A SUPERCOMPUTER FOR RADIO ASTRONOMICAL IMAGING

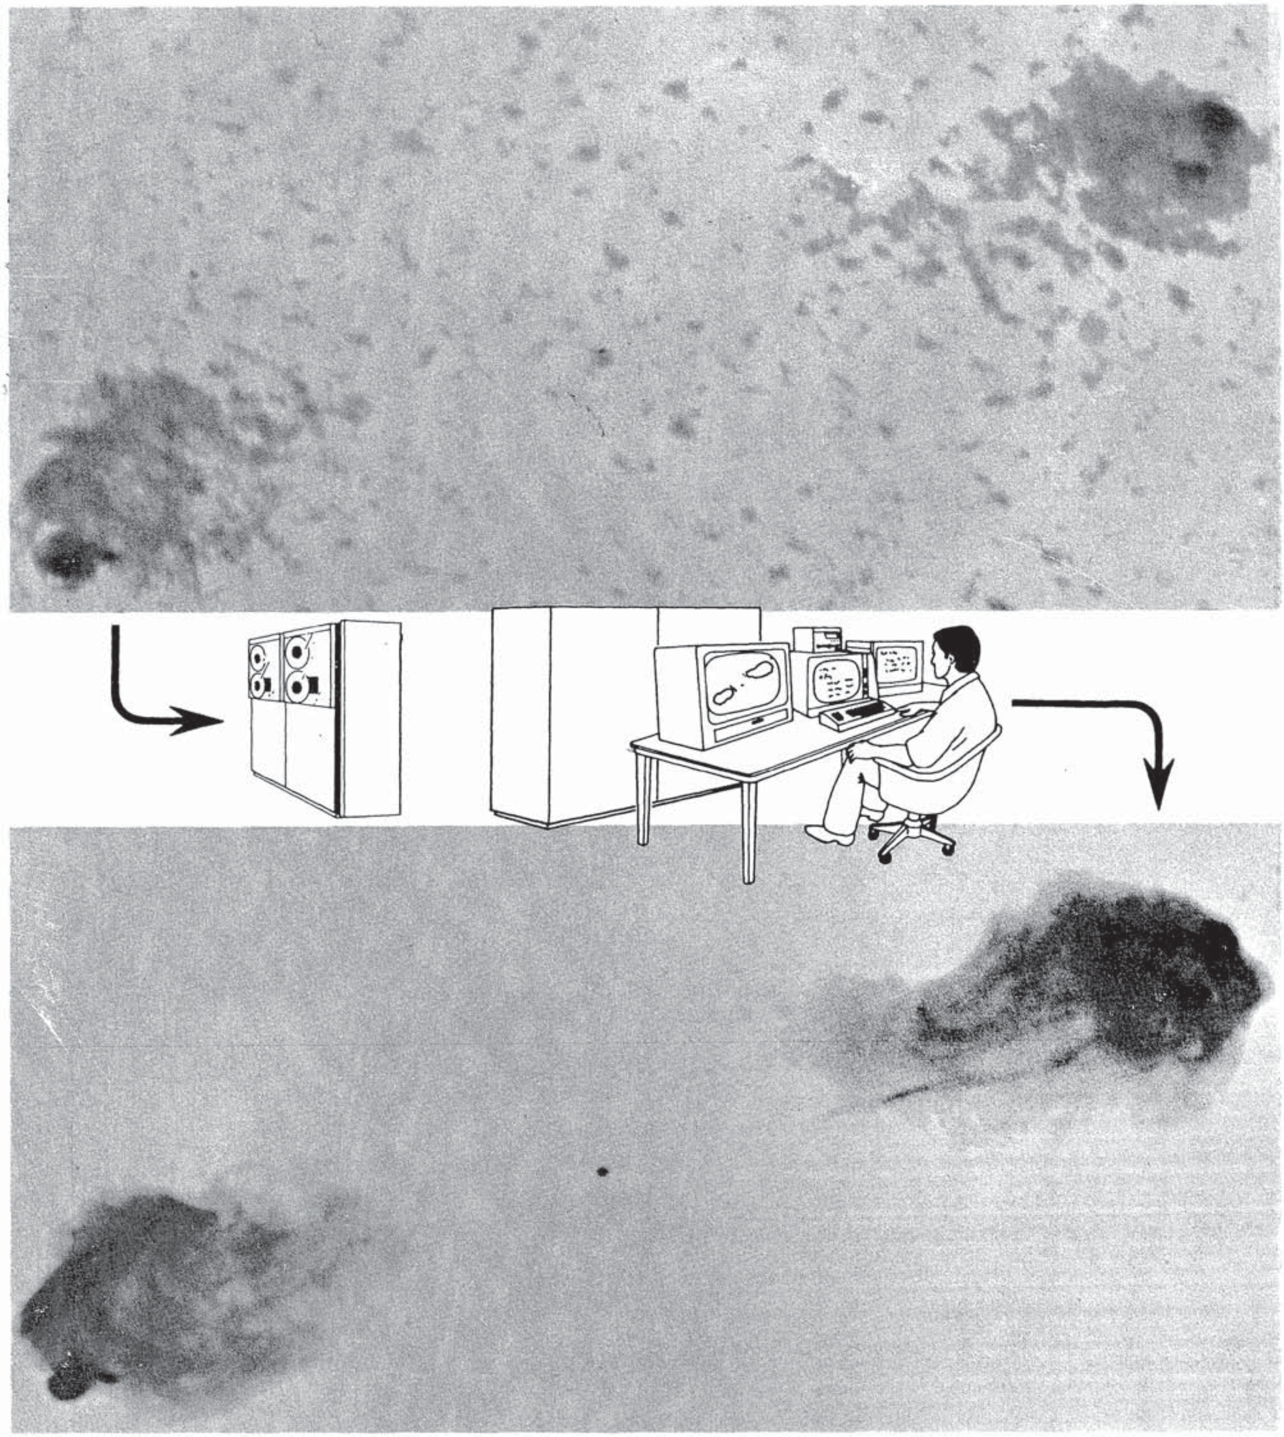

NATIONAL RADIO ASTRONOMY OBSERVATORY

OPERATED BY ASSOCIATED UNIVERSITIES, INC.,

UNDER CONTRACT WITH THE NATIONAL SCIENCE FOUNDATION

A SUPERCOMPUTER FOR RADIO ASTRONOMICAL IMAGING

A Conceptual Proposal
March 1985

Summary

In recent years, astronomers have made increasing use of synthetic aperture radio telescopes, such as the Very Large Array (VLA), to obtain ever higher resolution and sensitivity. Since the image forming optics of such instruments are computers, major improvements can be made even after construction of the telescope hardware is finished. This flexibility has led to a dramatic increase in scientific productivity since both the quality and quantity of the images produced by the VLA now greatly exceed the original goals.

This increase in productivity of the VLA has come at the cost of greatly increased computer usage, but the available computer resources have not kept pace with demands. At this time, only a small fraction of those exciting scientific investigations which are very computer intensive are being pursued.

In 1984, almost 600 observers from 140 different institutions used the VLA facilities to conduct their research projects. On the basis of these projects, and the demands for other projects which are not feasible with the present inadequate computer resources, the National Radio Astronomy Observatory (NRAO) has estimated future demand for data processing. Similar estimates have been made for the Very Long Baseline Array (VLBA) and for other anticipated projects. From this analysis, it is concluded that the future computer demands for image formation and enhancement will exceed NRAO's present total capacity by more than a factor of 25.

These requirements imply a need for an achieved, continuous computer capacity of at least 60 Megaflops. This power is 100 times greater than one of NRAO's current image processing systems and can be provided by one relatively large supercomputer. The supercomputer needs a large, fast access memory (290 Megabytes) and the best achievable disk I/O performance in order to facilitate the processing of the largest images. Total disk capacity must exceed 10 Gigabytes.

Over the past several years NRAO has invested a substantial effort to making a major portion of its software (the Astronomical Image Processing System) truly transportable. Thanks to this prior effort, NRAO is now in a position to make effective use of such a supercomputer on a short time scale.

The proposed supercomputer facility is estimated to cost $18.8 million, only 13% of the $140 million capital cost of the VLA (corrected to 1985 dollars), yet it will enormously enhance the VLA's scientific capabilities. The estimated annual operating cost is $5.6 million, $2.5 million above NRAO's current computer operations expenditures. An investment of this size will make the VLA into a new instrument capable of developing new frontiers in astronomy.

A Supercomputer for Radio Astronomical Imaging

Contents

Preface		1
CHAPTER 1	ASTRONOMICAL MOTIVATIONS & COMPUTATIONAL REQUIREMENTS	3
1.1	Motivation	3
1.2	Examples of Computer limited Problems	8
1.3	Astronomical Requirements	11
1.4	Computer Performance Requirements	11
1.5	Why a Supercomputer?	16
1.6	The New Astronomy	18
CHAPTER 2	THE PROPOSED FACILITY	22
2.1	Hardware Configuration	22
2.2	The Astronomical Image Processing System (AIPS)	25
2.3	Interactive Image Processing	26
2.4	Communications Needs	27
2.5	Archival Storage Needs	28
2.6	Physical Plant	29
CHAPTER 3	OPERATING THE SUPERCOMPUTER	30
3.1	NRAO's User Community	31
CHAPTER 4	COST ESTIMATES	33
CHAPTER 5	CONCLUSIONS	34
APPENDIX A	SCIENTIFIC REVIEW PANEL	35
APPENDIX B	NRAO COMPUTER ADVISORY GROUP	36
APPENDIX C	THE EXAMPLE CONFIGURATION	38
APPENDIX D	BIBLIOGRAPHY	46

A Supercomputer for Radio Astronomical Imaging

Preface

For several years it has been obvious that current computer capacity is insufficient to handle the projected load of the VLA operating at its full scientific capability. NRAO has responded to this problem by formulating a long-term plan to increase its computer capacity. This plan has been reviewed by a scientific review committee, which considered and supported the scientific objectives presented here, and an external Computer Advisory Group consisting of prominent professionals in the data processing field, who reviewed and supported the technical arguments in the plan. Appendices A and B list the members of these two groups and quote from their recommendations. In addition, NRAO staff members have visited a number of supercomputer installations and the major American vendors of supercomputers and array processors, in order to better define the requirements and achievable performance for the type of processing needed. This plan is a result of these investigations and the recommendations of the two review committees.

The overall long-term computing plan for NRAO encompasses three aspects:

- i) developing existing software for use with supercomputers and gaining experience in the use of supercomputers;
- ii) beginning regular use of a supercomputer to provide some additional interim data handling capacity for NRAO users; and
- iii) acquiring a supercomputer.

Part i of the plan is in progress. It was begun with experiments using the facilities of the Los Alamos National Laboratory. Currently, development work is being done on facilities made available through the supercomputer access program of the National Science Foundation's Office of Advanced Scientific Computing. Part ii of the plan is expected to begin before the end of 1985. The remainder of this document addresses the third part, the acquisition of a supercomputer for NRAO.

CHAPTER 1

ASTRONOMICAL MOTIVATIONS & COMPUTATIONAL REQUIREMENTS

1.1 Motivation

Since its completion in 1980, the National Radio Astronomy Observatory's Very Large Array (VLA) has made spectacular contributions to many branches of astronomy. With an angular resolution better than that of any earth-based optical telescope, it has imaged the emission from quasars and radio galaxies with unprecedented detail. The most dramatic among these results are the relativistic plasma jets emanating from galactic nuclei and powering giant radio lobes. Discoveries in normal galaxies include incipient radio jets, giant interstellar bubbles, and the first detailed studies of the radio light curves from supernovae. Images of our Galaxy also have been full of surprises: the nucleus is surrounded by spiral-shaped ribbons of gas and, on a larger scale, an enormous Galactic flare is traced by streams of relativistic plasma which seem to be wrapped in filaments of thermal material. Observations with both high sensitivity and high angular resolution have opened up new areas in the field of stellar astronomy. New discoveries have been made of quite unexpected radio emissions from stars; examples include highly polarized bursts from flare stars and close binaries, and "non-thermal" radiation from unusually hot or energetic electrons in stellar winds. Observations of atmospheres, rings and satellites of planets are complementing the research done by planetary space probes. Looking to the future, the VLA will provide astronomical images with resolution comparable to that of the Space Telescope. The VLA is shown in Figure 1-1 and described in [ref 1].

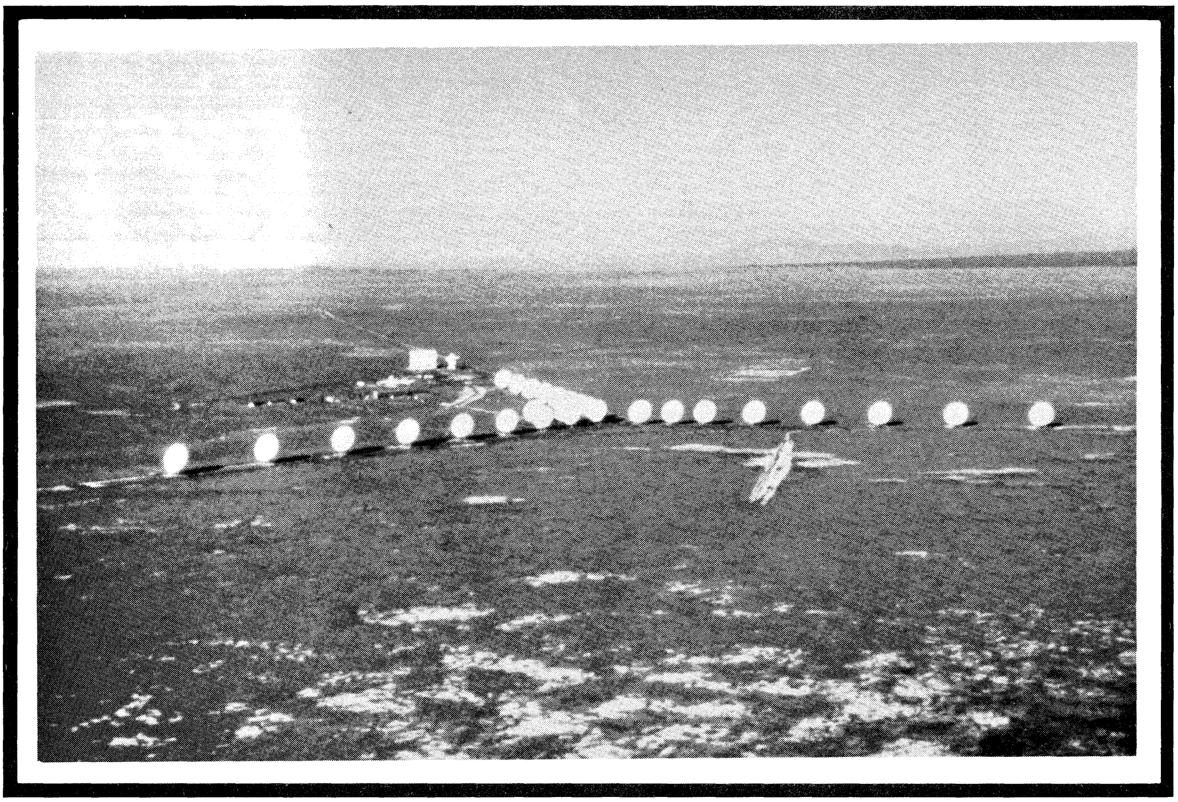

Aerial view of the Very Large Array
as seen from the northeast

Figure 1-1

A Supercomputer for Radio Astronomical Imaging

A conventional monolithic radio telescope with the resolution of an optical telescope would require an aperture many kilometers in diameter. It would be prohibitively expensive and probably impossible to construct. The development of the digital computer made it possible to achieve this resolution with "aperture synthesis" radio telescopes. These telescopes use arrays of relatively small antennas to measure the coherence of the wave front over a large area; an image is then formed in a digital computer which now becomes an essential component of the telescope. NRAO has one major synthesis telescope in operation, the Very Large Array (VLA). Another, the Very Long Baseline Array (VLBA), has been funded for construction to begin in 1985.

When the VLA was originally proposed in 1971, it was expected to produce images in which the ratio between the brightest features and the faintest believable features would be limited by various errors to 100:1 (the dynamic range). The two main effects which cause this limited dynamic range in synthetic arrays such as the VLA are: 1) incomplete measurement of the coherence of the wavefront due to the finite number of antennas and limited observing time; and 2) wavefront collimation errors due to the spatially and temporally varying refractivity of the atmosphere above the array (i.e. what optical astronomers call "bad seeing"). In recent years two new techniques, deconvolution [ref 2] and self-calibration [ref 3], have been developed which can correct for both of these sources of error. Their use has greatly increased the quality of the best VLA images from the 100:1 dynamic range specified in the original proposal to over 10,000:1, but at the cost of greatly increased computation time. The impact of self-calibration has been even greater in Very Long Baseline Interferometry where true imaging of objects, at resolutions up to a thousand times greater than the VLA's capability, is now possible.

The cover of this proposal illustrates the dramatic improvement in image quality achieved when these computational techniques are applied to an image of the archetypal radio galaxy Cygnus A [ref 4]. The results of three separate steps in this process are shown in Figure 1-2.

The upper image is the result of the normal processing and satisfies the original design specifications of the VLA. In the central panel the image has been corrected for the effect of the incomplete measurement of the wavefront by using a nonlinear deconvolution algorithm. It now shows the thin line of emission tracing the energy flow from the central core, which is the nucleus of the parent galaxy, to the upper lobe. Finally, the lower image has also been corrected by self-calibration for the effect of atmospheric "seeing" errors and for the first time the structure in the lobes is seen to be in a filamentary form. The wealth of unexpected detail discovered in the final image has dramatically enhanced its scientific value.

In addition to deconvolution and self-calibration, two further factors have contributed to the computer load: 1) the new image processing techniques enable VLA observers to produce many more images by using much shorter observing times than were thought to be useful ten years ago (VLA observers call such images "snapshots") and 2) the VLA hardware has been augmented to permit simultaneous recording of up to 512 different radio frequency bands. The net result of these enhancements is a peak data output rate almost a million times greater than that anticipated when the telescope was first proposed. Although the capacity of digital computers also increased greatly during this period, it did not increase by the same factor; consequently, for many experiments the data output is now intentionally limited and the science which could be performed with the VLA is seriously compromised by the available computational resources.

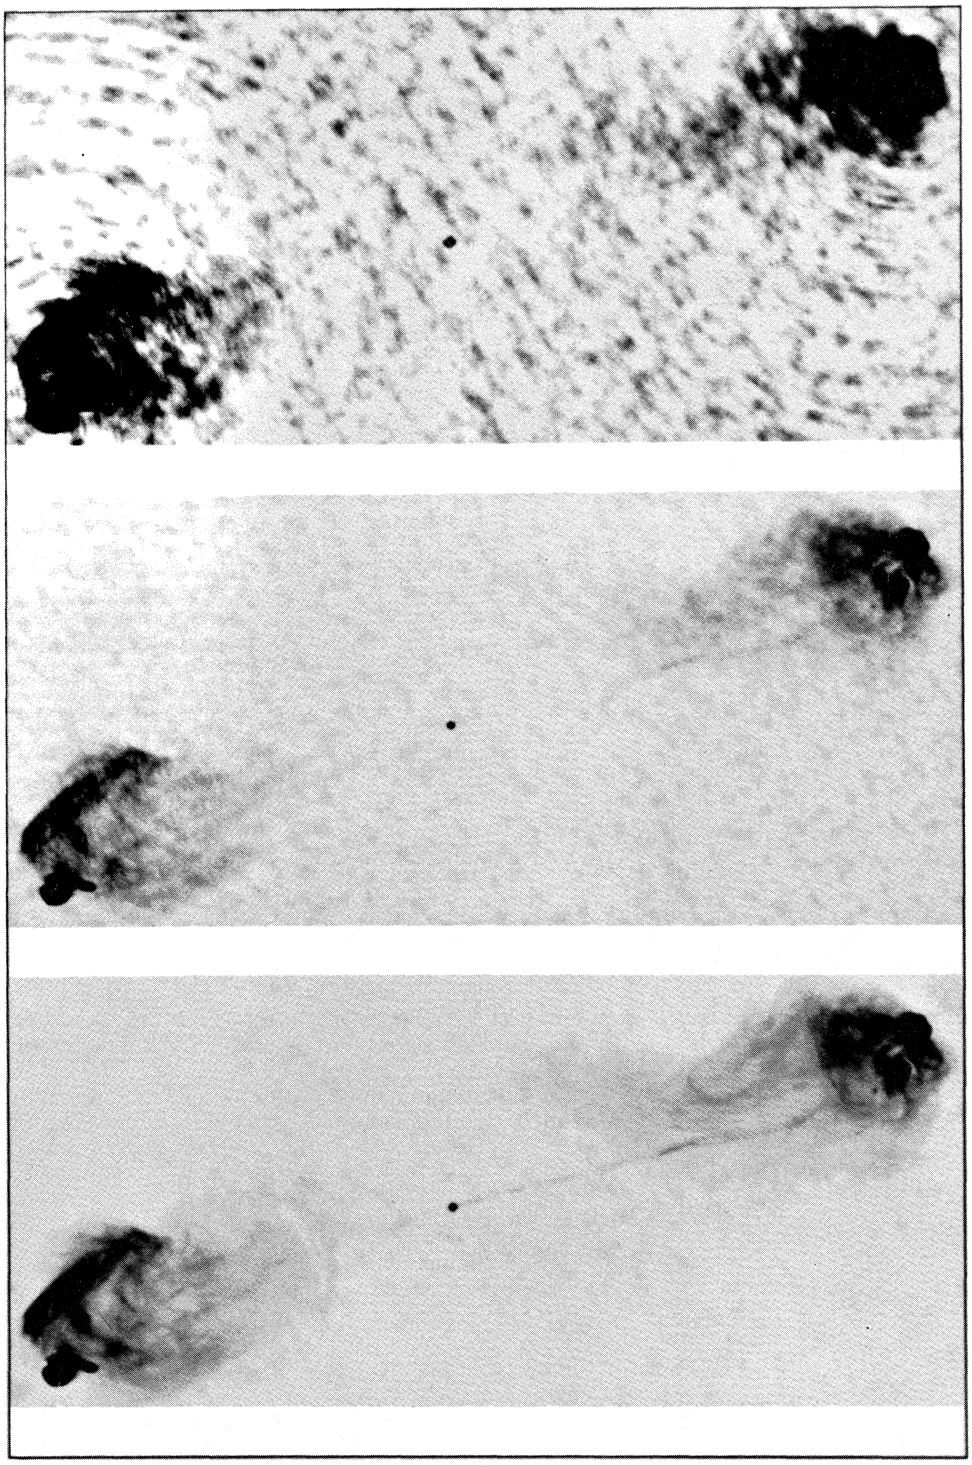

Figure 1-2

The powerful radio galaxy Cygnus A; (top) before any computer image enhancement, (center) after use of the deconvolution algorithm to correct for incomplete measurements, and (bottom) after both deconvolution and removal of atmospheric "seeing" errors.

As remarked in the Curtis report to the National Science Foundation [ref 5], lack of access to supercomputers has inhibited activity in certain scientific fields. This situation certainly occurs with synthesis radio telescopes: the best astronomers, conscious of computer limitations, will simply turn to other problems and the potentially most imaginative and most important problems in radio astronomy may not be studied. For observations which are just feasible with the present resources, the long turnaround time inhibits both experimentation with data analysis methods and investigation of unusual effects. The process of scientific discovery is a fragile one. Limited computer resources which increase the barrier between the scientist and his observations interfere with this process.

1.2 Examples of Computer Limited Problems

To make these points more concrete, it is useful to look at some specific examples based on actual projects which have been attempted with the currently available resources. The mapping of the remnant of the supernova explosion called Cassiopeia A is currently in progress [ref 6]. The scientific objective is to determine the distribution of intensity and the polarization of the radio emission with a resolution of 0.7 arcseconds over the whole remnant of size 6 arcminutes. This will be compared with the existing infrared, optical, and x-ray images to study the development of turbulence, and, after second epoch observations, to measure directly the motion of the turbulent features. The data result from a total of 26 hours observing, and comprise five million independent measurements of the wavefront coherence. The data are stored on twenty magnetic tapes. Creating a low resolution (2048 x 2048) image (Figure 1-3) used three days of computer time (exclusive use) on a current NRAO image processing system (a VAX-11/780 with an FPS AP-120B array processor). Creating the desired full resolution map (4096 x 4096) will take 12 days computing. Final correction of this image would require a further few months of time on the same dedicated computer. Projects of this scale could be processed on a supercomputer in less than the time it takes to observe them.

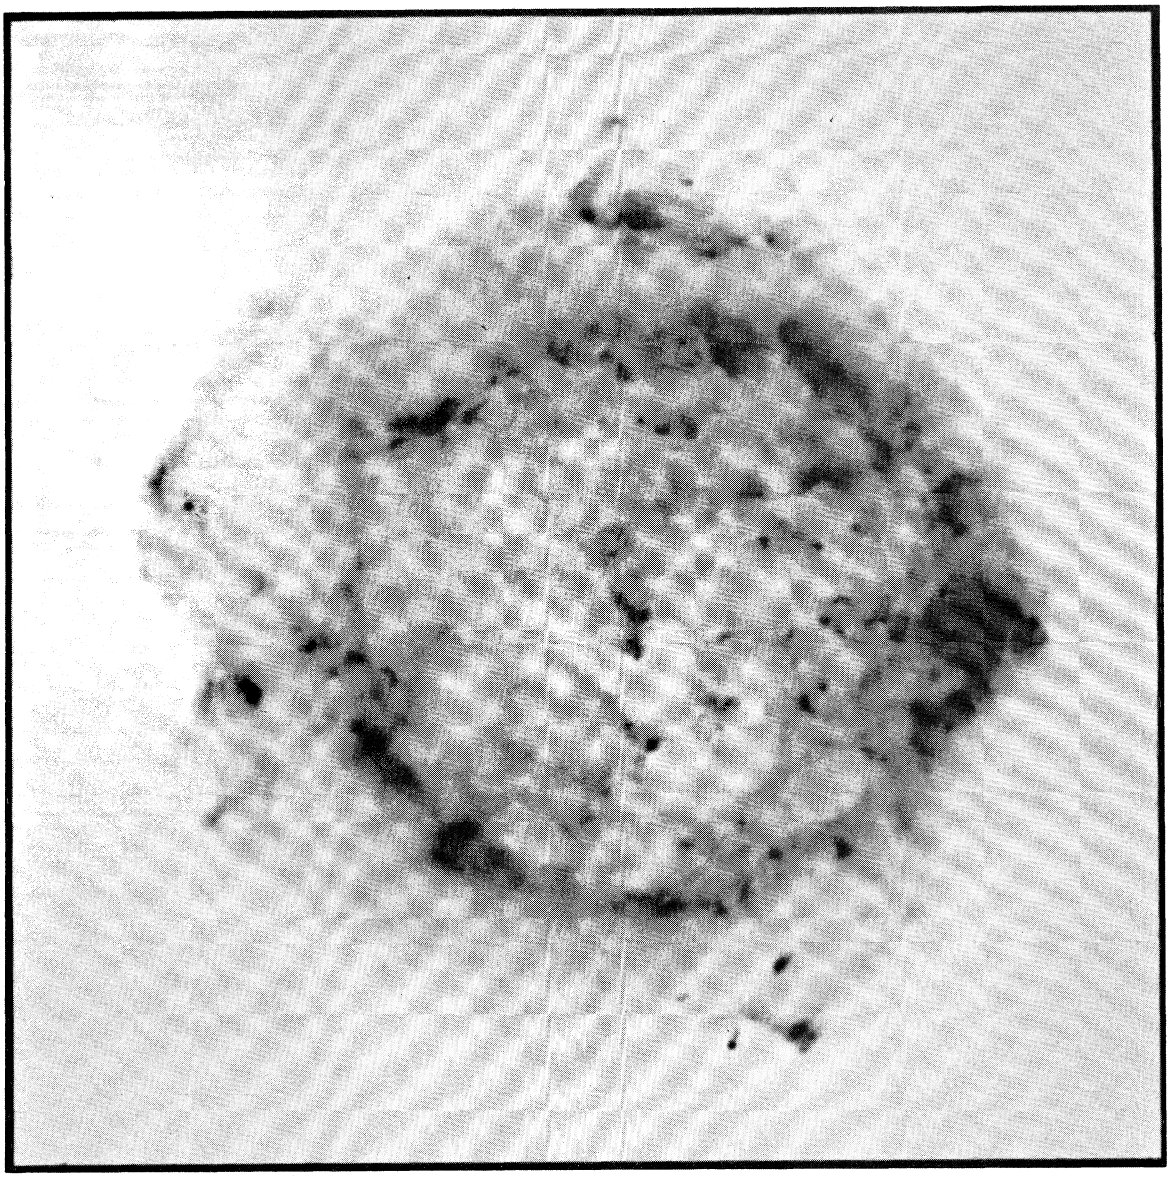

Figure 1-3

Remnant of Exploding Star in the Constellation of Cassiopeia

This is the result of deconvolution of a 2048 x 2048 image. It required three days computing time on a current NRAO image processing system, but has only one-third the possible resolution.

The second example involves determining the mass distribution in the central region of our Galaxy from the distribution of the velocities and the positions of a special type of star which emits a sharp radio frequency line due to OH maser emission in its atmosphere [ref 7]. In this type of observation, the OH maser stars are found by searching in a three dimensional cube (two spatial dimensions and one velocity dimension). With the present array hardware, the resulting images consist of a cube for each of twelve different fields, with each cube of dimension 1024 x 1024 x 512. The observing time for this project is five days, but it can only be computed in ten days using the present NRAO image processing system. When all of the capabilities of the VLA hardware are used, it will be possible to increase the depth of the cubes to 1024 pixels, and to obtain the cubes in each of two orthogonal polarizations simultaneously for two different OH lines. The same five days of observing time will then require 55 days computing time on the present NRAO image processing system! The final result of this experiment is a list of the positions and the velocities of the few hundred stars which are found in these cubes. Hence, the last step in the analysis leading to the mass distribution in the center of our Galaxy is a relatively simple one. In projects of this type, the extensive computations are only required as an intermediate step; the final result involves analysis of a very modest amount of data.

The preceding discussion and examples were based on experience with the VLA. Although there are other synthesis telescopes either now being funded, such as the VLBA, or which will be proposed in future, such as the millimeter wavelength array and the orbiting satellite arrays, such detailed analysis of their requirements cannot be made before observational experience is gained. Crude estimates indicate that these future telescopes will have requirements probably less than, but perhaps comparable to, that of the VLA [ref 8].

1.3 Astronomical Requirements

In order to make a detailed estimate of the future data processing load, NRAO has made an in-depth analysis [ref 9] of the projected scientific requirements of the VLA. This analysis of future research needs consisted of the following steps:

i) Listing all types of astronomical projects which currently use, or are projected to use, the VLA and tabulating the observational parameters relevant for estimating processing needs;

ii) Determining the required processing capacity on the assumption that the astronomical objectives should not be compromised by the computer, while accepting any other hardware limitations; and

iii) Estimating the fraction of observing time that would be spent on each project based on a projection of the current use of the VLA. This estimate is necessary to calculate the time-averaged throughput required, given the widely varying amount of processing required for different projects.

1.4 Computer Performance Requirements

In order to convert these astronomical requirements, which are specified in terms of parameters such as image size, into computer hardware requirements, NRAO has analysed [ref 10] the imaging process in terms of the number of floating point operations and the input/output (I/O) bandwidth required. The analysis involved a hypothetical computer of conventional architecture, with sufficient main memory to minimize the use of disk I/O for sorting, and with computational power of five million floating point operations per second (5 Megaflops). The results of this analysis have been scaled to other computers by the ratio of the Megaflops ratings. Figure 1-4 summarizes the results of these two investigations.

In Figure 1-4, the various categories of astronomical projects are indicated by blocks in which the abscissa is the computing time in hours required for each day of observing. The time scale is set for a supercomputer with a continuous achieved capacity of 60 Megaflops. Such a computer, on a long-term average, will just keep up with the projected VLA load. To scale this diagram to a current NRAO image processing system, the unit of time should be four days instead of one hour! The ordinate is the fraction of the observing time which would be devoted to each class of observation. Thus the time averaged throughput corresponds to an area in this diagram, and the capability, in terms of turnaround time for a given observation, is given by the extent of each block in the horizontal direction. The shaded area indicates the total capacity of the proposed 60 Megaflops (sustained average) supercomputer, and the dotted inset shows the total capacity of all of NRAO's present computer systems.

The information on turnaround time in Figure 1-4 has been redrawn in Figure 1-5 as a cumulative distribution. Note that the abscissa is on a logarithmic scale. Figure 1-5 illustrates that 40 percent of VLA astronomical projects can be handled in less than real time with the present NRAO image processing system, and explains why much good science is being done even with the limited computing resources now available. These VLA projects are ideally suited to the superminicomputer systems which are now in use both at NRAO and in universities. However, the diagram also shows that 30 percent of the projects would require more than 30 days central processing unit (CPU) time with a system this size, a figure which experience suggests would be totally impractical. The line scaled for a 60 Megaflops supercomputer shows that 80 percent of projects will be computed in less than real time with such a computer, and that 95 percent will be handled with a turnaround not much longer than real time.

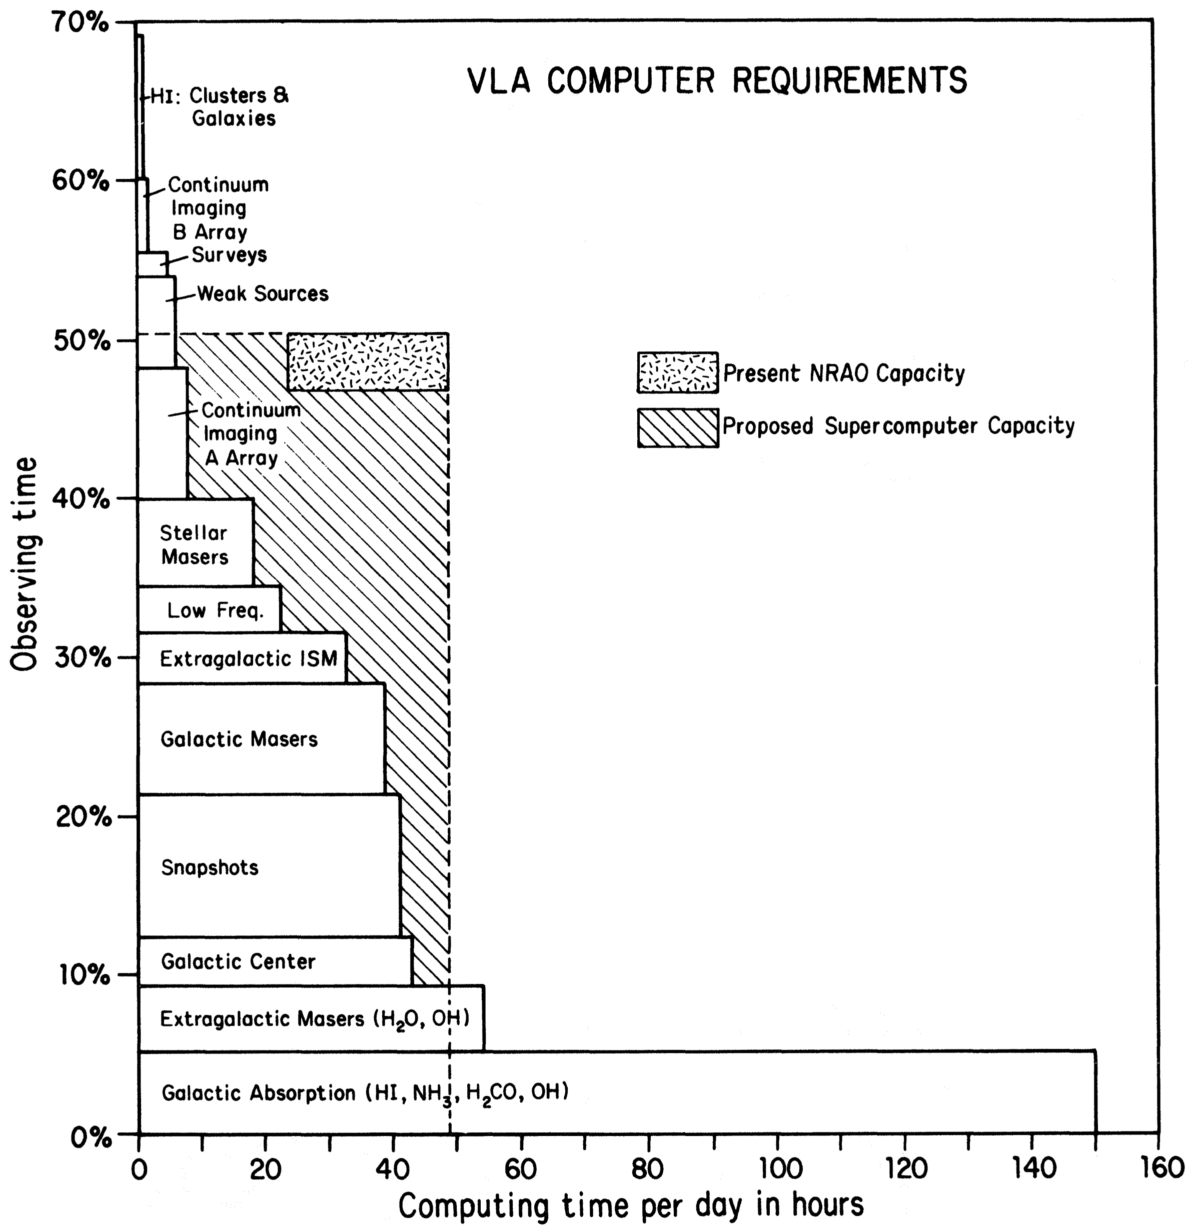

Figure 1-4

VLA Computer Requirements

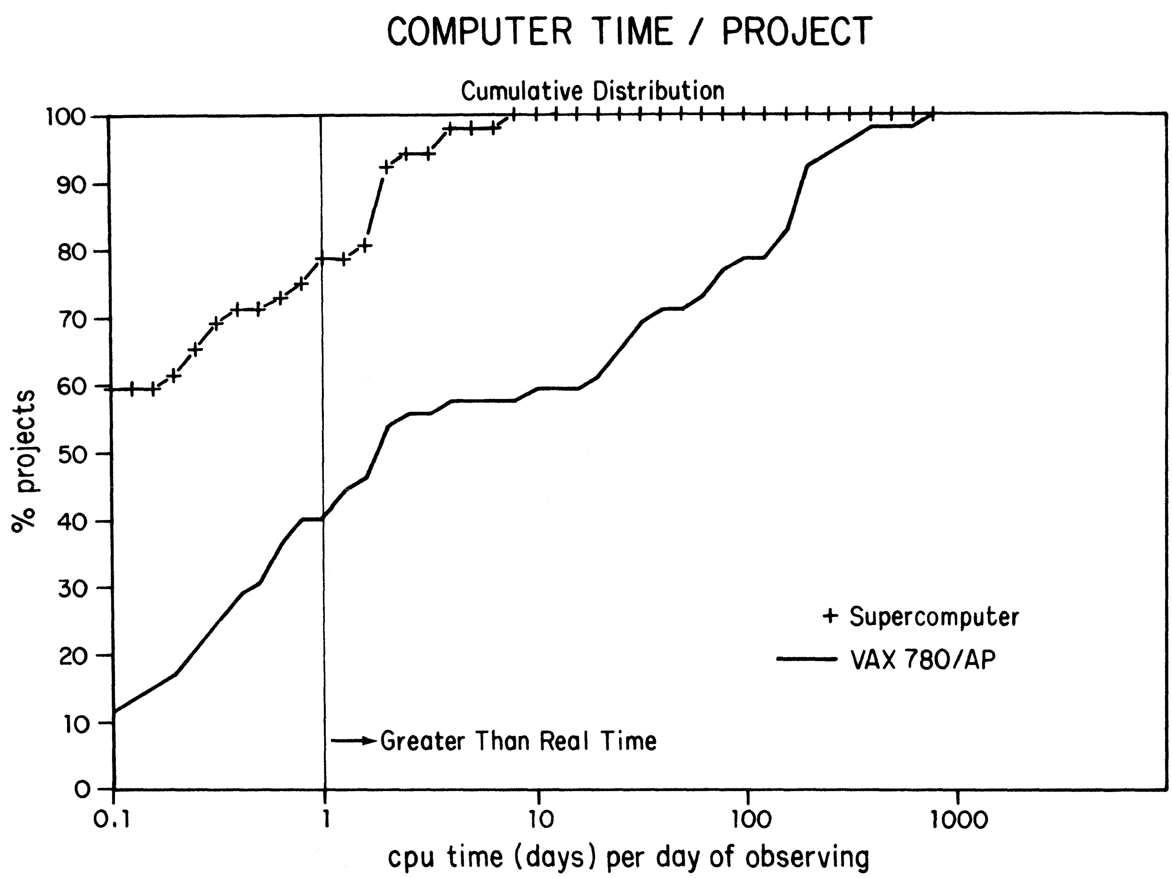

Figure 1-5

Cumulative distribution of the computer turn-around time per project.

A Supercomputer for Radio Astronomical Imaging

The formation of radio images from measurements of complex fringe visibility (synthesis mapping) is not only CPU intensive but also I/O intensive. Ability to deliver good I/O performance is just as important a criterion as CPU performance in evaluating computers for image processing applications. Datasets are large (as large as 1 GByte per user in worst cases); therefore, adequate disk space directly attached to the computer system is critical. A computer facility which does not have high performance tape drives directly attached to the computer will be at a disadvantage when transferring massive files between parts of the system. A computer which provides a large amount of fast memory (either main or auxiliary) for the temporary storage of files will have a special advantage in mapping applications, because numerous scratch files are created. Those that fit into the memory can be transferred at rates an order of magnitude faster than those resident on disk drives. NRAO's recent experience using supercomputers at Los Alamos Scientific Laboratory and at Vector Production in Los Angeles have strongly reinforced these conclusions regarding the importance of I/O and fast memory for image processing.

Provision of I/O bandwidth and compute cycles sufficient to operate the VLA at full efficiency will not be enough to make the new computer facility successful. If the system is to have its full impact on the amount and the quality of science done at the VLA, the astronomer must also be provided with display tools adequate to enable him to digest and interpret his data as easily and efficiently as possible. Because decisions on the processing steps to use are based on an assessment of the data, the current image and the astronomical objectives, this system must be interactive.

1.5 Why A Supercomputer ?

The scientific needs discussed above can be summarized as two computer requirements: 1) the new system must be capable of a sustained rate of 60 Megaflops, with high performance input/output and interactive digital image displays, and 2) it must be capable of analysing the majority of the proposed observations in a time comparable to the observing time. NRAO has considered three different technical approaches which might supply the needed computational power and the needed response time: a special purpose machine, multiple machines supplemented by array processors, and a supercomputer.

There is a considerable amount of research in progress which is directed toward the goal of creating very powerful computers from large numbers of relatively inexpensive integrated circuits operating in highly parallel modes. NRAO is not in a position to participate in the development of such experimental computers, and therefore these machines have been ruled out of NRAO's current plans. In its 1982 report, NRAO's Computer Advisory Group recommended that "... the plan should not depend on being at the leading edge of computer... technology", and in 1984 they reiterated this recommendation.

A major problem with multiple computer/array processor configurations is that they generally must be acquired from more than one vendor. Because these subsystems do not interface well, the performance of combined CPU plus array processor systems is limited. While a configuration of multiple independent CPUs, each with an array processor, appears to be capable of delivering the desired total computing throughput, it could not provide the necessary short turnaround time because it would not have concentrated power. A unified configuration involving multiple CPUs and array processors operating in parallel might achieve the short turnaround time which is needed, but its inflexibility would seriously inhibit the development of new data processing techniques. NRAO has had extensive experience with this multiple machine plus array processor approach; a valuable outcome of this experience is the lesson that was learned concerning the need for flexibility, and the very high price of complexity [ref 11]. NRAO believes that future dependence on array processors would impede the development of the VLA and synthesis telescopes in general.

Supercomputers are distinguished from other large, but conventional, machines by their logical organization (their "architecture"). For the purpose of this discussion, the defining feature of a supercomputer is the presence of one or more pipelined arithmetic units which are used to implement vector instructions. The combination of pipelining and parallel operations produces a considerable increase in computing speed, but is only practical when the same operation is to be done to a relatively large number of values. In general, synthesis mapping software already meets this basic requirement.

A supercomputer system has four great advantages for NRAO over multiple array processors (the other plausible technical option):

i) Its high performance allows the largest processing calculations to be performed in a reasonable amount of time. It thus allows fuller exploitation of the observational data and evaluation of alternative processing procedures;

ii) It provides a unified environment of software, CPU, arithmetic pipelines, memory, and I/O devices, all assembled and supported by a single vendor with the aim of providing the highest performance in a total system. This is important for synthesis mapping because of the heavy I/O and memory requirements;

iii) It supports high level languages using compilers which automatically utilize the pipelined vector hardware. This is a fundamental advantage, because it means that even experimental programs written by any user gain a substantial part of the pipelining advantage;

iv) It provides a direction for growth to allow for future needs; this point was especially stressed by NRAO's Computer Advisory Group at their 1984 meeting.

NRAO has concluded that acquisition of a supercomputer is the most direct and straightforward approach to its data processing problem, since the required sustained throughput of at least 60 Megaflops cannot be achieved by any other hardware, except through a great deal of replication. Furthermore, the dominance of vector operations such as the Fast Fourier Transform in the VLA task mix ensures that the special properties of a suitably configured supercomputer can be exploited effectively to complement the existing superminicomputer systems in the universities. NRAO's Computer Advisory Group (see Appendix B) said: "The committee endorses the plan to acquire a supercomputer system... We believe it is the most attractive option currently available and absolutely essential for the prosecution of the science."

1.6 The New Astronomy

High resolution continuum imaging of individual discrete objects such as Cygnus A (Fig. 1-2) and Cassiopeia A (Fig. 1-3) will provide a moderate load for the proposed supercomputer. A few such objects are being fully analysed with the present resources but major astronomical advances will come only when such high quality images can be routinely produced.

The VLA has been able to push the detection limits for cosmic radio emission far below anything previously attainable. This great sensitivity has allowed the VLA to exploit new astronomical fields such as stellar radio astronomy and radio astronomy of asteroids and planetary satellites (in addition to the traditional deep source surveys used to study cosmology and the evolution of the universe). At these low flux levels the background of stronger radio sources in the field of view create confusing sidelobes unless a corrected image covering the entire field is computed. Analysis of such projects is so time-consuming with present computer resources that sensitivity or frequency coverage is commonly sacrificed.

The extension of the operation of the VLA to lower radio frequencies will open up entirely new fields of astronomical research [ref 12]. Before the invention of the self-calibration algorithms described in Section 1.1, the ionospheric seeing made this part of the radio spectrum useless for high angular resolution imaging. However, this type of observation generates a severe computational load since the full field of view, containing many hundreds of sources, must be imaged and the required ionospheric corrections are a function of both position in the image and time. Although somewhat more sophisticated algorithms still have to be developed to handle this case, there is no doubt that computation of high quality images will be possible [ref 13]. These low frequencies extend the range of energies over which synchrotron radiating plasma can be detected and are particularly interesting because the older populations of relativistic particles become observable, and because the spectral curvature can be used to determine their age.

Another observing strategy which is too computationally intensive to be pursued in a serious manner at present is the "snapshot" style of observation, in which either large areas are covered with a mosaic of images, or the structures of large samples of objects are measured. The VLA is capable of generating thousands of these snapshots in a day of observing, each reaching objects fainter than were observable with most pre-VLA telescopes. A few highly successful projects have already been attempted; for example, the search for gravitational lenses and the fascinating high resolution mosaic showing all the unexpected detail in the radio emission from the central region of our Galaxy. Since new classes of radio emitting astrophysical objects are still being discovered, this mode of observation is sure to be very fruitful.

The highest computer loads are generated by spectral line observations, especially when large objects are imaged with both high spatial resolution and high velocity resolution. The most difficult of these observations occur when spectral lines of HI, OH, H_2CO, and NH_3 are seen in absorption against a background continuum radio source. In this case the signal-to-noise ratio allows detailed studies of the finest scale structure known to be present in the interstellar medium.

At the other extreme are observations of maser sources. Although these regions are individually too compact to resolve, they are spread over a large volume in space and velocity; this implies that enormous data cubes are needed to find them. The VLA could be very successful in searching for more of the recently discovered extragalactic Mega-masers (masers millions of times more powerful than any in our Galaxy), but these searches must cover the entire velocity range of a rotating galaxy and consequently involve very large data cubes.

Perhaps the largest predictable expansion in scientific productivity will come from the ability to make routine observations of the dynamics of galaxies and clusters of galaxies. These will become quite modest projects for a supercomputer (10 minutes CPU time, compared to days at present) so that research into galaxy dynamics, and hence mass distributions, will no longer be limited to a few objects. Given realistic mass estimates for larger samples,including many different types of galaxies, astronomers will be in a much better position to pursue the problem of the missing mass in the universe.

For every class of observation considered, the "tunnel vision" imposed by the current computing resources seriously reduces the chance of accidental discoveries of unusual or unsuspected radio emission outside the source of interest. Such serendipitous discoveries have played a crucial role in the development of astronomy (cf. Harwit's "Cosmic Discovery" [ref 14]). No one can predict the new phenomena which might be discovered in the future, but even today there are examples of discoveries which were missed in the past. The VLA was used several years ago to observe a field which contained the giant Galactic center flare which was later discovered by Yusaf-Zadeh, Morris and Chance [ref 15]. This fascinating feature was not noticed in the earlier experiment because the observers minimized computing needs by limiting the field of view to just meet their objectives. In this case, we were lucky; a subsequent observation made the discovery. How many other discoveries have already been missed due to computer limitations?

CHAPTER 2

THE PROPOSED FACILITY

2.1 Hardware Configuration

The core of NRAO's computer plan for 1985-1990 is the acquisition of a supercomputer. A schematic of the proposed facility is shown in Figure 2-1. This facility will be based on a central processing unit (CPU) capable of a peak performance of more than 200 Megaflops, which is intended to ensure that the system can deliver sufficient sustained performance (60 Megaflops) to meet NRAO's needs. The system will include large amounts of both fast access memory and disk storage (290 Megabytes and 10 Gigabytes, respectively).

The fast memory will contain large image arrays and will also substitute for disks for the storage of temporary files (sustained transfer rates from memory to memory will generally be at least ten times faster than from disks to memory). The large disk space is required for the large datasets which the speed of this machine will permit it to process. The total disk capacity must be several times that needed by any one user process (more than 1 Gigabyte per user for large problems) in order to permit the staging of new data from tape to disk, the processing of multiple job streams, and the staging of results from disk to tape and display devices.

Many operations involved in processing synthesis array data involve performing relatively few arithmetic operations on a large amount of data, more data than can be kept in fast memory. Thus data and intermediate products must be moved from the CPU to disk and back, often repetitively. It follows that the disk-to-CPU transfer rate of the supercomputer must be proportional to the speed of its arithmetic pipelines if its full performance potential is to be achieved. If the individual I/O channels are not fast enough, the total bandwidth required can generally be achieved with multiple concurrent I/O channels.

The small computers designated as peripheral interface computers serve the traditional role of front-end machines (interfacing terminals, special peripherals, workstations and communications systems to the supercomputer). The facility will have a range of user "workstations" with capabilities from small to large, roughly corresponding to the spectrum of problems. Because the workload has an exponential character (cf. Figure 1-5), the facility will have more medium workstations than large ones, and more small stations than medium ones. The large stations will utilize the peripheral interface machines, medium stations will consist of NRAO's existing supermini systems with additional peripherals, and the small stations will consist of supermicro workstations. Supermicro workstations can support image and graphics displays and thus provide both interactive user interface to the supercomputer for modest image analysis, and a high quality programming environment.

Even though supercomputer models are not numerous, several choices still exist. Among U.S. manufacturers, Cray Research offers several models and Control Data Corporation (CDC) offers its Cyber 205. Other U.S. manufacturers, such as IBM, may enter the supercomputer market soon. Among Japanese manufacturers, Fujitsu and Hitachi either have announced supercomputers or are expected to do so in the immediate future. An example configuration is based on a Cray X-MP; details are given in Appendix C.

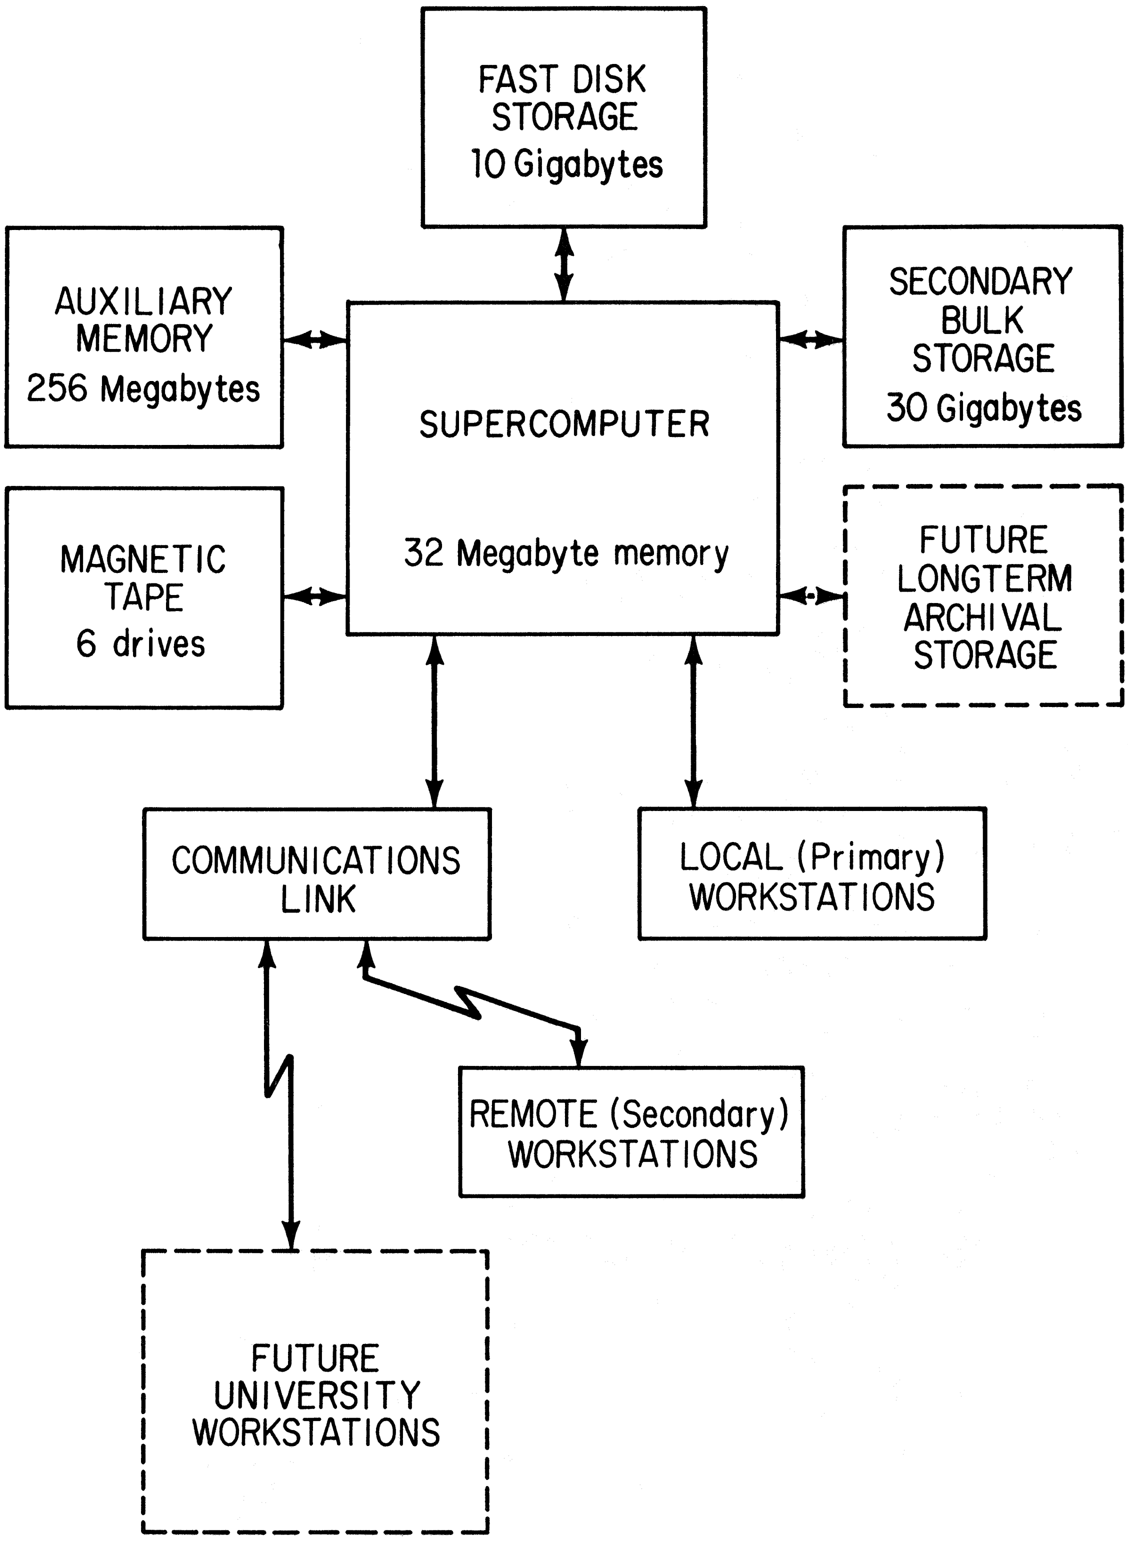

Figure 2-1

Schematic of the Proposed Facility

2.2 The Astronomical Image Processing System (AIPS)

NRAO is in an excellent position to use a supercomputer because a significant effort has already been expended during the past few years to make truly portable a major portion of NRAO's synthesis mapping software, the AIPS system. Thanks to this prior effort, NRAO is in a position to make effective use of a supercomputer for synthesis mapping on a short time scale.

AIPS has been installed on several different computer hardware architectures other than the VAXes on which it is used most widely, and under more than one operating system on several of the architectures. The portability of AIPS is due to the use of a portable dialect of FORTRAN, modular coding, isolation of machine- and operating system-dependent code, and parameterization of system and device characteristics. As far as possible, NRAO intends to maintain the ability of AIPS to run on all classes of computers, ranging from supermicrocomputers to supercomputers. All implementations of AIPS will continue to have the same user interface, with the application programs providing the same functionality and differing only in performance with the computer architecture.

Although array processor microcode is a critical element in the success of AIPS in its superminicomputer implementations, this microcode has been carefully confined to a limited number of modules, which makes it possible to replace these modules when implementing AIPS on new kinds of vector processing hardware, such as the supercomputer. In general, NRAO's synthesis mapping problems are suitable for supercomputers [ref 16]; NRAO's extensive experience with array processors provides a head start in vector computing. Fast Fourier Tranforms, which are a substantial part of NRAO's computing burden, are particularly well suited to vector computers.

2.3 Interactive Image Processing

As has been stressed already, the user needs access to his data and his images in forms suitable for extracting the maximum amount of science from the raw numbers. Digital image displays are a crucial part of the man-machine interface and are vital to the success of NRAO's plans for computing. Based on past experience with image displays, and the perceived needs of NRAO users, NRAO has planned for a range of capabilities in display systems. Large display workstations comprise a major portion of the requested display funds and will provide vector graphics, multi-plane color image displays, and hard-copy devices on-line to the supercomputer for highly interactive imaging tasks. Medium workstations are similar except that the vector graphics devices are replaced by simple raster scan devices; these workstations are similar to the display facilities currently in use at NRAO, but have the ability to handle larger images. The small workstations are designed to be used off-line, for example, for examining images already produced in final form. Appendix C details the display facilities required and the costs incurred to achieve the desired capabilities.

The two important components of data display are imaging (gray-scale pixels) and graphics (lines, points, and regions); both may be in monochrome or in color. There is an obvious difference in application: identical data may be displayed on the former as a "photograph", or on the latter as a graph or contour map. The two forms of display are complementary. Two-dimensional line graphics, e.g. contour maps and cross-sectional displays, do not convey much information about a large image but they are essential to provide quantitative information about smaller regions in the image.

Foremost in importance is the ability to display the final images in gray-scale form. The display of a single image of 512 x 512 pixels, which is the typical size used in the existing systems, presents no problems. Image display devices with storage for 2048 x 2048 pixels using monitors with 1024 lines are available. The proposed image display facilities will allow for the flexible display of images of this size. The facilities will give the astronomer an instant overview of the contents of his data, as well as enabling him to lift out details.

The VLA is capable of producing three-dimensional images of sizes of 2048 x 2048 x 512 pixels or larger. An image storage device currently under development at NRAO has the capacity to display slices through a 512 x 512 x 512 cube at a cinematographic rate, but for the longer term, even more sophisticated hardware for the selection and display of such data is required. Three-dimensional graphical display systems are available today and are proposed for the large workstations.

2.4 Communications Needs

The ideal communication system would allow NRAO to make available the full capabilities of its supercomputer facility, including interactive image processing, to all NRAO users, independent of their location. The use of interactive image processing is critical in NRAO's application, since both the specification of parameters to optimize image enhancement and the choice of astronomical analysis procedures are dependent on inspection of the images throughout the reduction process. In order to estimate the digital communication bandwidth which is required to support interactive image processing, NRAO distinguishes two classes of service: "local" and "remote".

Local users obtain responsiveness in their image display and graphics hardware which is limited only by the data rates of the local area network which interfaces the supercomputer to the peripheral interface machines and workstations. In particular, the largest "cubes" of spectral line data can be transmitted and displayed quickly for local users.

The performance goal for remote service must be a realistic compromise, limited by both the current economics and the current technology of communications links [ref 17]. NRAO has adopted a peak burst rate of 400 kilobits/sec as its minimum performance goal; such links will be able to write 1024 x 1024 8-bit images into remote digital image displays in 20 seconds. NRAO expects that a rate of 100 kilobits/sec will still be useful in the near term, but 10 kilobits/sec definitely will not. The example system for budgetary estimates assumes that NRAO's central research facility in Charlottesville, VA, will be the initial remote site for the supercomputer facility.

Because the supercomputer will be located close to the computers used to control the VLA telescope (see Section 3.1), it will be possible to link the two devices for a direct transfer of data in real-time. The present plan does NOT include that connection as a critical element, but it is a possible future option. The data rates which would apply are an average rate of 80 kilobits/sec, and a peak rate of 320 kilobits/sec. The chief attraction of such a connection would be to implement interactive observing modes for a subset of VLA experiments. The remote level of service defined above is sufficient both to transmit the raw data for such a subset to the supercomputer and to return the processed images to the observer at the telescope.

It is desirable that NRAO's users be able to have interactive access to the supercomputer from their home universities. NRAO will develop its communications support for the remote class of service in such a way that it can be replicated and extended to give its university user community access to the supercomputer if and when the economics permit.

2.5 Archival Storage Needs

If the VLA were unconstrained by computer processing power, the projected mix of observing programs would produce an estimated output data rate of 800 Megabytes per day. In addition, several versions of the final images will be produced; both the original data and the final images must be archived. Table 2-1 summarizes this situation. Even with compression of the images by a factor of four, the total archive storage requirement will be 2.2 Gigabytes per day. The VLBA will at least double this estimate. A required yearly net storage capacity of 1600 Gigabytes results.

Table 2-1
VLA Archiving Requirements

Data Type	Current Gbytes/day	Estimated Gbytes/day	Maximum Gbytes/day
Raw data	0.2	0.8	3.2
Images	0.4	1.4 *	5.6 *

* Assumes data compression by a factor of four

The VLA currently archives all observed raw (uncalibrated) data using conventional magnetic tape. The physical space required to store thousands of reels of tape is already a problem for the VLA. NRAO has developed its own high-density archiving system based on video cassette recorder technology, but only as an interim solution. NRAO's plans assume that VLA data will be archived on optical disks within the next few years, and that the capital cost of the appropriate hardware will be small enough to be absorbed by the capital investment funds specified in the supercomputer operating budget. The costs for the actual recording media are included separately in the estimated annual operating costs [ref 18].

2.6 Physical Plant

A supercomputer requires a great deal of power and air conditioning, as well as structurally sound floors and damage protection systems. The computer rooms must have adequate space for the peripheral interface computers, other peripherals, operations and maintenance personnel, data archive vaults, and users. The physical security of the valuable equipment must be assured. NRAO does not currently have a computer room with all of the required facilities at any of its sites. Therefore, new building space and environmental equipment will be planned in conjunction with the planning for the combined VLA/VLBA array operations center in Socorro.

CHAPTER 3

OPERATING THE SUPERCOMPUTER

There are several assumptions built into this discussion; if any of these is not true, more manpower will be needed than discussed below, and capital and/or operational costs will be substantially higher:

i) NRAO will own and operate its own supercomputer facility for the processing of data from its synthesis telescopes;

ii) NRAO will cease to operate several of its present synthesis mapping facilities. Three existing superminicomputer image processing systems will be retained to serve as workstations for the supercomputer and as stand-alone image processors;

iii) The supercomputer will be located in the vicinity of the VLA, its primary data source, and will be housed in an extension of one of NRAO's existing buildings, either the VLBA headquarters building or the VLA control building.

NRAO intends to operate the supercomputer center with fewer personnel than is customary at other such installations. The users of the facility will be astronomers who will be using existing software packages. Software development will mostly be done by NRAO staff programmers, mathematicians, and scientists. Thus NRAO expects that it will not need as many user assistants, consultants, and operations support personnel as it would if it were operating a general purpose computing center. Table 3-1 below gives estimates of the manpower required to operate the proposed supercomputer facility; NRAO's synthesis mapping computing staff, now 33, will increase to 50 persons. Some of the existing manpower will be reassigned from existing processing facilities. The manpower plan includes a total of four mathematical analysts (an increase of two) since research into image processing algorithms is an important component of the plan to acquire and operate a supercomputer for synthesis mapping.

Table 3-1
Manpower Requirements

Area	CV now	VLA now	Total now	Needed	Increase
Management	1	2	3	4	1
Secretarial	0	0	0	2	2
Systems Analysts	2	1	3	6	3
Prog. Analysts	6	6	12	12	0
New Project Progs.	0	0	0	3	3
Math. Analysts	1	1	2	4	2
Hardware	1	2	3	4	1
Operation Support	2	2	4	7	3
Operators	5	1	6	8	2
	==	==	==	==	==
Total	18	15	33	50	17

3.1 NRAO's User Community

NRAO's charter is to support basic research in radio astronomy, not computer science. NRAO exists to serve a large user community of visiting astronomers who submit proposals for use of the NRAO telescopes (in this case the VLA and VLBA). After successful peer review, telescope time is allocated and logistic support provided for acquisition and reduction of data (the observational research needs of NRAO's scientific staff are served on the same terms). Accordingly, as pointed out earlier, NRAO will operate the supercomputer chiefly in support of the aperture synthesis image processing needs of these astronomers.

A Supercomputer for Radio Astronomical Imaging

NRAO's users are a diverse group widely distributed throughout the USA and the rest of the world. In 1984, 480 separate research projects were conducted at the VLA by 600 observers from 140 different institutions. These users normally come to the VLA to set up their observational program and to use the NRAO computer facilities for part or all of their data processing. During this phase the NRAO scientific and computing staff provide substantial support in the form of advice on processing strategies and general help in using the software packages provided, and provide intensive help for novice users. In turn, the NRAO staff receives feedback on the astronomical needs, and on operational or technical limitations of the facilities. If the observing programs are straightforward and the data rates allow it, VLA facilities and staff support some remote observing and processing via telephone lines. At least 20 of the user institutions now have their own superminicomputer-based systems, running the Astronomical Image Processing System (AIPS) software package supplied by NRAO, which are used for the later stages of data reduction.

To maintain good communication between the user community and the NRAO research and engineering staff, and to provide observers with interactive access to the power of the supercomputer, the proposed facility should be located near the VLA operations center.

CHAPTER 4

COST ESTIMATES

The tables below summarize the estimated costs for NRAO to acquire and operate a supercomputer facility. The estimates are based on the example configuration; more details are discussed in Appendix C. All prices are list prices (in 1985 dollars), and do not include any discounts.

Table 4-1
Capital Investment

Supercomputer and peripherals	$15,000,000
Display Systems	1,400,000
Physical Plant	1,000,000
Communications equipment	400,000
Contingency (6%)	1,000,000
	===========
Total	$18,800,000

Table 4-2
Incremental Annual Operating Costs

Capital investment	$ 1,000,000
Personnel costs	800,000
Supercomputer operations	700,000
	===========
Total	$ 2,500,000

CHAPTER 5

CONCLUSIONS

As a result of VLA enhancements, a million times more data is being produced than was contemplated when the instrument was designed. An additional investment of a small fraction of the capital cost of the VLA will buy enough computing capacity to remove the processing bottleneck and make the VLA a new instrument capable of opening up rich new frontiers in astronomy.

For the first time the full capacity of the hardware can be exploited and many scientific projects not currently feasible will be started. Snapshot surveys of large areas will make new discoveries commonplace - discoveries like the giant Galactic Center flare. It will become possible to analyse large statistical samples of objects of the type currently confined to single sources such as the radio galaxy Cygnus A (front cover). Almost all types of spectral line observations will expand dramatically. As the measurement of the dynamics of galaxies becomes a straightforward observational project, the tortuous route to the elusive missing mass may become easy. Most exciting of all are the projects as yet unconceived because they lie in directions where the spectre of never-ending data processing has turned away many good researchers.

APPENDIX A

SCIENTIFIC REVIEW PANEL

This committee of radio astronomers met in Charlottesville in October 1983 to review the NRAO computer plans, including the concept of a supercomputer for synthesis mapping at NRAO.

Visitors:

Dr. J. Taylor, (Chairman)	Princeton University
Dr. K. Johnston,	United States Naval Research Laboratory
Dr. M. Kutner,	Rensselaer Polytechnic
Dr. P. Palmer,	The University of Chicago
Dr. M. Reid,	Harvard-Smithsonian Center for Astrophysics
Dr. R. Wilson,	Bell Telephone Laboratories

NRAO staff:

Dr. A. Bridle
Dr. R. Brown
Dr. B. Clark
Dr. D. Hogg
Dr. H. Liszt
Dr. C. Walker

This committee's recommendation (October 1983) regarding NRAO's computing plan was:

> "... the [VLA] is known to be capable of addressing important scientific problems for which, at present, sufficient number-crunching capability is unavailable. We are convinced of the need for computing capacity at least in the 'small supercomputer' range... construction of the VLBA will, of course, add to this need. We urge NRAO to proceed with a detailed plan for achieving this kind of computing capacity in the most expeditious manner possible."

APPENDIX B

NRAO COMPUTER ADVISORY GROUP

The NRAO Computer Advisory Group was convened at the request of NRAO to review and comment on NRAO's computing plans. The group met at the VLA site in March, 1982 and in Green Bank in September, 1984.

Dr. K. King	(Chairman) Vice Provost for Computing, Cornell University
Dr. A. Brenner	Head, Computing Department, Research Division, Fermi National Accelerator Laboratory
Dr. W. Brouw	Director, Netherlands Foundation for Radio Astronomy
Dr. P. Green, Jr.	Member, Corporate Technical Committee, Thomas J. Watson Research Center (IBM)
Dr. P. Kunz	Staff Physicist, Stanford Linear Accelerator Center (SLAC)
Dr. H. McDonald	Member, Technical Staff, Computer Technology Research Laboratory, Bell Telephone Laboratories
Dr. S. Patterson	President, Technical Computation Systems, Inc.
Dr. P. Patton	Vice President and Program Director of Parallel Processing, Microelectronics and Computer Technology Corporation (MCC)
Dr. F. Ris	Senior Manager of Computation-Intensive Systems, Computer Science Department, Thomas J. Watson Research Center (IBM)

A Supercomputer for Radio Astronomical Imaging

The following text is quoted from the relevant parts of the summary section of the Advisory Group recommendations (1984) regarding NRAO's long-range computer plan:

"After reviewing the documents containing information on requirements, options, and the plan and hearing the presentations, the committee came to the following conclusions:

- "The committee endorses the plan to acquire a supercomputer system... We believe it is the most attractive option currently available and absolutely essential for the prosecution of the science."
- "We believe it is important for the NRAO to gain supercomputer experience as quickly as possible. We recommend that the NRAO go to the NSF now to get time at one of the National Supercomputer Resource Centers. The NRAO could use of the order of 10 hours in 1985 and hundreds of hours in 1986-87."
- "The committee believes that the best location for the proposed NRAO supercomputer is Socorro with a high bandwidth data link to the VLA. We believe that the supercomputer should be located close to the principal data source."

APPENDIX C

THE EXAMPLE CONFIGURATION

The example configuration for budgetary estimates consists of a Cray X-MP/14 processor (single processor), equipped with substantial external memory, high-speed input/output channels, two peripheral interface computers supporting a communications link to at least one other site, 10 Gigabytes of on-line disk storage, a large mass storage subsystem and high-performance digital image display workstations. This facility would run under Cray's COS operating system, and the NRAO AIPS application software package would provide the main software tool for the synthesis mapping applications. The example configuration is illustrated in Figure C-1.

NRAO is confident that this configuration would satisfy its needs for performance and operating environment. The implication of this example is not that a Cray X-MP is the best one for NRAO's needs at this time, but rather that the capital and operating costs of configurations provided by competitive vendors would be approximately the same as for this example system.

At full speed, the Cray X-MP will execute 220 Megaflops but it is a rare application that can make use of that speed, except in very small bursts. After the initial setup, a series of 2048-point complex Fast Fourier Transforms can be performed on this machine at the rate of 1 millisecond each or roughly 130 Megaflops. This is the upper limit of this machine under the VLA workload.

The 256 Megabyte Solid-state Storage Device (SSD) is needed to augment main memory because the type of work to be conducted on this machine is heavily dependent upon efficient I/O. A single 2048 x 2048 image would occupy all of the main memory, and some current algorithms for image enhancement involve as many as ten different arrays of the same size as the image. To some extent the SSD acts as an extension to main memory, but it is not addressable at the word level. The SSD is read and written as if it were a super-fast disk drive. The Cray X-MP also includes 1 Megaword (8 Megabytes) of buffer memory which forms a common pool of memory space accessible to the I/O processors of the system. The main purpose of this memory is to minimize idle CPU time caused by I/O blockage, but it can also be used as another type of fast auxiliary storage.

Almost all of the disks are connected directly to the Cray X-MP rather than to the peripheral interface computers. The Cray model DD-49 disk units have a maximum capacity per drive of 1,200 Megabytes; 8 drives would be needed to provide the 10 Gigabytes of fast access disk storage required by the VLA. The peak transfer rate of these drives is 10 Megabytes per second. More than 30 Gigabytes of additional mass storage space with somewhat slower access are provided. NRAO would prefer all of this mass storage to have fast access if such a configuration is available at a competitive price.

Table C-1

Supercomputer

Item	Cost
Supercomputer, including 32 Megabytes of memory, I/O interfaces, software, installation, etc.	$ 7,250,000
Additional 256 Megabytes of memory (SSD).	3,300,000
Common file system, including CPU, 35 Gigabyte mass-store unit, staging disks, six tape drives.	1,775,000
10 Gigabytes of on-line disk storage.	1,200,000
Peripheral interface computers (three).	900,000
Interactive software development workstations (twenty).	300,000
Interfaces to communication equipment.	275,000
	===========
Total	$15,000,000

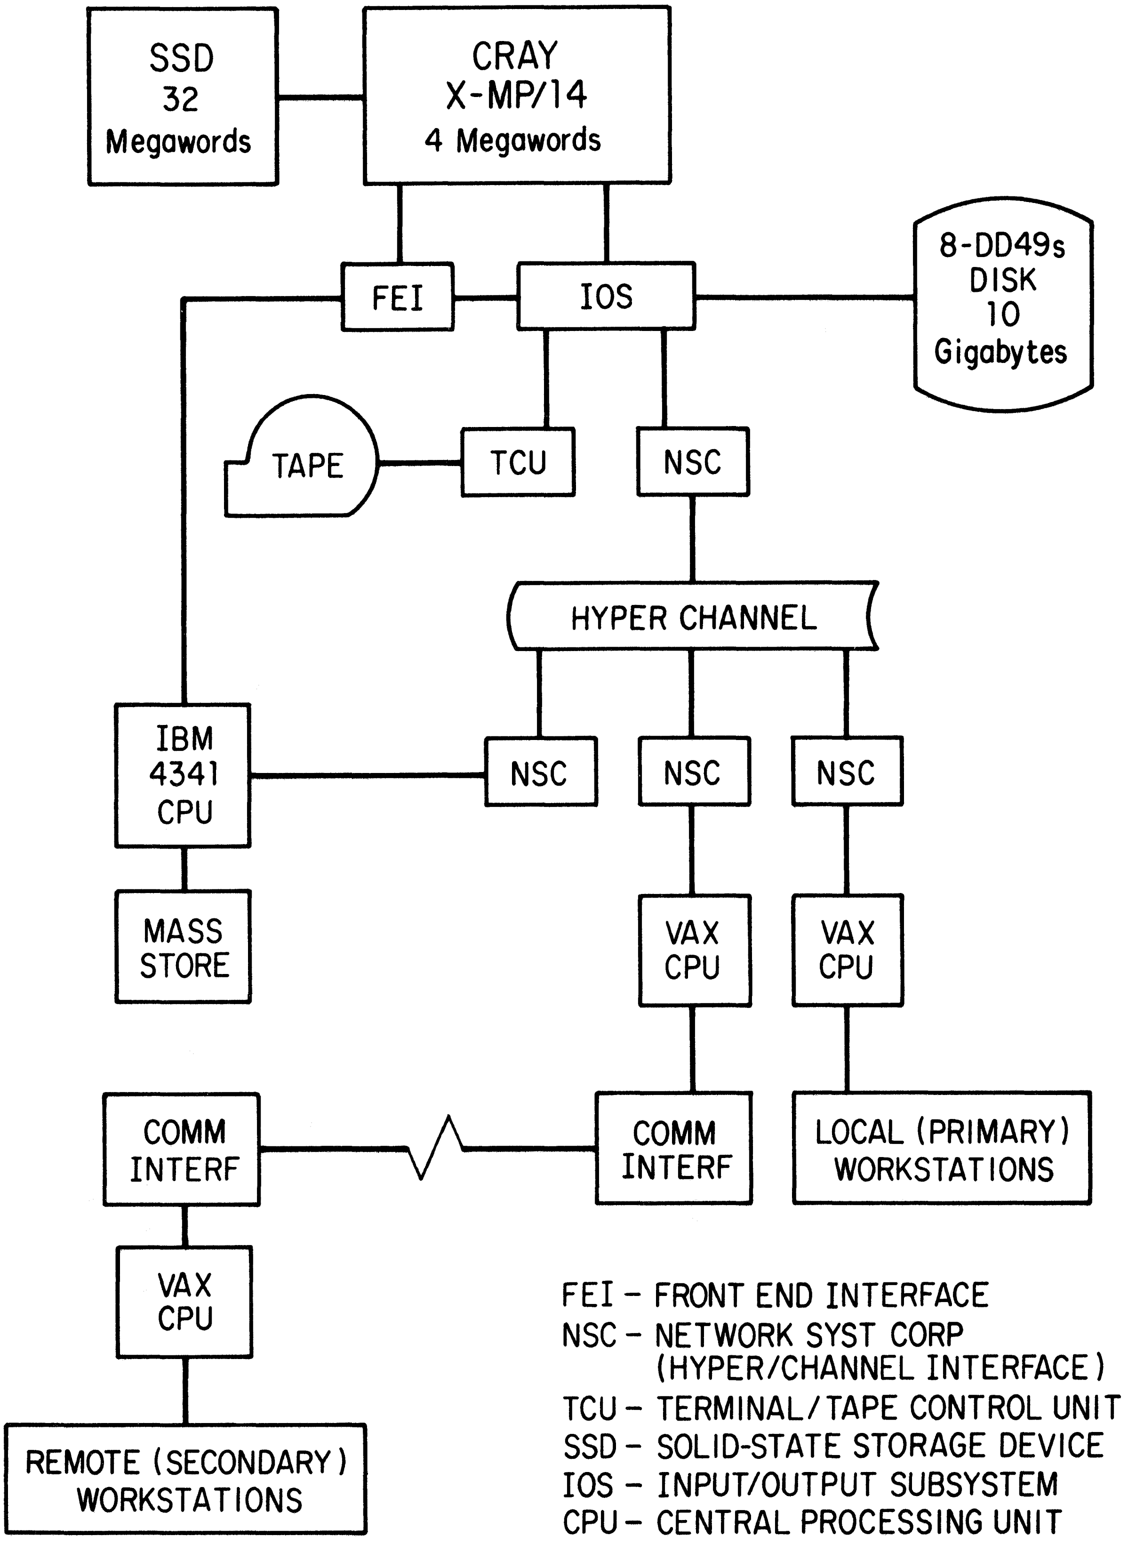

Figure C-1

Cray X-MP/COS Configuration

The large image display workstations in the example system (Table C-2) are International Imaging Systems (IIS) model 600; the three-dimensional graphics displays are Evans and Sutherland model PS-340. These stations will be enhanced as needed with NRAO image storage units (or commercial equivalents) and control panels to make them as powerful as possible. The medium workstations are the IIS model 75; these stations include 1024-line two-dimensional color raster-graphic displays. Several of these image displays will be enhanced with image storage units. The small workstations consist of supermicro-based systems such as those manufactured by SUN Microsystems with monochrome raster-graphics displays. These systems will have local disk drives. Most of the small stations have 512-line digital image displays of limited capability.

Table C-2

Image Display Workstations

No.	Item		Cost
2	Large Workstations at $265,000 each		$ 530,000
	including: 1024-line Image Display	$120,000	
	3-D Graphics Display	80,000	
	Peripherals, Misc.	65,000	
3	Medium Workstations at $140,000 each		420,000
	including: 512-line Image Display	$ 55,000	
	2-D Graphics Display	20,000	
	Peripherals, Misc.	65,000	
10	Small Workstations at $25,000 each		250,000
	including: supermicros with 2-D multi-window graphics; some with 512-line displays.		
2	High Resolution Film Recorders at $100,000 each		200,000
			==========
Total			$1,400,000

The communications equipment for the example system (Table C-3) consists of a satellite link providing a full duplex data rate of 224 kilobits/sec between Socorro, NM and Charlottesville, VA (Vitalink Communications Corp.). As with the supercomputer itself, the costs of other possible approaches are expected to be similar to those shown here.

Table C-3

Communications

Item	Cost
Socorro (dish, interfaces, etc.)	$ 160,000
Charlottesville (ditto)	150,000
Shipping & Installation ($25K/site)	50,000
Planning, Licensing, Misc.	40,000
	=========
Total	$ 400,000

The physical plant expenses (Table C-4) represent the incremental costs of adding the supercomputer facility to either the VLBA control building in Socorro or the control building at the VLA; therefore, land costs are not included. Office space requirements are only for the 17-person increment (cf. Table 3-1), not for all of NRAO's synthesis mapping manpower.

Table C-4

Physical Plant

Item	Area sq ft	Unit Cost	Total Cost
Building space for:			
CPUs and Peripherals	2,500	$100	$ 250,000
Offices, Halls, etc.	3,600	60	215,000
User Display Areas	1,000	100	100,000
Maintenance Shops	900	100	90,000
Data Archive	500	60	30,000
Cooling/Power Eqpt.	500	60	30,000
Cooling (100 Tons)	-	-	115,000
Power Conditioning (150 kVA)	-	-	110,000
Fire Protection (Halon)	-	-	60,000
	=====		==========
Total	9,000		$1,000,000

Operation of a supercomputer will cost much more than even that of a large mainframe, and more than the combined costs for running NRAO's present computers. However, increasing computer power by an order of magnitude costs much less than an order of magnitude more to operate. Table C-5 shows that NRAO's computer operation costs with the supercomputer facility are estimated to increase to just over 5.6 million dollars, 2.5 million more than the current operating costs. Note that the increase due to the supercomputer is partially offset by decreases due to removing several existing systems at the VLA and in Charlottesville.

The increase in the investment budget is based on the assumption that it should continue to be about 8 percent of the capital investment in order to minimize the effects of obsolescence and the growth of the processing load.

Table C-5

Annual Operating Costs

Category	Current	Increment	New Total
Salaries, Benefits, Travel	$1,660,000	$ 800,000	$2,460,000
Capital Investment	500,000	1,000,000	1,500,000
Contract Maintenance	600,000	300,000	900,000
Materials and Supplies	200,000	50,000	250,000
Communications	40,000	150,000	190,000
Power	100,000	80,000	180,000
Archiving	50,000	120,000	170,000
	==========	==========	==========
Total	$3,150,000	$2,500,000	$5,650,000

APPENDIX D

BIBLIOGRAPHY

Copies of NRAO documents are available on request from:

Computer Division
National Radio Astronomy Observatory
Edgemont Road
Charlottesville, VA 22901

1. "The Very Large Array: Design and Performance of a Modern Synthesis Radio Telescope", by P.J. Napier, A.R. Thompson and R.D. Ekers, Proc. IEEE, Vol. 71, No. 11, 1295-1322, November 1983.

2. "Image Restoration (and the CLEAN Technique)", by T.J. Cornwell, in "Synthesis Mapping, Proceedings of the NRAO-VLA Workshop", ed. A.R. Thompson and L.R. D'Addario, Chapter 9, 1982.

3. "Adaptive Calibration", by F. Schwab, Proc. S.P.I.E., Vol. 231, p. 18, 1980.

4. "The Jet and Filament in Cygnus A", by R.A. Perley, J.W. Dreher and J.J. Cowan, Ap. J., Vol. 285, p. L35, 1984.

5. "A National Computing Environment for Academic Research", by K.K. Curtis, National Science Foundation, July 1983.

6. Paper in preparation, by R. Tuffs, R. Perley, P. Angerhofer, M. Brown and S.F. Gull, 1985.

7. "OH/IR Stars Within 50 pc From the Galactic Centre", by A. Winnberg, B. Baud, H.E. Matthews, H.J. Habing and F.M. Olnon, Ap. J. Lett., 1985 (in press).

8. "Need for a Large Computer", by R.C. Walker, VLB Array Memo No. 365, July 1984.

9. "Astronomical Requirements for Future VLA Processing", by R.D. Ekers, E.B. Fomalont and F.N. Owen, VLA Scientific Memorandum No. 150, September 1983.

10. "A Computer Plan for the VLA", by R. Duquet, G. Hunt and R. Burns, VLA Computer Memorandum No. 168, September 1983.

11. "Whatever Happened to the Pipeline?", by R. Duquet, VLA Computer Memorandum No. 172, August 1984.

12. "A Proposal for a Low Frequency Array Located at the VLA Site", by R.A. Perley and W.C. Erickson, VLA Scientific Memorandum No. 146, April 1984.

13. "Relaxing the Isoplanatism Assumption in Self Calibration; Applications to Low Frequency Radio Interferometry", by F. Schwab, VLA Scientific Memorandum No. 151, October 1983.

14. "Cosmic Discovery", by M. Harwit, Basic Books, 1981.

15. "Large, Highly Organized Radio Structures Near the Galactic Centre", by F. Yusaf-Zadeh, M. Morris and D. Chance, Nature, Vol. 310, p. 557, 1984.

16. "Gridding Synthesis Data on Vector Machines", by D.C. Wells and W.D. Cotton, NRAO AIPS Memo No. 33, January 1985.

17. "Communications Technology and Costs", by D. Wells, NRAO Computer Planning Note No. 47, February 1985.

18. "Numerology and VLA Data Archiving", by D.S. Retallack, NRAO Computer Planning Note No. 48, February 1985.