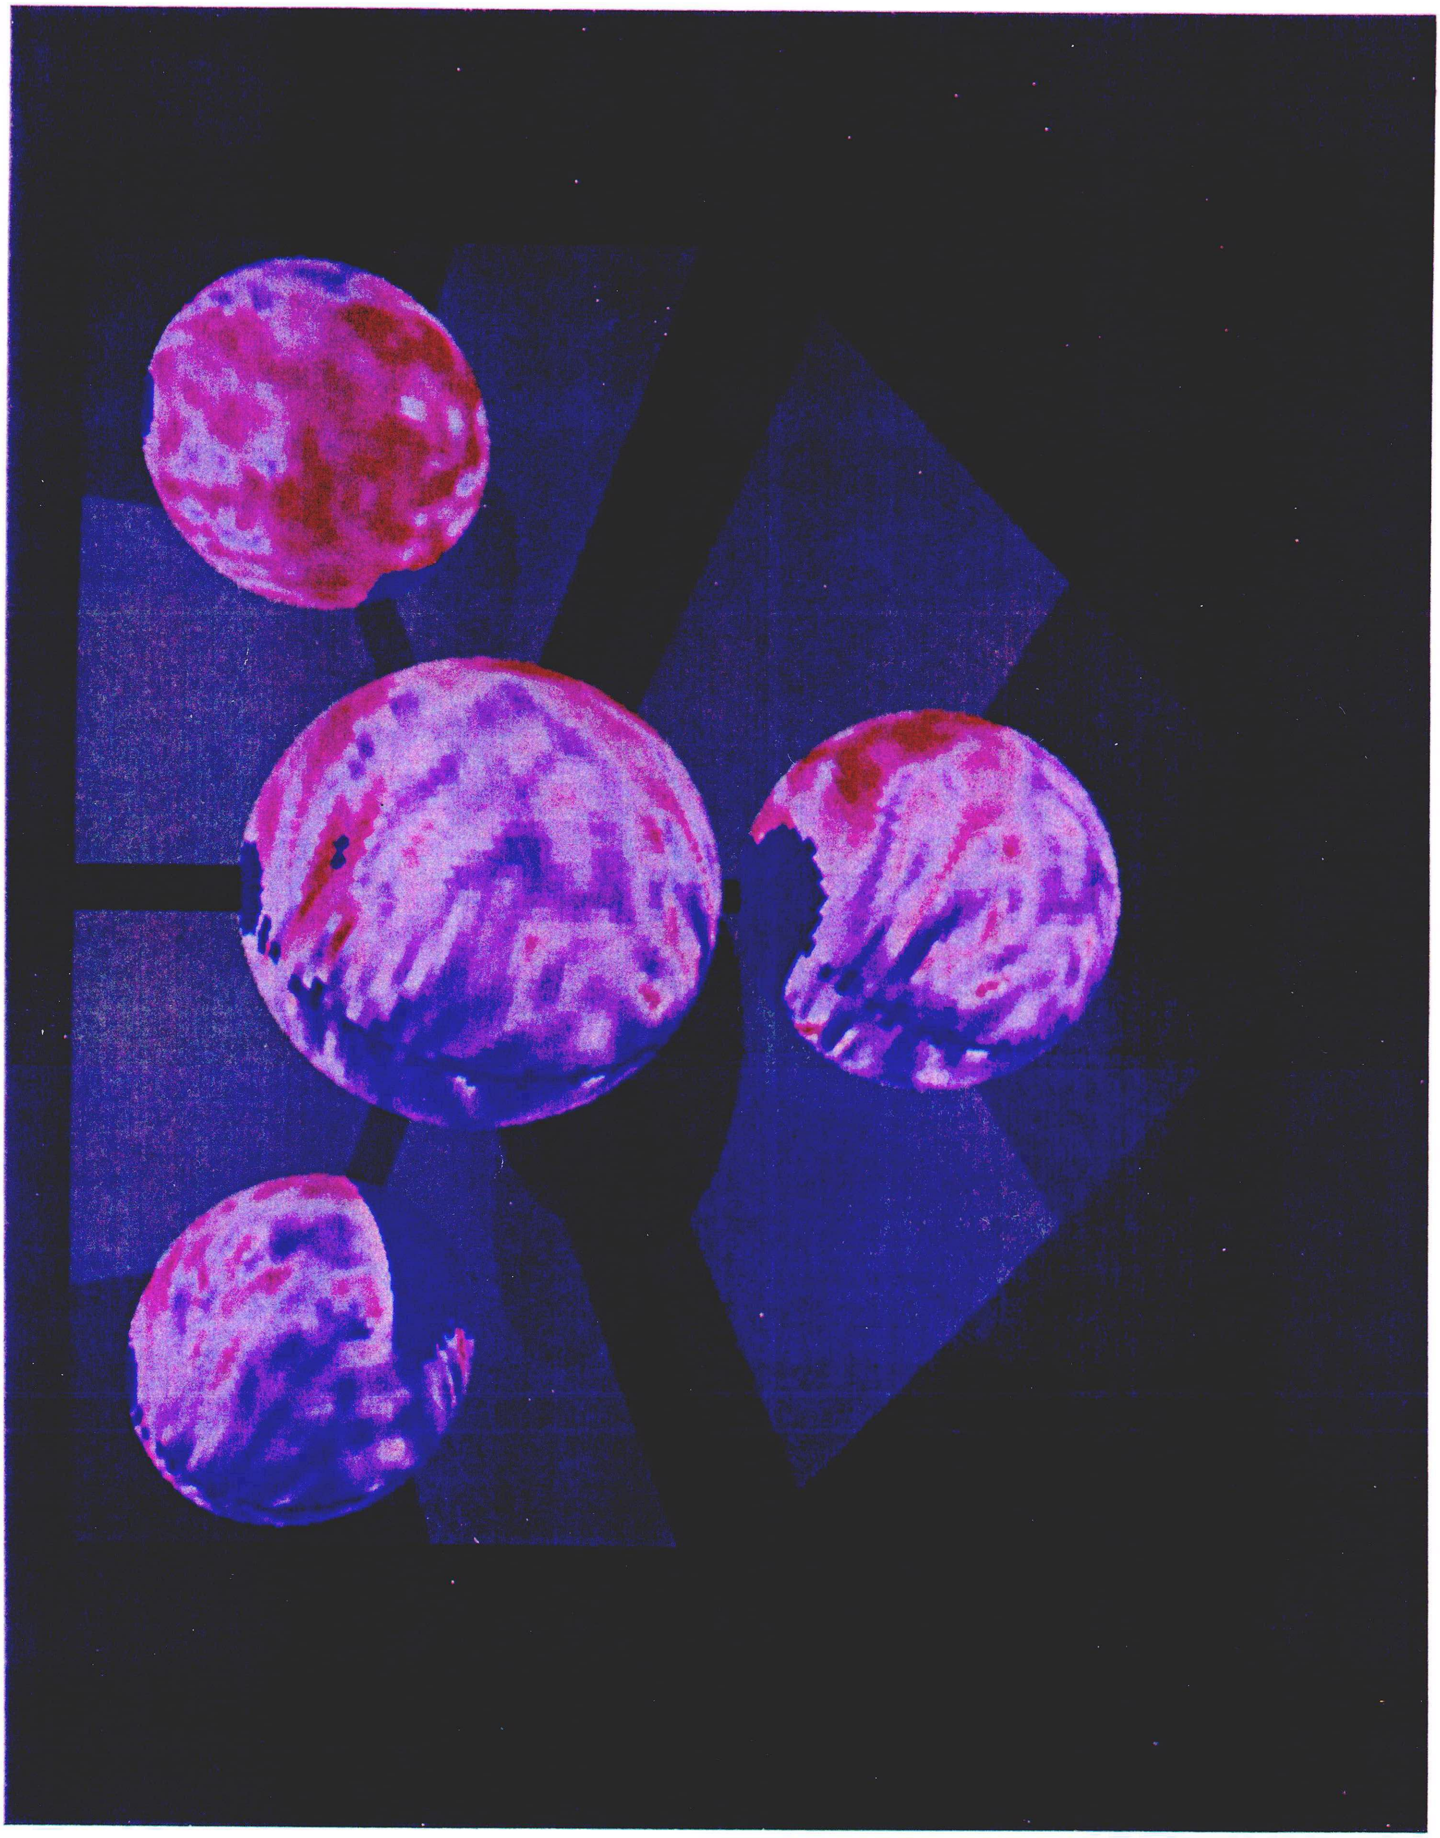

Figure 2

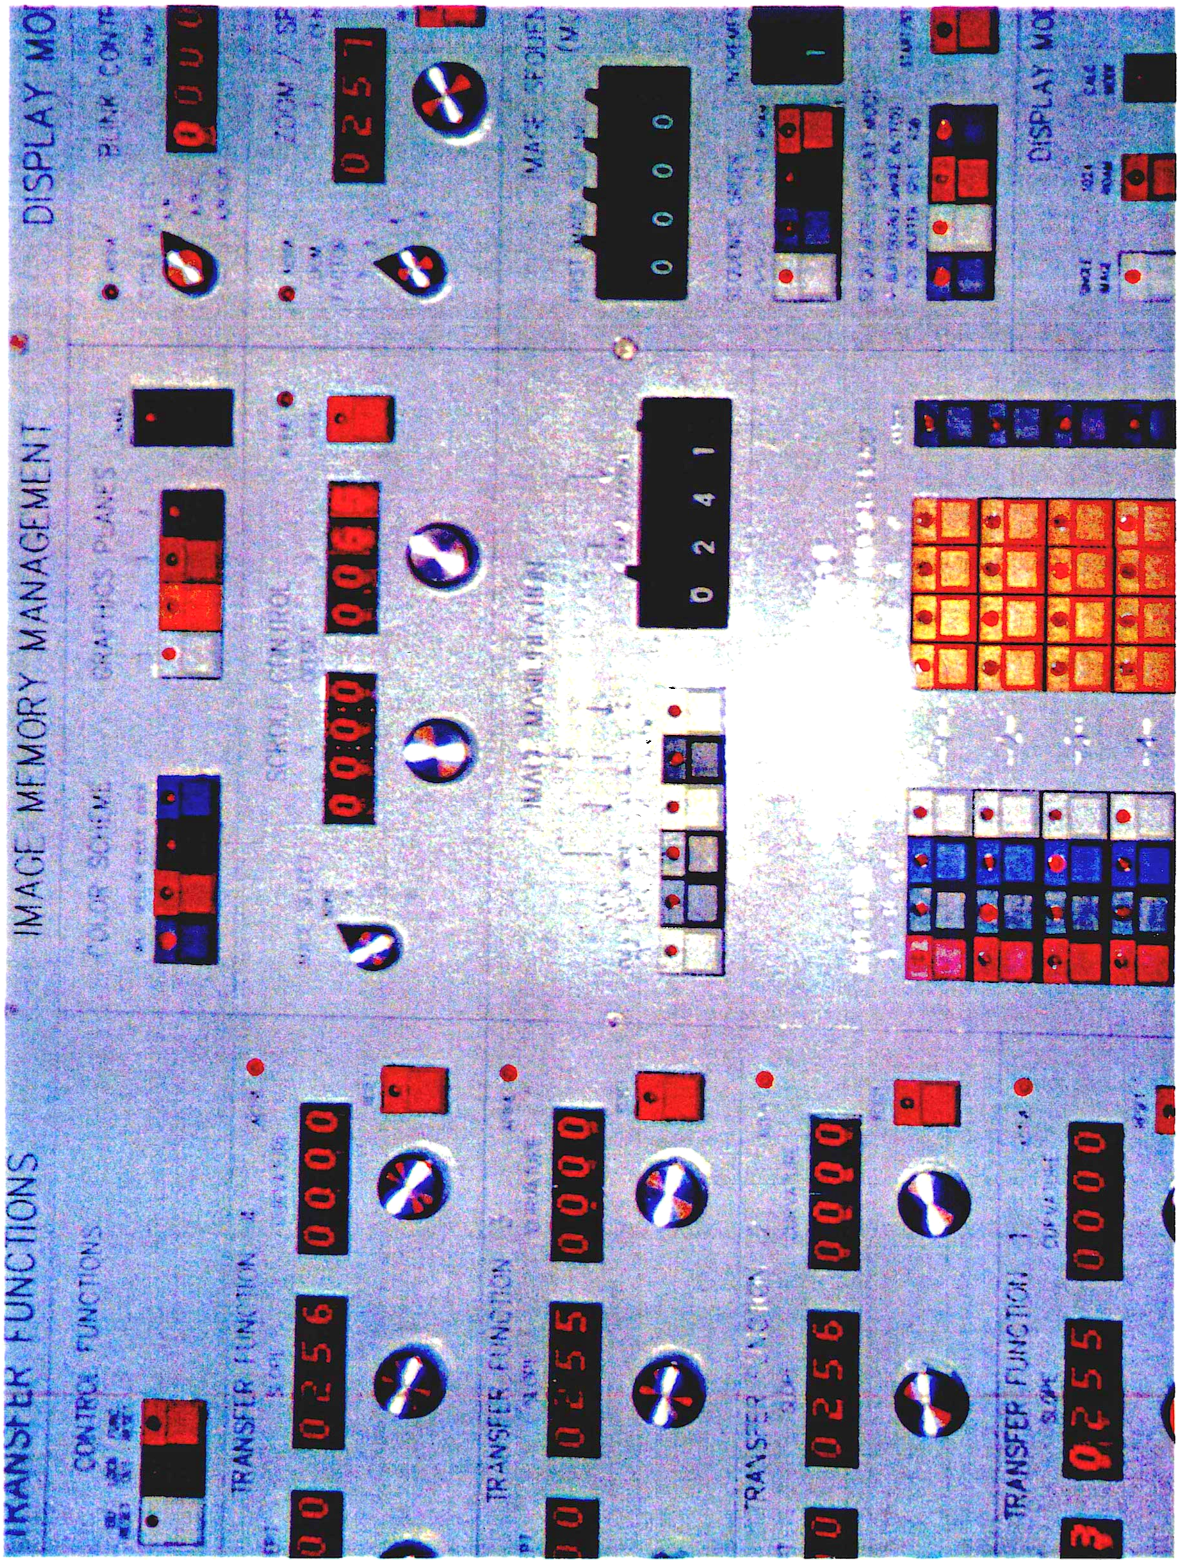
IMAGE MEMORY MANAGEMENT
COLOR SCHEME
GRAPHICS PLANES
CONTROL FUNCTIONS
TRANSFER FUNCTION 1

Figure 3

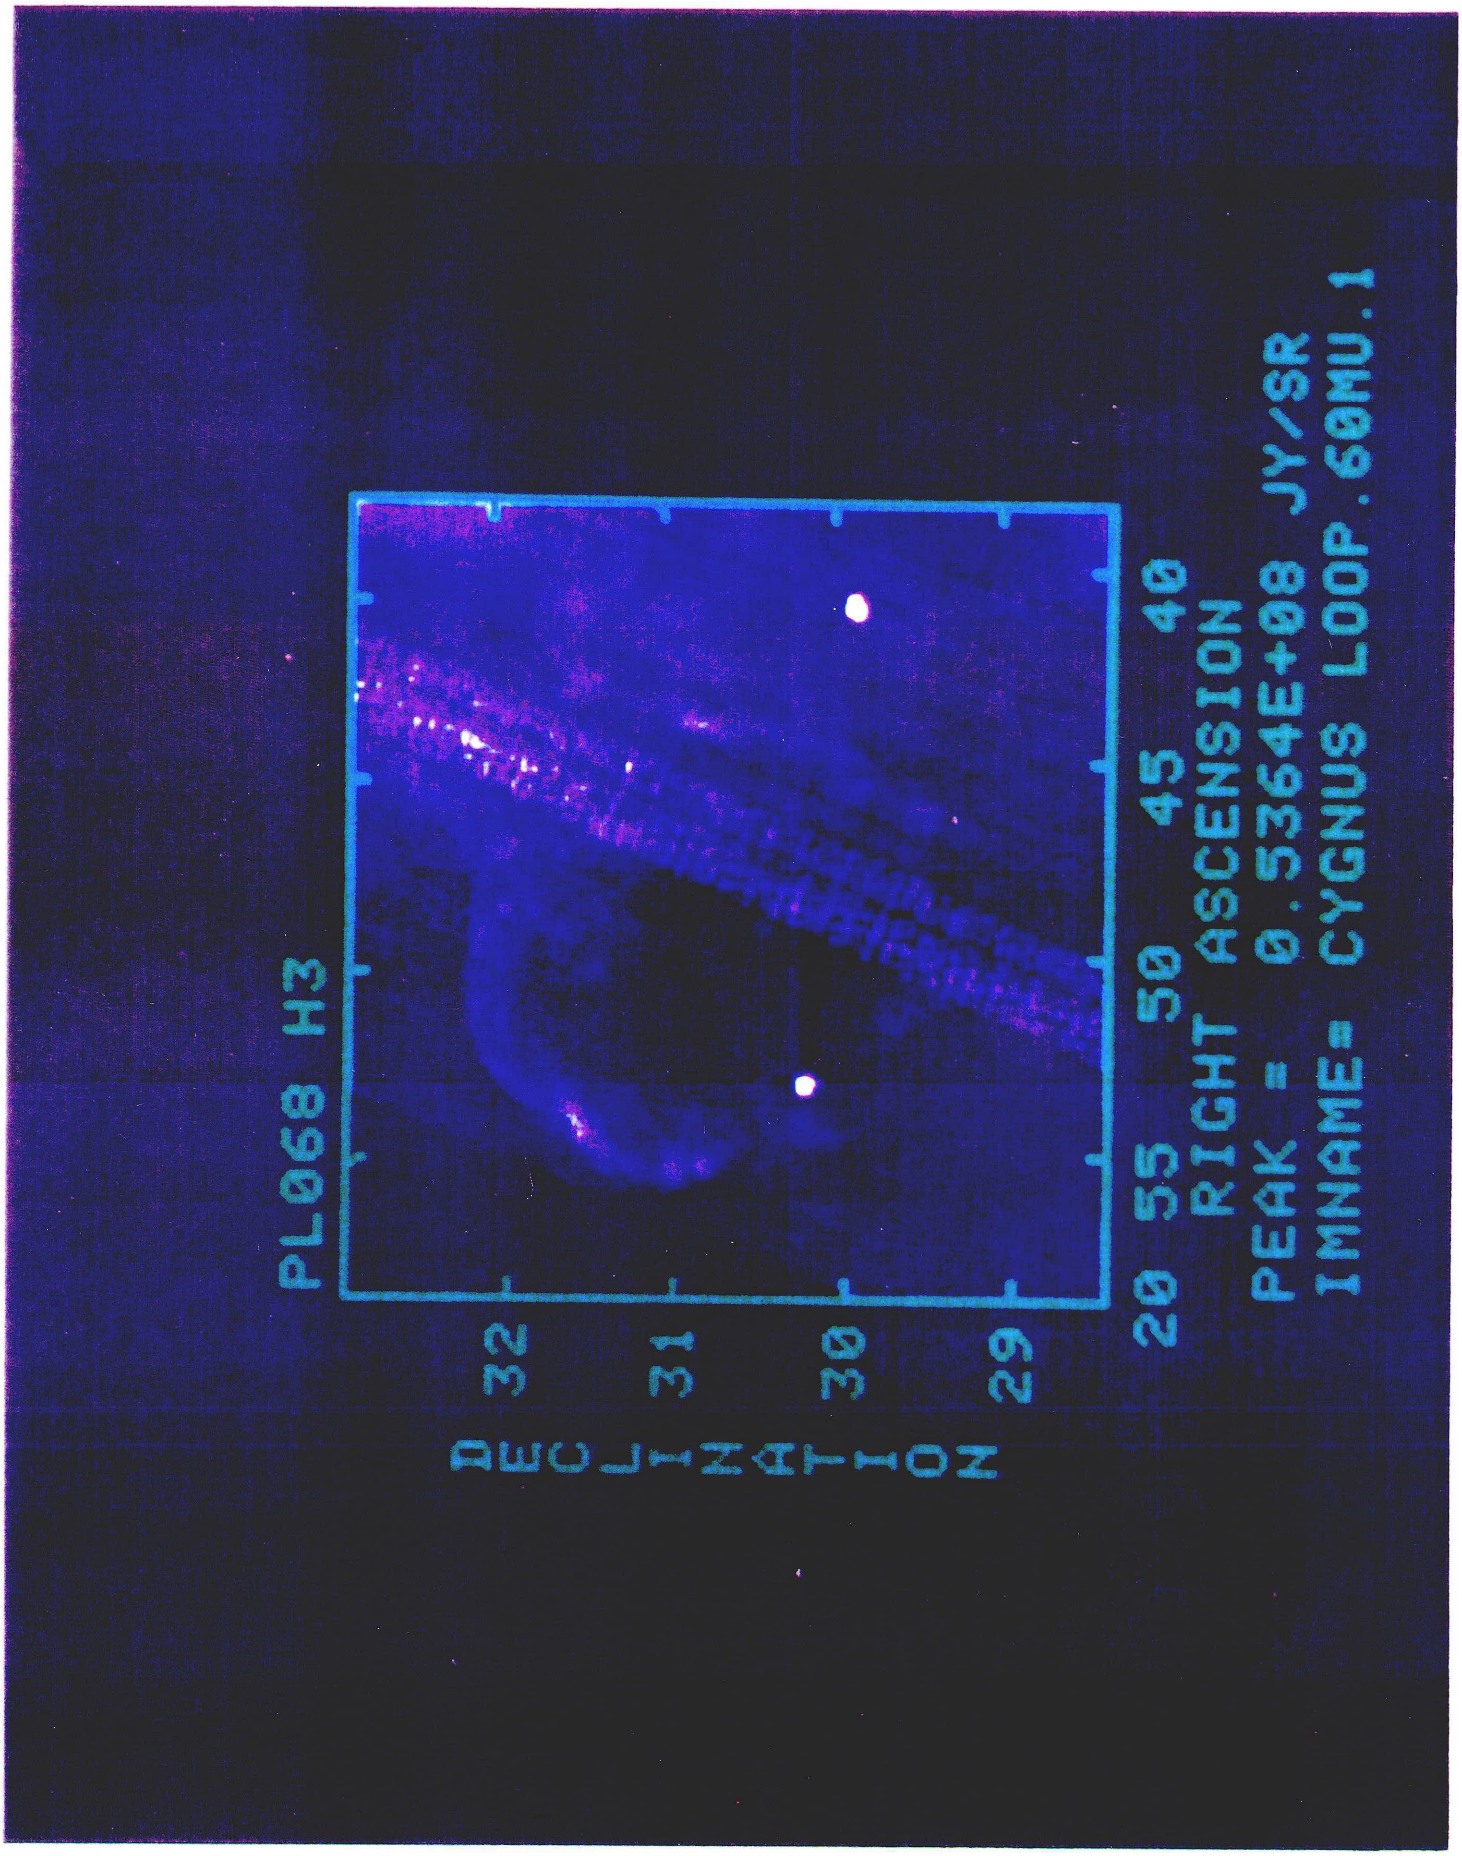
PL068 H3
RIGHT ASCENSION
20 55 50 45 40
DECLINATION
32
31
30
29
PEAK = 0.5364E+08 JY/SR
IMNAME= CYGNUS LOOP.60MU.1

Figure 4

IX. VITAE

Goddard Space Flight Center

Dr. Jan M. Hollis is currently Assistant Chief of the Space Data and Computing Division, NASA/GSFC. He is an astronomer with approximately fifteen years of experience, predominately regarding spectral line mm-wave detections of interstellar molecules, which help determine the physical conditions within galactic cloud complexes, and optical line and radio continuum imaging of stars with surrounding nebulae. Dr. Hollis has an extensive background in astronomical image processing because of his research interests and his former position (1982-1988) as Head of the Science Operations Branch (Code 684) which provides image-processing support for the International Halley Watch, the ASTRO mission Ultraviolet Imaging Telescope, the Hubble Space Telescope High Resolution Spectrometer, and the International Ultraviolet Explorer satellite. Currently, Dr. Hollis is the Cosmic Background Explorer project manager within Code 630, which will be responsible for producing the level 2 analyzed science data sets (i.e., sky maps) from level 1 project data sets. Dr. Hollis is the author of approximately sixty manuscripts in refereed journals such as the Astrophysical Journal, the Astronomical Journal, Icarus, the Review of Scientific Instruments, Astronomy and Astrophysics, the Origins of Life, and Nature. Recent publications over the last five years have concentrated on extracting science from radio imagery taken with NRAO's Very Large Array in New Mexico and from optical imagery taken with CCD devices at Kitt Peak National Observatory.

Dr. Ryszard L. Pisarski is an astrophysicist with the Space Data and Computing Division, NASA/GSFC. He has over six years of experience in the analysis of astronomical data and detector development, which were gained while doing research and supporting various flight projects. His primary field of interest is X-ray astronomy, which includes analyzing supernova remnants and their interaction with the surrounding interstellar medium. Currently, Dr. Pisarski is responsible for the evaluation of the scientific analysis software that will be used in the imaging analysis of X-ray data from the joint U.S./German satellite ROSAT. He also was involved in extensive modelling of instruments, the instrument response, and deconvolution of the data obtained. Such instruments include X-ray detectors aboard sounding rockets and the Einstein Observatory satellite, IR CCDs developed at the Naval Research Laboratory, and the Far Infrared Absolute Spectrophotometer aboard the COBE satellite. Dr. Pisarski frequently publishes his research in the Astrophysical Journal, SPIE, Transactions on Nuclear Science, and various conference proceedings.

Mr. Edward C. Sullivan, a supervisory mathematician, is currently the acting Head of the Science Operations Branch (Code 684) within the Laboratory for Astronomy and Solar Physics, NASA/GSFC. He has approximately 30 years of experience in computer science, numerical analysis, graphics, differential equations, data management, image processing, and scientific archive construction. For example, Mr. Sullivan recently designed and implemented an on-line software system to make International Ultraviolet Explorer satellite spectra available to the general scientific community through the National Space Science Data Center. He has published in such journals as Physical Review A, Planetary Space Science, Icarus, the Journal of Geophysical Research, and Computer Physics Communications.

Dr. Richard A. White is currently an Associate Scientist with the Applied Research Corporation, Landover, Maryland. He is an astronomer with approximately 11 years of experience in extragalactic astronomy. His interests lie primarily in the influence of environment on the form, content, evolution, and state of radio and x-ray activity of galaxies and clusters of galaxies. Dr. White has an extensive background in image processing through his use of wide variety of optical and radio telescopes and x-ray satellites, including KPNO CCDs, the VLA, and the Einstein Observatory. He has developed image processing programs, utilizing IRAF, IDL, and, most particularly, AIPS. He has recently applied this knowledge in support of image analysis and model fitting for the IRAS group at Goddard Space Flight Center. Currently, Dr. White is supporting the COBE Science Working Group at the GSFC in a variety of roles, including image processing using AIPS, zodiacal light modeling, and the production of the Analyzed Science Data Sets; he also supports the Principal Investigator for the Faint Object Spectrograph on the Hubble Space Telescope in the area of image processing using IRAF and IDL. Dr. White is the author of many papers published in the Astronomical Journal, the Astrophysical Journal, and various conference proceedings.

National Radio Astronomy Observatory

Dr. Robert Burns is head of the Computer Division at the National Radio Astronomy Observatory, has over 20 years experience in data processing in astrophysics, and has administrative oversight of the AIPS project. With an educational background in information theory, his early work involved the extraction of signals in noisy environments in both astrophysics and geophysics. Dr Burns has published in Geophysics, Astronomy and Astrophysics, and the Astrophysical Journal and routinely reviews articles for various journals including Computers in Physics. He has organized workshops in astrophysical data processing, image processing, and supercomputing and has contributed to many others. He has also served on review committees for various national and international facilities including the National Center for Supercomputing Applications, the Netherlands Foundation for Radio Astronomy, and the National Optical Astronomy Observatories. His current research interest is in alternative approaches to high-performance computing.

Dr. William D. Cotton is a System Scientist in the NRAO Computer Division. He is an astronomer with thirteen years of experience in imaging objects using the technique of Very Long Baseline Interferometry (VLBI), beginning with his post-doctorate at the Massachusetts Institute of Technology. During this time, he began a program in collaboration with Dr. Steven Spangler (University of Iowa) doing multi-frequency studies of low frequency variable sources; this project has continued to the present and is deeply involved with the NRAO AIPS project described elsewhere in this document. His other primary research interest at NRAO is in the active cores of extragalactic radio sources. Dr. Cotton is a member of the American Astronomical Society (AAS), the International Astrophysical Union (IAU), and the International Union of Radio Science (URSI). He has published in the Astrophysical Journal, the Astronomical Journal, Monthly Notices of the Royal Astronomical Society, and various conference proceedings.

Dr. Eric W. Greisen is currently an Associate Staff Scientist in the Computer Division of the National Radio Astronomy Observatory. He is an astronomer with nearly twenty years experience in the study of interstellar matter, primarily the small-scale structure of interstellar neutral hydrogen. These studies have contributed information on the state of interstellar clouds, the birthplaces of new stars, and the structure of our Galaxy. Such research has required Dr. Greisen to become expert both on radio aperture synthesis techniques and on astronomical data, particularly image processing. Over the past ten years, Dr. Greisen has been the chief designer and manager of NRAO's Astronomical Image Processing System (AIPS), described previously in this proposal. Dr. Greisen was also one of the principal designers of the FITS data format, which has now become both an IAU and a NASA standard. Dr. Greisen has published in several journals including the Astrophysical Journal, Astronomy and Astrophysics, SPIE Proceedings, and the IAU Astronomical Image Processing Circular. He is author or editor of a considerable body of NRAO publications on AIPS and image processing.

Dr. Donald C. Wells is a System Scientist in NRAO's Computer Division, and is currently designer for the correlator software for the Very Long Baseline Array. He is an astronomer with twenty years experience in the application of all forms of digital data systems technology to astronomical research, including seven years of real-time control and data acquisition projects and thirteen years in astronomical image processing. Dr. Wells was the architect of the Interactive Picture Processing System (IPPS) at NOAO (1974-1979), and co-designer of the Flexible Image Transport System (FITS) format (1979). He was Chairman of the Working Group for Astronomical Software of the AAS for six years (1981-87). In recent years, he has concentrated on planning and coordinating the installation of mini-supercomputers, workstations and networks at NRAO.

X. BUDGET

	FY89[1]		FY90[1]		FY91[1]	
	NASA	NRAO	NASA	NRAO	NASA	NRAO
Salaries[2]						
CS Code 630 (1 man-year)	0		0		0	
NRAO (1 man-year)[3]		20		40		40
Travel						
CS Code 630	0		0		0	
NRAO	7		3		3	
Hardware/Software[4]	135		138		138	
Hardware/Software Maintenance	50		50		50	
Publications/Graphics[5]	0		1		2	
Affiliation Overhead						
GSFC	8		8		8	
NRAO[3]		28		56		56
NASA TOTAL	200		200		200	
NRAO[3] TOTAL		48		96		96

Notes:

[1] GSFC/Code 630 will manage the budget, monitor the work progress at both NRAO and GSFC, and will pass funds to NRAO for hardware/software purchases, and maintenance costs.

[2] The COBE project will provide manpower, either civil service or contractual, as required for project-specific software development.

[3] NRAO funds are currently not available; a specific NSF request will be made for the funds or such funding will be sought from other sources.

[4] Hardware/software will be purchased by NRAO; two MSWS systems of the same type will be competitively procured; prices reflect an expected 30% discount and are calculated on a lease-to-purchase basis, assuming a 10% annual fee. Each system is comprised of: 1 x MSWS, 4 x 0.78 GigaByte Striped Disk, 1 x 32 MB memory, 1 x 1280x1024 color display, 1 x 1/4" cartridge tape, 2 x 9-track 6250 bpi tape unit, 1 x programmable analog control panel, 1 x FORTRAN compiler, 1 x C compiler.

[5] Publications of the results of the collaborative effort are projected to be published in refereed scientific journals.