



# ARRAY TELESCOPE COMPUTING PLAN

*Proposal to the NSF  
September 1987*



*The National Radio Astronomy Observatory is operated by Associated Universities Inc., under contract with the National Science Foundation.*

Cover: The radio galaxy M87 and its well-known "jet" seen at three different stages of the image production process:

(upper) The data accumulated by the VLA--sampled over the synthesized telescope on the Plains of San Agustin.

(middle) The image from the sampled data with significant background substructure caused by the nonuniformity of the VLA synthesized telescope aperture.

(lower) The final radio image after correction for known distortions introduced by the imperfect synthesized telescope and for variable distortions caused by atmospheric turbulence.

# ARRAY TELESCOPE COMPUTING PLAN

National Radio Astronomy Observatory

September 1987

The National Radio Astronomy Observatory (NRAO) is operated by Associated Universities, Inc., under contract with the National Science Foundation.

## TABLE OF CONTENTS

|  |    |
|--|----|
| Preface . . . . .  | ii |
| 1. ASTRONOMICAL MOTIVATION . . . . .                               | 1  |
| 1.1 Role of the Computer in Array Science . . . . .                | 1  |
| 1.2 Examples of Computer-Limited Problems . . . . .                | 4  |
| 2. THE REQUIREMENTS . . . . .                                      | 12 |
| 2.1 Astronomical Requirements . . . . .                            | 12 |
| 2.2 Computer Performance Requirements . . . . .                    | 13 |
| 2.3 Other Requirements . . . . .                                   | 15 |
| 3. ESSENTIAL PARTS OF THIS PROPOSAL . . . . .                      | 17 |
| 4. THE PROPOSED NRAO FACILITY . . . . .                            | 20 |
| 4.1 What is the Appropriate Type of Computer? . . . . .            | 20 |
| 4.2 The Proposed Facility . . . . .                                | 23 |
| 4.3 Mass Storage and Archive System . . . . .                      | 26 |
| 4.4 Operations Personnel . . . . .                                 | 27 |
| 4.5 Existing Facilities . . . . .                                  | 27 |
| 5. COMPUTING AT USER INSTITUTES . . . . .                          | 28 |
| 6. SUPERCOMPUTER ACCESS . . . . .                                  | 30 |
| 6.1 Purpose . . . . .  | 30 |
| 6.2 Supercomputer Experience and Overview . . . . .                | 30 |
| 6.3 Future Use of Supercomputers . . . . .                         | 31 |
| 7. SOFTWARE AND ALGORITHM DEVELOPMENT . . . . .                    | 33 |
| 7.1 The Astronomical Image Processing System (AIPS) . . . . .      | 33 |
| 7.2 Optimization . . . . .   | 36 |
| 7.3 Manpower Needs . . . . .                                       | 36 |
| 7.4 Algorithm Research . . . . .                                   | 38 |
| 8. BUDGET . . . . .  | 40 |
| Bibliography . . . . .   | 43 |
| Appendix A: Chronology of NRAO Array Telescope Computing . . . . . | 45 |

## Preface

The Very Large Array (VLA) and the Very Long Baseline Array (VLBA, now being constructed by NRAO) are both Fourier synthesis radio telescopes. Such telescopes achieve the high-resolution imaging capability of a single large antenna by using an array of small antennas to sample the electric field at discrete locations within an aperture that corresponds to the large antenna. An image of the radio sky is then formed by Fourier transforming the correlations among the sampled data. This technique relies heavily on computers. Because the computers form the images, they are a crucial part of the "optics" of the telescope. Since the VLA was built, advances in image processing algorithms have enormously increased the power of Fourier synthesis telescopes. These advances enable us to remove, to a large extent, the main instrumental and atmospheric effects that degrade radio images. As well as forming the images, the computers provide Fourier synthesis telescopes with sophisticated "adaptive optics" that can compensate for the effects of instrumental fluctuations, and for ionospheric and tropospheric "seeing". This software "adaptive optics" is critical for the VLBA, which could not form images without it.

The history of the VLA—from design, through construction, to full-scale use—has been one of rapid advances. The sensitivity of the VLA to weak signals has been markedly improved by replacing its original parametric up-converter receivers with FET and HEMT amplifiers. The imaging speed of the array has been increased by data processing techniques that derive scientifically useful images from brief observations. The image quality—as measured by the dynamic range of the image—has greatly benefited from the development of self-calibration techniques. These advances have dramatically increased the VLA's scientific capability. Both the quality and the quantity of images it produces now greatly surpass the original goals, as Table 1 shows.

Table 1: Development of VLA Imaging Power

|                                     | Goal<br>1969 | Achieved<br>1980 | Achieved<br>1986 |
|-------------------------------------|--------------|------------------|------------------|
| Speed (images per day)              | 3            | 100              | 200              |
| Image Size – Routine (pixels)       | 128 × 128    | 512 × 512        | 1024 × 1024      |
| Image Size – Maximum (pixels)       | 512 × 512    | 1024 × 1024      | 4096 × 4096      |
| Spectral Line Channels (full array) | —            | 8                | 32*              |
| Dynamic Range – Routine             | 100:1        | 500:1            | 2,000:1          |
| Dynamic Range – Maximum             | 100:1        | 2,000:1          | 100,000:1        |
| Maximum Sensitivity (mJy)           | 0.1          | 0.05             | 0.005            |
| Resolution (arc seconds)            | 1            | 0.1              | 0.07             |

\*The number of spectral line channels available with the full array will again increase, from 32 to 512, when a new on-line computer system becomes fully operational in 1987.

Each increase shown in Table 1 has required computing resources beyond those originally anticipated. The growth in demand for computing resources has outstripped our ability to provide them within NRAO's annual operating budgets. (Appendix A gives a short chronology of array telescope computing at the NRAO.) Only a small fraction of the exciting but exceptionally computer-intensive scientific investigations can be supported. Useful observing capabilities that are designed into the array hardware are now withheld from users, solely to avoid overloading the data reduction computers.

In 1983, we began working on a plan that ultimately recommended a dedicated supercomputer for processing VLA data—*A Supercomputer for Radio Astronomical Imaging* (1985). At that time, only a supercomputer offered the needed computational speed. For various reasons, that plan was not realized. This proposal presents a new plan.

Our first step in formulating the new plan was to reassess the computing loads for the VLA and the VLBA, in terms of our present computing capacity. This assessment reveals that proper handling of the projected computing load of both instruments would require almost 90 times our present capacity. About 70% of this load, however, will come from a few extremely computer-intensive spectral line projects, making up only about 10% of the expected observing proposals. We propose that the data for the highest-priority projects of this kind be analyzed at supercomputer centers, especially those operated under the NSF Supercomputer Access Program, using code that is optimized for such extreme problems.

We now believe that the computing needs of the other 90% of the observing load can best be met at the NRAO and at our users' home institutions by taking advantage of recent developments in "mini-supercomputers" and workstations. Since submitting the 1985 plan, the NRAO has installed two first-generation mini-supercomputers (Convex C-1's). The C-1 runs a standard UNIX\* operating system and supports real-time interfaces to image display hardware. These machines are serving us well. Discussions with several manufacturers about their plans for second-generation mini-supercomputers lead us to believe that a modular system based on mini-supercomputers can handle the image processing that should be done at the NRAO for all but about 10% of the expected VLA and VLBA observing proposals. We therefore propose an incremental enhancement of the computing resources at the NRAO. This will allow manageable, multi-year funding and will also let us realize the benefits of new computer technology quickly.

We will also maintain active support of array telescope computing at our users' institutions, through continued development and export of portable software. By using this software to benchmark and certify new machine architectures as they appear on the market, we will continue to obtain experience from which to assist our users' institutions with procurement decisions, and with installation and maintenance of our code. *Our aim will be to preserve the roughly equal balance of array telescope data processing capacity that now exists between the NRAO and our users' institutions, while the total capacity increases by a factor of about fifteen.*

Funding for additional personnel is as important as capital funding if this plan is to succeed. The additional personnel are essential to support development of new and improved algorithms, to maintain and optimize the software, to increase the NRAO support of computing at non-NRAO centers, and to operate expanded computing facilities at the NRAO.

---

\* UNIX is a trademark of Bell Laboratories.

The VLA and VLBA together represent an investment of over 200 million dollars (in 1987 funds). For about 5% of that investment, spread over three years and partially funded through the VLBA construction plan, we can ensure that the full imaging power of both arrays can be exploited by all observers.

# 1. ASTRONOMICAL MOTIVATION

## 1.1 Role of the Computer in Array Science

Because the wavelengths of radio waves are much longer than those of visible light, a conventional monolithic radio telescope would need to have an aperture many kilometers in diameter to equal the angular resolution of an optical telescope. It would be prohibitively expensive and probably impossible to construct. In the 1960's, advances in digital computers made it possible to achieve this and higher resolutions with Fourier synthesis radio telescopes. These telescopes use arrays of many small antennas to sample the coherence of a wave front over a large aperture. An initial image of the radio sky is then formed by Fourier transforming the data in a computer. The computer is thus an essential component of the telescope optics. The technique is similar to measuring the hologram of the radio sky on the ground, and reconstructing the image of the sky in the computer. NRAO now operates one major Fourier synthesis telescope, the Very Large Array (VLA). Construction of a second, the Very Long Baseline Array (VLBA), began in 1985. These arrays (Figure 1-1) are described in detail in [1].

Since its completion in 1980, the VLA has made spectacular contributions to many branches of astronomy. Its sub-arcsecond resolution has allowed it to make images that reveal unprecedented detail in the radio emission from quasars and radio galaxies. The most dramatic among these results are the discoveries of relativistic plasma jets linking active galactic nuclei to giant radio lobes. Studies of normal galaxies have revealed incipient radio jets and giant interstellar bubbles; they have also produced the first detailed radio light curves from supernovae, and have demonstrated the stripping of neutral hydrogen from galaxies in the Virgo cluster. VLA images of our own Galaxy have also contained many 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. The high sensitivity and angular resolution of the VLA have opened up new and unexpected areas of research in stellar astronomy. These include studies of highly polarized bursts from flare stars and close binaries, and "non-thermal" radiation from unusually hot or energetic electrons in stellar winds. The VLA has also been used to probe the delicate balance between inflow and outflow that results in the formation of new stars—winds, both ionized and atomic, sweep outward from their poles while fresh material falls inward onto their equators. The VLA has pinpointed the radio-emitting characteristics of features of the solar transition region and corona—magnetic loops, prominences, coronal holes, new and old active regions—giving new insight into solar plasma processes. Observations of the atmospheres, rings, and satellites of planets, and of outgassing from Halley's Comet, have complemented research done by planetary space probes. Looking to the future, the VLA will provide images of many classes of astronomical objects with resolution comparable to that of the Space Telescope.

The 27-antenna configuration of the VLA was settled on in 1969; the array was then expected to produce images in which the ratio between the intensity of the brightest feature







Figure 1-1. The VLA in its most compact configuration (upper panel) and the Pie Town, NM, VLBA antenna (lower panel).

and that of the faintest "believable" one\* would not exceed 100:1, mainly because of two limitations described below. Today, we can do much better with the same 27-antenna array.

The first limitation arises from incomplete sampling of the arriving wavefront. The sampling is incomplete because both the number of antennas and the observing time are finite. The incompleteness leads to sidelobe patterns in the "dirty" images that are made directly by Fourier transforming the observed data. Fortunately, powerful computing algorithms can remove the effects of these sidelobes, yielding deconvolved ("clean") images of much higher quality. An algorithm known as CLEAN is used most often, but others such as the Maximum Entropy Method (MEM) are playing an increasingly significant role. These algorithms all utilize credible *a priori* constraints on the image, such as its finite size, or (where appropriate) its positivity.

The second limitation is imposed by the earth's atmosphere. Irregularities in the refractivity of the troposphere or the ionosphere, on scales smaller than that of the array, distort the arriving wavefront before it is sampled. The irregularities vary both in space and in time, producing the radio equivalent of what optical astronomers call "bad seeing". In multi-antenna arrays such as the VLA and the VLBA, the number of instantaneous samples greatly exceeds the number of antennas. Thus it is possible to estimate the atmospheric errors and to remove much of their effect. This technique is known as "self-calibration", or "hybrid mapping" [2,3]. It is an off-line analog of the use of deformable mirrors in adaptive optics for imaging at visual wavelengths. It is also closely related to the concept of "structure invariants" in X-ray crystallography. The technique is computer-intensive in that it is iterative, and each iteration must include an application of a deconvolution algorithm.

Use of deconvolution and self-calibration techniques has enormously increased the quality of the best VLA images. Although a dynamic range of only about 100:1 was specified in the 1969 proposal to build a 27-antenna VLA, dynamic ranges of over 100,000:1 have been obtained in a few projects by using the spectral line system for multichannel continuum imaging. This thousandfold improvement has come at the cost of greatly increased data volume and computation time. Dynamic ranges of thousands to one are routinely obtained, even for complex sources containing little or no unresolved structure. At least half of all discoveries made with the VLA would not have been possible without deconvolution and self-calibration, and many results benefit more subtly from them. Self-calibration has been even more valuable to Very Long Baseline Interferometry (VLBI). There, it has permitted true imaging of radio-emitting objects at resolutions up to a thousand times greater than that of the VLA. In 1981, this produced the first direct and unambiguous demonstration of "superluminal" (apparently faster than light) expansion of radio features in a quasar [4].

The cover of this proposal shows how these computational techniques produce and improve an astronomical image, here that of the radio galaxy Virgo A [5]. Three stages in the imaging process are illustrated. The upper panel represents the first stage—sampling the coherence of the wavefront. This produces an array of correlations whose intensities and distribution form the pattern shown. The central panel shows the image resulting from the minimal Fourier-transform processing that satisfies the original design specifications of the VLA. The final panel shows an image that has been deconvolved and self-calibrated to correct for incomplete

---

\* This ratio is a rough—and admittedly subjective—measure of the quantity we describe as the *dynamic range* of the image throughout this proposal.

sampling of the wavefront and for errors produced by atmospheric "seeing". In this image, the extended emission regions are seen to break up into filamentary detail. The detail apparent in the final image dramatically enhances its scientific value, by revealing unexpected filamentation.

In addition to deconvolution and self-calibration, three further factors contribute to the VLA's computer load:

- the new image processing techniques enable observers to produce images much more rapidly and from much less data than had been envisaged, using much shorter observing times than were considered useful ten years ago (the so-called "snapshot" observing mode);
- releasing the full capabilities of the VLA hardware will permit simultaneous recording of up to 512 different radio frequency channels;
- a "mosaicing" algorithm has been developed that allows imaging of larger fields of view than had been thought possible for Fourier synthesis telescopes.

The peak data rate from the VLA can now be much greater than that anticipated when the instrument was conceived. Although the capacity of the computers in use at the NRAO has also increased greatly since that time, it has not increased by nearly as much. Consequently, for many VLA experiments the data output is now intentionally limited. The scientific potential of the VLA is thereby seriously restricted by the available computational resources.

The Curtis report to the National Science Foundation [6] noted that lack of access to adequate computational facilities has inhibited activity in certain scientific fields. We see this happening with synthesis radio telescopes. Important problems in radio astronomy go unstudied simply because the computers cannot process the data load associated with them. Studies that are *just* feasible with the present resources are also adversely affected by the long turnaround times needed for adequate analyses. This discourages VLA observers from experimenting with different approaches to analyzing their data, and from adequately investigating unexpected results.

The process of scientific discovery is fragile. Anything that increases the barrier between the scientist and the observations interferes with this process, and may obstruct it altogether. Limited computer resources are now the primary obstacle to realizing the full potential of the VLA.

## 1.2 Examples of Computer-Limited Problems

We now give some concrete examples of projects that are hampered by current computing limitations, and of the scientific gains that could be made by removing these limitations.

### 1.2.1 High-resolution continuum imaging

There are many hundreds of extended cosmic radio sources that are bright enough to be imaged in detail with the VLA with a resolution of 0.2 arc seconds or better. These objects include quasars, radio galaxies, active and normal galaxies, supernova remnants, and HII regions. Because of computational limitations only a few per cent of all such objects have been observed using the full imaging capabilities of the VLA. These include the two strongest radio sources in the sky, Cassiopeia A and Cygnus A. In both cases the images have resulted in important astronomical advances.

The entire remnant of the supernova explosion known as Cassiopeia A, a region over 300 arc seconds across, has been imaged with the VLA at 0.2 arc seconds resolution (Figure 1-2). A sequence of such images, obtained since 1979 and recently published in *Nature* [7], has brought new insights into the evolution of the supernova remnant. These came from studying the outward motion of the conical structures, revealing that these structures can be interpreted as perforations of the remnant due to dense high-velocity ejecta. The magnetic field orientations in the remnant also suggest the propagation of shocks associated with these perforations. It became practical to compute the deconvolved image shown in Figure 1-2 only when (a simplified version of) our software was run on the Cray X-MP supercomputer at Digital Productions with time made available through the NSF Supercomputer Access Program.

Figure 1-3 shows a VLA image of the powerful radio galaxy Cygnus A, whose filamentary structure and jet [8] were revealed only after many iterations of deconvolution and self-calibration. The discovery of regions with exceptionally high Faraday rotation in this object [9] came from comparing high-quality images of the polarized emission at several observing frequencies. This discovery explains previously perplexing lower-resolution observations, and shows that large-scale magnetic fields are present in the intracluster medium around the radio source.

While the few such projects that have been undertaken with the present resources have produced spectacular results, their data reduction demands have created almost unacceptable congestion in the existing computing facilities at the NRAO. This inhibits many users from attempting similar projects; responsible scheduling of resources also prevents us from accepting proposals for such projects routinely.

### 1.2.2 High-sensitivity imaging

The VLA has pushed the detection limits for cosmic radio emission far below those previously attainable. This great sensitivity allows VLA observers to exploit new fields of radio astronomy, such as studies of galactic stars, asteroids, and planetary satellites. It also allows them to push deeper into traditional areas such as surveys of the distant extragalactic sources used to study cosmology and the evolution of the universe. When studying extremely faint sources, it is essential to remove all the artifacts caused by strong radio sources in the field of view. Correcting for responses to these sources over the entire field of view is so time-consuming, with present computer resources, that sensitivity or frequency coverage is commonly sacrificed. Once again, a few projects of this character have been successfully processed with the available resources, but at the cost of great congestion. The full capability of the VLA for such work cannot be realized until the computing resources are significantly expanded.

### 1.2.3 Low frequencies

Extending the frequency coverage of the VLA to as low as 327 MHz and 75 MHz will open new fields of radio astronomical research [10]. Before the invention of self-calibration algorithms, the ionospheric "seeing" rendered this part of the spectrum useless for high-resolution imaging. Although more sophisticated algorithms still have to be developed to handle the worst ionospheric conditions effectively, there is now no doubt that high-quality images will be





Figure 1-2. A VLA image of the supernova remnant Cassiopeia A; this 4096 × 4096 image was deconvolved using NRAO software in the Cray X-MP supercomputer at Digital Productions in Los Angeles.



7

Figure 1-3. VLA image of the radio galaxy Cygnus A, showing hot spots, previously unknown filamentation throughout the lobes, and the first clear evidence for a knotty jet linking the central radio source to the lobe at the right of the image.

obtainable [11]. Observations at low frequencies will however generate a severe computational load, since the full field of view—containing many hundreds of sources—must be imaged at once and the required ionospheric corrections are a function both of position in the image and of time. Imaging at these low frequencies will extend the range of energies over which synchrotron radiating plasma can be detected, and will be sensitive to steep-spectrum sources that were previously undetectable with the VLA, e.g. millisecond pulsars. In extragalactic sources, the older populations of relativistic particles will become observable.

The feeds and receivers necessary to support low-frequency observing are only now being installed on the VLA. The multichannel observing capability that is crucial to this work is also about to become available. The computing burden of such projects has therefore yet to be felt by other VLA observers. It will be large (Section 2).

#### 1.2.4 High-speed observing

Another observing technique that is too expensive to exploit fully with the present computing resources is the “multiple-snapshot” mode. This can be used to cover large areas of sky with a mosaic of images, or to catalog the structures of large samples of objects. The VLA can generate over a thousand “snapshot” images in a day of observing, each reaching radio sources fainter than were observable with most pre-VLA radio telescopes. Because the very short observing times used for snapshots lead to high sidelobe levels, such projects rely heavily on deconvolution to produce useful images. A few have already been successfully attempted. An extensive snapshot search for gravitational lenses is in progress. So far, it has resulted in the discovery of two new lens systems [12]. A high resolution snapshot mosaic has shown unexpected detail in the radio emission from the central region of our Galaxy. Since new classes of astronomical radio source are still being discovered, this mode of observation is certain to be fruitful.

#### 1.2.5 Galactic absorption lines

The highest computer loads are generated by spectral line observations, especially when extended fields of view must be imaged with both high angular resolution and high velocity resolution. The most difficult of these observations occur when spectral lines of HI, OH, H<sub>2</sub>CO, and NH<sub>3</sub> are seen in absorption against a background continuum radio source; the high signal-to-noise ratio then allows detailed studies of the fine structure of the interstellar medium. At present, only a few projects of this class can be scheduled without overwhelming the available resources. The sixteen-fold expansion of the number of spectral line channels that will be provided by the new VLA on-line computer system will increase the demand for such observing and for galactic emission-line work. This expansion will greatly reduce the ability of our present computers to serve other users while spectral line data are being processed.

#### 1.2.6 VLA studies of maser sources

Observations of maser sources present a different, though equally severe, computational burden. Although maser regions are individually too compact to resolve, they are spread over

large volumes in space and in velocity. Large data "cubes", with two dimensions of direction and one of velocity, must be searched to find them. The VLA is a powerful tool for searching for the recently-discovered extragalactic "Mega-masers" (masers millions of times more powerful than any in our Galaxy); the searches must, however, cover the entire velocity range of a rotating galaxy and consequently involve extremely large data cubes.

VLA observations of masers in our own Galaxy also generate heavy computing loads. For example, the mass distribution in the central region of our Galaxy can be inferred from the distribution of the velocities and the positions of a type of star whose atmosphere emits a sharp radio frequency line from an OH maser [13]. In a project that was undertaken with the present array hardware, the maser stars were found by searching a  $1024 \times 1024 \times 512$  cube for each of twelve different fields. Processing of five days' observing took several years on various minicomputer systems outside the NRAO. When the current restrictions on VLA data output are removed later this year, cubes will be obtainable in each of two orthogonal polarizations, simultaneously, for two different OH lines. The same five day observation will then require about 120 days' dedicated computing time on one of our Convex C-1 image processing systems! The impact on other observing programs would be so great that such a project could not, in practice, be scheduled.

#### 1.2.7 VLBA studies of maser sources

An important project for the VLBA involves determining distances to H<sub>2</sub>O maser sources using the method of statistical parallaxes. This technique compares the line-of-sight velocity dispersion (as determined from Doppler shifts of the maser lines) with the relative angular velocities of the maser spots (determined by relative astrometry over a time baseline of about a year). Such observations have already been painstakingly made for a small number of sources. As a result, for example, the distance to the Galactic Center has been determined to be  $7.1 \pm 1.5$  kpc—somewhat less than the value obtained by less direct techniques [14]. These measurements can be extended, by observing many more sources at time intervals of a few months, to establish the scale of our Galaxy more precisely, and to determine the distances to other galaxies (by observing the maser sources embedded within them). Eventually such measurements may lead to an estimate of the Hubble Constant that is independent of the many assumptions underlying present distance estimates.

Another example of a VLBA project that is extremely computationally expensive is the determination of the magnetic fields in the envelopes of red giant and supergiant stars through the observation of the Zeeman splitting of OH masers. Such an observation involves recording both hands of polarization for the 1612, 1665, and 1667 MHz maser transitions of OH. The computing resources needed would have to cope with generating, analyzing, storing, and displaying twelve cubes of dimension  $\sim 1024 \times 1024 \times 512$ , one for each frequency and each Stokes parameter. Such a task requires resources far exceeding the present NRAO computing capacity, but would return a most useful astronomical result.

#### 1.2.8 Dynamics of galaxies and clusters

Observations of the 1420 MHz spectral line of neutral hydrogen can be used to measure



the radial velocity field of a galaxy, and hence to probe its dynamics and mass distribution. These observations extend what can be done by optical spectroscopy, since the neutral hydrogen covers the galaxy more uniformly and extends much further from its center than the optical line emission—giving a better estimate of the total mass of the system. Typical observations result in a data cube  $1024 \times 1024 \times 64$  channels deep, similar to Figure 1-4 [15].

It is especially necessary when performing this type of dynamical analysis to use data in which all spatial scales are represented. To do this with the VLA, data from several VLA configurations must be combined and processed together. With the current resources it is exceptionally difficult for an astronomer to get both enough disk storage space and enough computer time for this type of project. For example, the basic processing for one galaxy observed with the VLA's B, C and D configurations now takes about two months. The process is inefficient now because large data bases have to be moved on and off the system and the processing disrupts users working on other projects. Because of the processing difficulties only a few galaxies have been observed so far, and none has been analyzed with more than one-third of the maximum available angular resolution.

The VLA will become a much more powerful tool for determining the dynamics of large numbers of galaxies and clusters of galaxies if the total computation time per object is significantly reduced. Research into galaxy dynamics, and hence into mass distributions, would be much more valuable if the long turnaround time for data reduction did not restrict it to just a few objects. Realistic mass estimates for large samples of galaxies, including many different Hubble types, would significantly affect studies of the missing mass in the universe.

### 1.2.9 Serendipity

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 region of interest. Such serendipitous discoveries have played a crucial role in the development of radio astronomy [16]. No one can predict which new phenomena might be discovered in the future, but there are numerous examples of discoveries that were almost missed in the past. For example, the VLA was used to observe a field which contained the giant Galactic Center flare, several years before the flare was discovered [17]. This important and puzzling feature was not noticed in the earlier experiment because the observers responded to computing limitations by restricting the field of view to the bare minimum that met their objectives. In this case, a later observer made the discovery. How many others have been missed because of computer limitations?




Figure 1-4. Stereo three-dimensional solid surface representation of 1420-MHz observations of the atomic hydrogen in the galaxy M81. The "short axis" is radial velocity; the main warp represents the rotation of the galaxy. The three-dimensional effect may be seen either by using a stereoscope or by focusing one's left eye on the left image and the right eye on the right image.

## 2. THE REQUIREMENTS

### 2.1 Astronomical Requirements

Each year, more than 600 observers from 140 different institutions use the VLA for their research. We have used the mix of scientific goals represented by their proposals to project the scientific requirements of the VLA user community for the near future. Our analysis [18] of these requirements consisted of the following steps:

- Listing all astronomical projects which currently use, or are projected to use, the VLA and tabulating the observational parameters that determine their processing needs;
- Calculating the processing capacity that would be required for each type of project so that the astronomical objectives not be compromised by the computer, while accepting other hardware limitations; and
- Estimating the fraction of the VLA's total observing time that would be spent on each type of project, based on a projection of current VLA usage. This estimate is necessary to calculate the time-averaged throughput required, because the amount of processing required varies greatly from project to project.

We first made these projections in 1983, and reviewed them again in 1987. As we anticipated in 1983, the proposal pressure for spectral line work and for the lower frequencies has increased. The needs that were projected in 1983 closely approximate these trends and are therefore still accurate for planning purposes.

The computing load for the VLBA has been assessed by similar means, but the results are less certain than for the VLA as they must be extrapolated from the current activity within *ad hoc* VLBI networks. The maximum data rate from the VLBA will be artificially limited to 0.5 Mbyte/sec (20 Gbyte/12 hr); the typical data rate should be closer to 1 Gbyte/day, similar to that of the VLA. Most of the subsequent data processing functions for the VLBA will be the same as for the VLA, though the emphasis will be different. The VLBA will be used with additional antennas, located world-wide, for some experiments; the number of baselines, the data volume to be handled, and the image size and complexity will all be significantly increased on these occasions. *The computing load will increase correspondingly.*

The continuum computing load of the VLBA should be only about 10% to 20% of that of the VLA. Several factors balance to produce this result:

- Some of the wide field imaging problems that produce heavy computing needs for VLA continuum projects are absent with the VLBA. "Confusing sources" are rarely a problem, owing to delay- and fringe-rate attenuation; in addition, the source regions that are bright enough to be detected by the VLBA are relatively small.
- Although images from the VLBA will generally be smaller (i.e., composed of fewer pixels) than those from the VLA, many more iterations of self-calibration will be needed for most VLBA observations, and bandwidth synthesis will be used more often.
- We do not expect the VLBA to be used much in snapshot mode, whereas snapshot observations make up a considerable fraction of VLA observing. However, use of non-

VLBA antennas in some experiments, and use of bandwidth synthesis, may make snapshot observing more attractive. The number of images produced for some projects might then increase by a factor of 10 to 20.

The spectral line computing demands imposed by the VLBA are much more severe than those of its continuum work. The fields of view (in resolution elements) required to observe maser sources can far exceed anything contemplated for the VLA, while the needed velocity resolutions and velocity ranges are similar. The worst VLBA cases are so demanding of computer resources that preliminary survey observations with lower-resolution telescopes such as the VLA will be mandatory, to restrict the VLBA observations to those small "cubicles" known to contain interesting emission. Even allowing for this, the aggregate spectral line computing load of the VLBA is expected to be similar to that of the VLA.

## 2.2 Computer Performance Requirements

To convert the astronomical requirements, specified by parameters such as image size, into computer hardware requirements, we have analyzed [19] the imaging requirements of the main classes of VLA and VLBA observing in terms of the number of floating-point operations and the input/output (I/O) bandwidth required. From this, we can convert the astronomical requirements into estimates of the required computing capacity.

Table 2 relates the *full-time* computing power that is required to handle *each* major class of project, at its projected fraction of the observing mix, to our present *total* capacity for VLA data processing. That total capacity is the equivalent of three Convex C-1's. It comprises the VLA's share of the computing done by two actual Convex C-1's and by four VAXes\* linked to Floating Point Systems Array Processors, plus the entire computing load of the VLA "Pipeline". These seven computing systems run a mixture of transportable AIPS software (Section 7.1) and highly optimized array processor microcode. The throughput achieved by these existing systems provides a good benchmark for expressing the astronomical requirements in terms of available computing performance. When expressed in these terms, *the sum of the entries in Table 2 directly measures our present shortfall in computing power, relative to the combined requirements of the VLA and the VLBA.*

Table 2 shows that processing the entire synthesis array computing load in real-time would require about 90 times the total computing capacity now in use at the NRAO. The projected VLA computing load is about 52 times present capacity, while that for the VLBA is about 37 times present capacity. The VLA thus contributes about 60%, and the VLBA about 40%, of the combined load.

The computing load projected for the VLBA in Table 2 is significantly greater than that estimated in the 1982 proposal to build the VLBA. In making the new estimate, we have allowed for improvements to the VLBA hardware specifications and in data processing algorithms since 1982. The increased computing load projected here is associated with significantly greater scientific capabilities than those specified in the 1982 proposal. In what follows, however, we work within the original 1982 VLBA budgetary estimates for post-processing computers. The proportion of the proposed budgets (Section 8) that we assign to the VLBA construction and operations also roughly matches the proportion of the overall computing load that we expect

---

\* VAX is a trademark of Digital Equipment Corporation.

to be generated by the VLBA. The increased computing power that can now be obtained from the original VLBA post-processing budget therefore roughly matches the increased computing demand expected from the VLBA, as well as the increased scientific expectations for the finished instrument. The 1982 VLBA proposal predicted that such increases would occur but it could not, of course, specify them as fully as we do here.

Table 2: Summary of Computing Requirements

| Type of Project                               | Estimated Capacity Needed<br>(relative to present capacity) | Sub-<br>Total |
|---|---|---------------|
| (a) VLA Continuum Observations:               |   |               |
| Full Field Imaging (A array, all frequencies) | 2.4   |               |
| Weak small objects, 327 and 75 MHz            | 4.3   |               |
| Snapshots (200 sources/day), all frequencies  | 7.2   |               |
| All other continuum projects, incl. mosaics   | 1.6   |               |
| Sub-total                                     |   | 15.5          |
| (b) VLA Spectral Line Observations:           |   |               |
| Extragalactic emission, all lines             | 6.4   |               |
| Galactic absorption, all lines                | 19.7  |               |
| Galactic OH and H <sub>2</sub> O masers       | 9.1   |               |
| All other VLA spectral line projects          | 1.1   |               |
| Sub-total                                     |   | 36.3          |
| (c) VLBA Continuum Observations:              |   |               |
| All projects                                  | 2.2   |               |
| Sub-total                                     |   | 2.2           |
| (d) VLBA Spectral Line Observations:          |   |               |
| H <sub>2</sub> O masers, HII regions          | 8.3   |               |
| H <sub>2</sub> O masers, proper motions       | 9.8   |               |
| Galactic OH masers                            | 7.1   |               |
| OH Mega-masers                                | 7.4   |               |
| All other VLBA spectral line projects         | 2.7   |               |
| Sub-total                                     |   | 35.3          |
| Total Computing Requirement                   |   | 89.3          |

The computing load for the VLA is dominated by projects that require capabilities that have long been present in the VLA hardware but which will first be made available to observers late in 1987. The VLA spectral line hardware has a factor of 16 more spectral channels than are now made available, and it also allows polarimetry in spectral line mode. This capability affects continuum observations as well as spectroscopy, as the multichannel spectral line system is needed to avoid "bandwidth smearing" in wide-field continuum imaging projects. *Small projects that now can coexist with spectral line data reduction at the VLA may be overwhelmed by it in the future.* The new 327 MHz and 75 MHz systems will also significantly increase the computing requirements of continuum observers.

The current imbalance between the available capacity and the computing needs of VLA observers has been managed by restricting the goals, and the frequency of scheduling, of the most computer-intensive projects. The imbalance is partly self-correcting, as many prudent scientists simply refrain from proposing hopelessly computer-intensive projects—but it will worsen once the capabilities released by the new on-line computer system come into demand. In the short term (1988) we will manage the imbalance by scheduling only those VLA proposals for which the present computing resources at the NRAO are adequate, or whose proposers have access to adequate resources elsewhere. In the longer term, the *total* computing capacity that is available for VLA (and VLBA) data processing—at the NRAO, at its user institutions, and elsewhere—must be increased to levels closer to those shown in Table 2. Failure to provide this increase will make the imbalance between the intrinsic capabilities of the VLA hardware and the available computing capacity increasingly clear to the user community. This situation will become steadily more difficult to accept, as the scientific capabilities of the hardware will be demonstrated by those few observers who are fortunate enough to have access to large computer resources.

## 2.3 Other Requirements

### 2.3.1 Interactive Data Displays

The VLA and the VLBA can produce three-dimensional image cubes with sizes of  $2048 \times 2048 \times 512$  pixels or larger. The generation of these images, and their refinement using image deconvolution and self-calibration techniques, involve recursive procedures. The astronomer will adopt processing strategies for each iteration of these procedures based on the results of the previous iteration, bearing in mind the specific astronomical objectives. Data which appear to be subtly contaminated may be temporarily excluded, and reintroduced at a later calibration step. The time required to reach a final, stable solution can be minimized if the astronomer can interact with the procedures as they execute, using displays of intermediate results to judge whether control parameters should be adjusted. The display of digital images throughout such analyses is indispensable. This interactivity demands high-bandwidth communications between the display and the main processing elements.

### 2.3.2 Archiving and Mass Storage

If the VLA were unconstrained by computer processing power, the projected mix of ob-

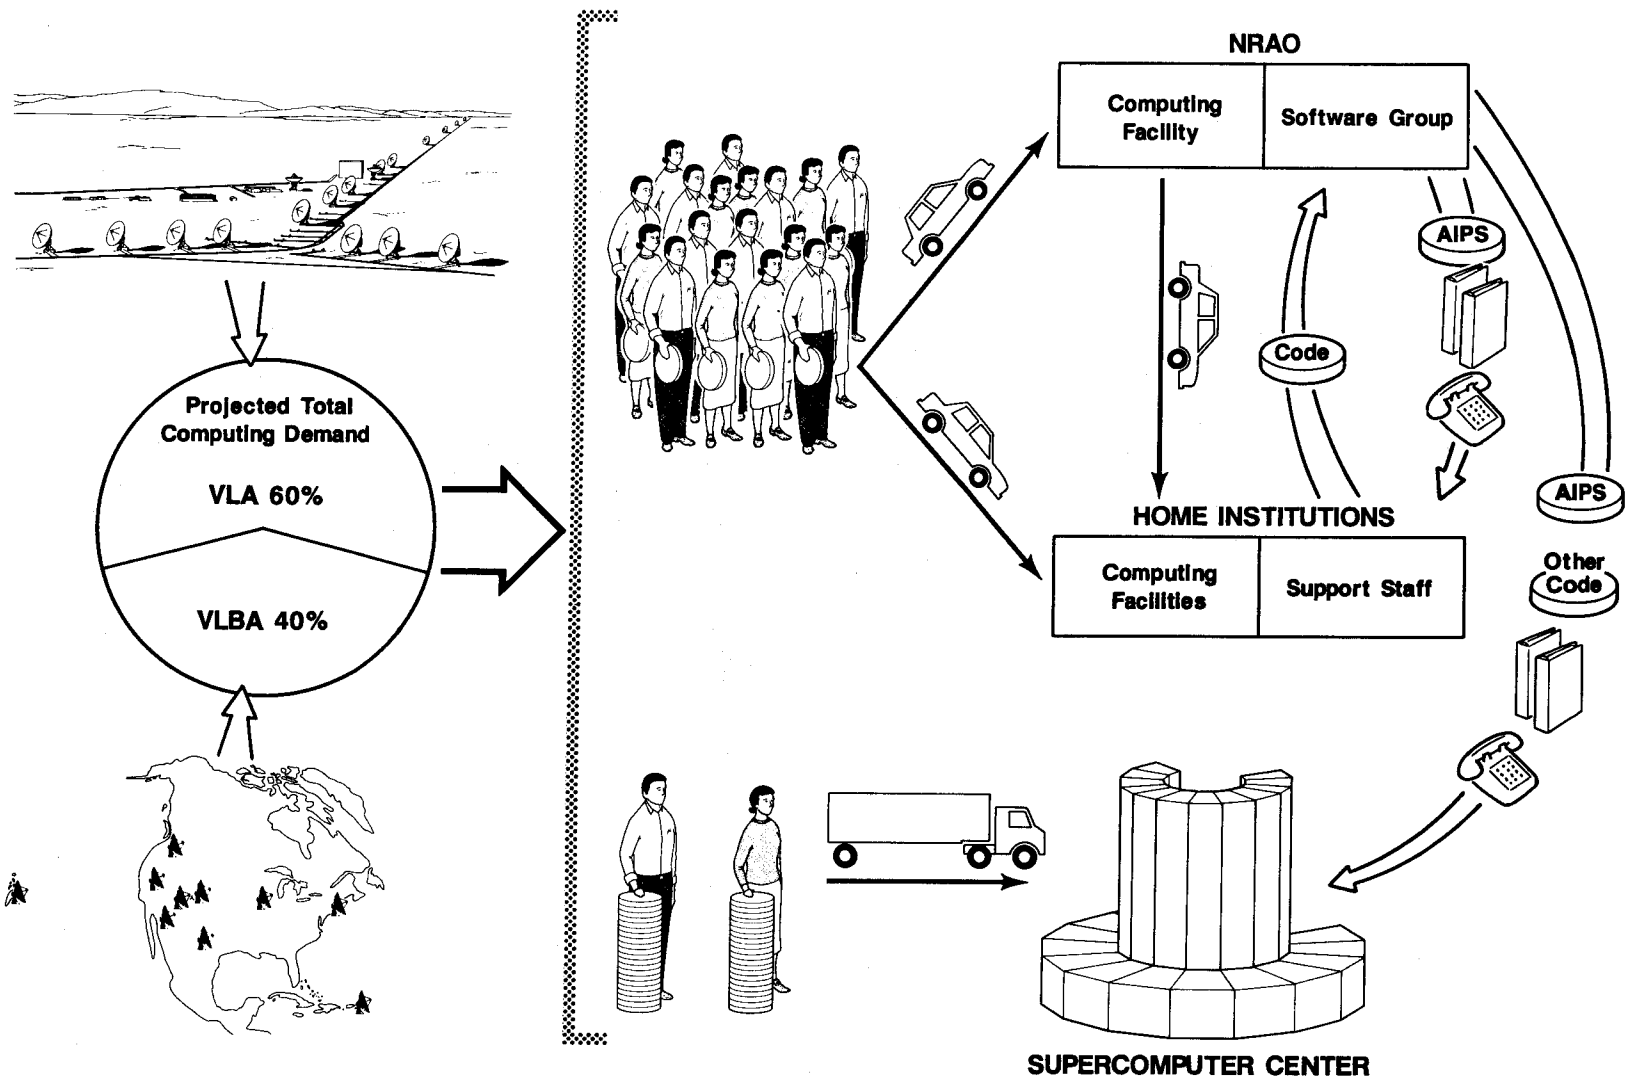
serving programs would produce an estimated output data rate of 800 Megabytes per day. In addition, several versions of the final images would be produced; both the original data and the final images must be archived. The archive storage requirement for the VLA will be from 1 to 2 Gbytes per day. The VLBA will at least double this estimate (and will greatly exceed it for short periods when operated in the peak burst mode). A required yearly net storage capacity of order 1500 Gigabytes results. Large blocks of this data stream must be downloaded to the computing system on demand. Processed images must also be stored for later analysis. This requires a mass storage/archive system capable of transferring large amounts of data without interfering with ongoing processing.

### 3. ESSENTIAL PARTS OF THIS PROPOSAL

Table 2 implies that a single machine to handle the combined VLA and VLBA computing requirement would need to be extremely powerful—and probably prohibitively expensive. We therefore propose a solution that will allow processing of about 90% of the anticipated VLA proposal load, and about 99% of the VLBA proposal load, either at the NRAO or at our users' home institutions. The rest of the load will be handled, with NRAO (and NSF) support, at general-purpose supercomputer centers. Our plan has three essential parts, illustrated in Figure 3-1.

- The 10% of the expected proposals that generate 70% of the computing workload (mainly, but not exclusively, galactic spectral line observations) will be treated separately from the rest. They will be handled in one of two ways, depending on the priority that is established for them within the radio astronomy community. The highest-priority projects in this class will be processed, with support from NRAO, in the most powerful machines available, e.g. in supercomputers at national and regional centers. The rest will simply be deferred (i.e. will not be scheduled on the telescopes) until either the cost of their computing becomes more manageable or their proposers can arrange their own access to adequate computing resources. By so treating 10% of the observing proposals, we reduce the total *dedicated* computing requirement for the VLA and the VLBA from 90 times present capacity to about 27 times present capacity.
- We will aggressively continue developing hardware and software solutions to the array image processing problems that can be readily exported to our user community. We are in an excellent position to do this because we have a long-standing commitment to making most of our image processing code (AIPS) readily portable. As a result of this, the VLA computing load is now distributed almost equally between NRAO and its users' home institutions [20]. Our goal will be to preserve this distribution, while increasing the efficiency of the software through optimization for widely-available hardware. This will require that additional NRAO manpower be dedicated to the support of computing outside NRAO, and that the NSF (and other agencies which support computing at user institutions) continue to fund hardware capable of image processing for the VLA and the VLBA at those institutions. Three valuable benefits accrue from this part of the plan. First, we support a significant fraction of our users' data reduction where they can do it most thoughtfully—at their home institutes, free from the artificial deadlines and the constraints implicit in a short visit to the NRAO. Second, by sharing a portable software environment with user institutions we can encourage our users to develop code that can be imported to the NRAO (and thus ultimately be made available throughout the user community as part of the AIPS package). Third, we reduce the total computing requirement at the NRAO itself by half—from about 27 times present capacity to about 14 times present capacity. This reduces the share of the problem that must be handled by hardware installed at the NRAO to a level that can be addressed by modular, multi-year funding of a loosely-coupled network of second-generation mini-supercomputers.





**Figure 3-1.** An illustration of the main features of this proposal. The off-line computing requirements of the VLA and the VLBA will be met by computers at the NRAO, at users' home institutions, and at supercomputer centers. The data processing needs of about 90% of the users will be met at the proposed NRAO facility and/or at their home institutions. A few users with extremely computer-intensive projects will take their work to supercomputer centers. The NRAO will directly support computing outside the observatory by exporting and documenting AIPS and other code, and by communication with staff at our user's institutions and at supercomputer centers. Code developed at user institutions will also be imported for inclusion in the NRAO-distributed packages.

- We will purchase, over a three-year period, four second-generation mini-supercomputers and associated workstations whose aggregate power will be about 14 times present capacity. We expect that, by 1991, advances in the machine performance and in optimization of our software will mean that these four machines will allow VLA and VLBA users to obtain images of the necessary quality in a reasonable time for all but the 10% most computer-intensive projects. This proposed increase in computing capacity at the NRAO will also approximately match the increase in the data rate that will appear when the new VLA on-line system becomes operational in late 1987.

The proposed NRAO facility will have three main purposes. First, it will be able to handle the data reduction for about 45% of the observing time of the VLA and the VLBA. It will therefore be able to handle most of the data processing requirements of the users who do not have access to significant processing capabilities outside the NRAO. At present, only those privileged users who have access to large computers for long periods can do the more difficult projects that this facility will handle. Second, it will provide the environment needed for continuing development of software and algorithms so that the NRAO can effectively support the remaining data processing at the supercomputer centers and at user institutions. Third, it will provide an environment wherein NRAO staff can continue to evaluate new hardware that may be applicable to the array image processing problem both at the observatory and elsewhere.

We believe that this proposal has flexibility in several important areas.

- By distributing the computing burden over a wide spectrum of machines, it permits matching the distribution of the hardware to the distribution of the computing load per project. The most computer-intensive work can be done in the most powerful machines, if necessary with highly optimized code that extracts maximum performance at some sacrifice in user-friendliness or convenience. The least computer-intensive work can be done in small machines, where convenience to the user and interactive response may be emphasized more than concentrated number-crunching power.
- By not requiring all the funding to be provided to one central site or at one time, we can use the NRAO's investment in software portability to avoid committing major funding to any one machine architecture while capabilities are rapidly evolving. An architecture that appears attractive now may have unforeseen disadvantages in the future as the computing burden and algorithms evolve. This proposal makes it easier to track the rapid growth of capability in the computing industry.

The following sections detail our plans for the four main elements of this proposal: (1) the proposed new computing facility, (2) NRAO support of computing at its users' institutions, (3) access to supercomputer centers, and (4) software and algorithm development and optimization.

## 4. THE PROPOSED NRAO FACILITY

### 4.1 What is the Appropriate Type of Computer?

#### 4.1.1 General requirements

The formation of radio images by Fourier synthesis is both CPU- and I/O-intensive. Good I/O performance is as important as fast CPU performance in our applications, where datasets can be as large as 4 Gbytes per user. Adequate disk storage directly attached to the computer system is therefore critical. Computers with a large amount of fast memory (either main or auxiliary) for temporary storage of files have special advantages for imaging—many scratch files are needed and files that fit into the memory can be transferred at rates an order of magnitude faster than would be possible if they resided on disk drives. Our experience using supercomputers at Los Alamos Scientific Laboratory and at Digital Productions has strongly reinforced this view of the importance of I/O and fast memory for image processing.

To maximize the amount and the quality of science done with the synthesis arrays, the image processing system(s) must also provide the astronomer with display tools that promote efficient comprehension and interpretation of the data. The data may be seriously degraded by different types of correctable instrumental errors, whose signatures on the data may differ only subtly. Because of this, and because of the complexity of the images, the process of data reduction often is—or should be—*exploratory*. Thus it is essential that our computing systems be highly interactive.

We have considered several architectures that might supply the needed computational power, I/O performance, and interactivity. We now discuss them in turn.

#### 4.1.2 Special-purpose computers

It is clearly possible to design and build special processors on a one of a kind basis for specialized processing, as is often done for signal processing. The VLA correlator is one example, and the NRAO is involved in another with the development of the VLBA correlator. We are therefore well aware of the advantages of this approach. It is viable only when a few—rigidly defined—algorithms (with long life expectancies) serve most of one's needs.

In 1976 (see Appendix A) the NRAO commissioned a design study for an unique special-purpose computer for VLA imaging. This processor would have used laser optics to perform the initial two-dimensional Fourier transforms, with only enough accuracy for the 100:1 dynamic range specification of the 1969 proposal. Although this design seemed attractive initially, it was abandoned, in favor of the VLA "Pipeline" (see below), late in 1977. Had it been built, the processor would have been obsolete by the early 1980's, as it would not have been able to compute images with the high dynamic ranges that could then be obtained via self-calibration.

As emphasized earlier, algorithms for Fourier synthesis imaging are still evolving rapidly, and this evolution is producing valuable increases in the scientific capabilities of Fourier synthesis telescopes. We see no prospect that our present algorithms are "final". Indeed, we propose to invest more effort in algorithm research (Section 7.4). Special-purpose hardware would

therefore be a particularly poor choice for our applications at this time.

There has been much research aimed at creating exceptionally powerful computers via massive parallelism. The number of highly parallel machines on the market has also increased significantly in the last few years. Although such architectures might bring enormous hardware cost advantages to applications for which they can be programmed, most remain confined to research environments or to special-purpose applications. They are not yet developed to the point where they can meet our overall computing needs while providing the needed flexibility. They also have a high ratio of software cost to hardware cost, and any image processing software that we might develop for them could not be exported to all our user community. In its 1982 report, the NRAO Computer Advisory Group recommended that our "... plan should not depend on being at the leading edge of computer ... technology", and in 1984 they reiterated this recommendation. Highly parallel machines have therefore been ruled out of our current plans.

#### 4.1.3 Array processors

Array processors attached to conventional CPU's are cost-effective for serial operations on data streams. Their limitation is that they are best suited to a restricted class of problems requiring number crunching. We have extensive experience with array processors; indeed, array processors hosted by mini-computers have provided most of the computer-intensive processing at NRAO in the last decade. Even with highly optimized code, the average throughput of array processors in our application was typically a small fraction of their theoretical capacity, because of I/O bottlenecks. Systems constructed from components supplied by many vendors were also difficult to diagnose and to maintain.

More significantly, our experience with array processors has demonstrated to us the need for flexibility and for ease of coding. While array processors appear capable of delivering needed computing capabilities at modest cost, their inflexibility could seriously inhibit development of new data processing techniques. Future dependence on attached array processors might therefore impede the development and use of image processing algorithms suitable for the VLA and the VLBA.

Both supercomputers (Section 4.1.4) and mini-supercomputers (Section 4.1.5) have important advantages over attached array processors for our applications. They provide unified environments of software, CPU, arithmetic pipelines, memory, and I/O devices—designed for the highest performance in the total system. They support high level languages with compilers that automatically use the pipelined vector hardware. This is significant, as it allows even experimental programs to realize a substantial part of the pipelining advantage.

#### 4.1.4 Supercomputers

For our purposes, "supercomputers" can be defined as the fastest general-purpose computers available. Supercomputers are, in principle, attractive for our application because they can perform the largest processing calculations in a reasonable time. This led us in 1985 to propose that a supercomputer would be a good solution for our image processing needs. Since then, we have accumulated more experience with supercomputer systems.

The most intensive interaction was with a Cray X-MP system run by Digital Productions, in Los Angeles. We obtained 400 service units\* of computing time between October 1985 and June 1986 and installed a substantial part of AIPS and a stand-alone Maximum Entropy deconvolution program. There were a few bugs in the software, and numerous inconveniences, but a workable user system for large image processing problems was close at hand when funding for this facility was discontinued.

This experience made it clear, however, that the normal supercomputer environment is not in harmony with all the data reduction requirements of synthesis array telescopes. Many needed functions that are easy to implement in VMS\*\* or UNIX systems are difficult to implement on Cray systems. These include: file management on tape, communication with running tasks, and interactive display of the data and of intermediate steps in the processing. This experience and the emergence in 1985 of the more flexible and affordable "mini-supercomputers" (Section 4.1.5) have convinced us to restrict our use of supercomputers to the reductions for which their CPU performance is critical. Section 6 describes our plans for future use of supercomputers.

#### 4.1.5 Mini-supercomputers

Mini-supercomputers represent a technology that did not exist at the time of our 1985 proposal and which was therefore not evaluated there. These machines currently range in power from 8% to 25% that of a single processor Cray X-MP, as measured by standard benchmark programs. In the mini-supercomputer industry, the performance for a given price is now approximately doubling every 18 months. The number of companies competing in this market leads us to expect that this growth will continue, despite any "shakeout" that may occur. Several manufacturers have assured us that by 1989 mini-supercomputers will be available with performance comparable to (i.e., about 75% of) a single processor Cray X-MP.

NRAO currently operates two first generation mini-supercomputers, both Convex C-1s, one in Charlottesville and one at the VLA. These machines have proven to be well suited to our image processing needs. Benchmarking our code on competitive equipment leads us to believe that mini-supercomputers from other manufacturers also would be satisfactory. We thus believe that we will be able to take advantage of industry competition.

Mini-supercomputers offer several advantages over traditional supercomputers. They provide more flexible environments for interactive image processing. They run the UNIX operating system, which offers portability. They typically support more memory than the traditional supercomputers. They cost much less (both to acquire and to operate) than supercomputers and are available from a variety of manufacturers, making competitive procurement possible. Common peripheral devices are used. Mini-supercomputers can achieve the crucial high disk I/O bandwidth by disk "striping", i.e. by spreading logical files across several disk spindles. The lower cost of their disks makes the large disk capacities required for our applications more

---

\* A *service unit* is a unit of computer resources defined by a complicated function of the operating system's monitored statistics, including CPU usage, memory residency, and I/O activity. One service unit is approximately equal to one hour of CPU time in a single processor Cray X-MP.

\*\* VMS is a trademark of Digital Equipment Corporation.

affordable. It is also easier to interface special devices, such as image displays, and the site requirements for power and cooling are much less.

The 32-bit architecture of mini-supercomputers is adequate for our dynamic range needs. A dynamic range of 100,000:1, the present VLA limit, requires no more than 17 bits and a few guard bits. With an additional 10 bits for sign and exponent, 32 bits suffice for most applications. In this respect, the exclusive use of 64-bit words in supercomputers is a disadvantage of their architecture, as it forces our data to occupy twice the memory that is needed for the required precision.

Mini-supercomputer systems have drawbacks, however. The anticipated performance of individual systems is not adequate for our most extreme needs, making the use of some number of them necessary. It might then be difficult, if not impossible, to bring the total computing power to bear on a single problem. Also, the power in an individual system is usually achieved not by clock speed or vectorization alone, but by additional strategies—often involving some form of concurrency. These strategies are generally invisible to the user, but some limitations are imposed on the algorithms to extract full performance from the machine. Finally, it is unlikely that all the (more than a dozen) companies which intend to market in this area will survive. This imposes some risk on any potential customer. These disadvantages can be overcome with prudent management.

## 4.2 The Proposed Facility

We consider the compelling argument for the use of multiple mini-supercomputers at the NRAO to be that they can provide the computing power needed for about 90% of the array telescope observing proposals in an affordable, flexible, general-purpose environment. This environment will foster software development, which could be as important as the initial computing capacity (Section 7.2). The modularity of a plan based on several mini-supercomputers allows the procurements to be spread over several years; this permits assembly, and later augmentation, of the system in an evolutionary manner that can track developments in mini-supercomputer technology.

Figure 4-1 gives the proposed configuration for a loosely coupled network of four mini-supercomputers. Three will be located in Socorro, and one in Charlottesville.

### 4.2.1 The main processors

The main processing element in the proposed system will be a mini-supercomputer having a capacity about 75% that of a single-processor Cray X-MP. This element itself may be made up of several concurrent processors; but the concurrency must in general be invisible to the user and to the programmer. Each processing element should have approximately one Gigabyte of fully randomly addressable memory. This memory will be adequate for large image arrays and for temporary file storage (sustained transfer rates from memory to memory will generally be at least ten times faster than from disks to memory).

Twenty Gigabytes of disk will be provided on each system. The large disk storage is dictated by the requirement that the total disk capacity be many times that needed by any one user process (more than 1 Gigabyte per user for large problems) to permit the staging of new

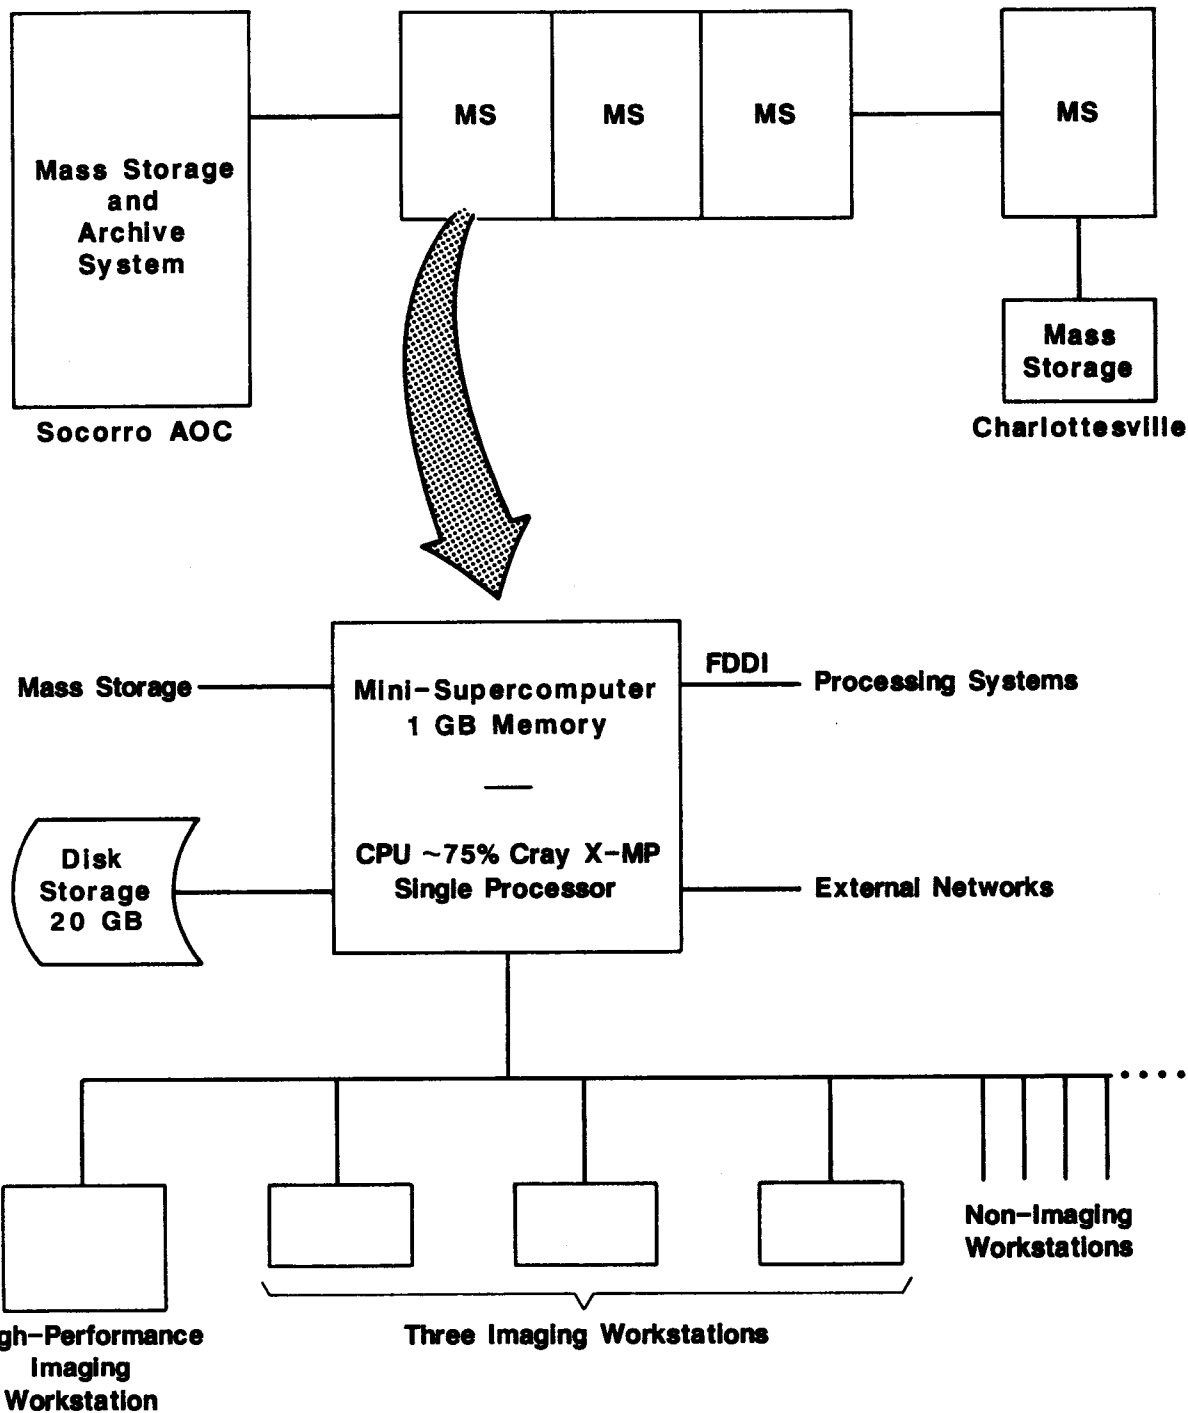


Figure 4-1. The configuration we propose for a loosely coupled network of four mini-supercomputers (MS). The lower panel shows detail of the individual elements depicted in the upper panel.

data from tape to disk, the processing of multiple job streams, and the staging of results from disk to tape units and display devices.

Many synthesis imaging tasks perform relatively few arithmetic operations on more data than can be kept in memory for several simultaneous users. Data and intermediate products must therefore be moved from the CPUs to disk and back, often repetitively. The disk-to-CPU bandwidth must be high enough to exercise the arithmetic pipelines continuously if full performance is to be achieved. This requires about 50 Mbyte/s per processing element. This may, in practice, require disk striping (e.g. using parallel disks running at about 10 Mbyte/s each), but the fundamental goal is adequate bandwidth, however it is obtained.

FORTRAN and C are the required languages. UNIX or VMS is the preferred operating system. The processors will be loosely coupled, to help system maintenance and to enable file transfer between processors.

#### 4.2.2 Interactive image processing and workstations

Foremost in importance is the ability to display the final images in grey-scale form. Image display devices with storage for  $8192 \times 8192$  pixels using  $2048 \times 2048$  displays are available. The proposed image display facilities will allow for the flexible display of  $2048 \times 2048$  images, giving astronomers an instant overview of their data, as well as enabling them to discern and enhance interesting detail.

The facility will have a range of user workstations, with capabilities roughly corresponding to the spectrum of problems. One high performance imaging workstation and three lower performance imaging workstations will be the main user interfaces for each processing element. A larger number of low-cost workstations, probably without image displays, will also be provided.

The high performance workstations are expected to be so-called "graphics supercomputers". Two vendors will announce machines in this class in the Fall of 1987. These machines will have computing and I/O performance nearly equal to that of the first-generation mini-supercomputers, tightly integrated with a workstation environment (high-resolution monitor, windows, mouse). The display is many bit-planes deep, so that these machines can be used as intelligent digital image displays. Large random access memories and fast disk systems will allow rapid exploration of the large digital images, especially spectral line cubes, which will be computed in the mini-supercomputer complex. Hardware support for graphics rendering, plus stereo imaging of 3-D data, under interactive user control, will provide an ability to visualize images that will be well matched to the computational ability of the complex.

Initially, conventional file transfer protocols and standard Ethernet will connect the display workstations to the processing elements. As the workstation and communications technology matures, we expect that data on disk drives in the complex will be shared using the NFS (Network File System) protocol, and that connections between the processing elements and the workstations will probably use the FDDI (Fiber Distributed Data Interface) standard. The use of workstations, which connect to the processing elements using industry standard protocols, will provide freedom in the location of the image display and analysis facilities, and will enable the complex to exploit evolving technology with minimum disruption.



### 4.2.3 Hard copy devices

Hard copies of the images are needed, with fast turnaround for intermediate results and with high quality for final, or nearly-final, results. Three classes of hard copy are necessary: (1) black and white or color copies of the displayed images, obtained from the monitor video signal, (2) laser printer copies of graphics and grey-scale plots, and (3) high quality final images to be written on an image recorder.

### 4.2.4 Communications

Ideally, a communications system would allow us to make the full capabilities of the facility, including interactive image processing, available to all NRAO users, wherever they are located. The performance goal for remote service must however be a realistic compromise, limited by current economics and by current technology of communications links. A peak burst rate of 400 kb/sec is desirable for single image planes; such links will be able to write 1024 x 1024 8-bit images into remote digital image displays in 20 seconds. We expect that a rate of 100 kb/sec will still be useful in the near term, but 10 kb/sec definitely will not. Fast transfer of spectral line data cubes will be difficult until T-1 (1.5 Mb/sec) capability is widely available.

VLA and VLBA users should have interactive access to computers at NRAO, as well as at NSF supercomputer centers, from their home institutes. We will therefore develop our code for image display workstations so that they can be operated over the NSF Supercomputer Network and thus benefit remote users as well as those at the facility itself. (The high-performance workstations are likely to be an attractive option as the entire installation at many of our user institutes.) Although the initial performance of the network will probably be marginal for remote image processing, anticipated improvements should make remote operation feasible for a growing subset of the applications. To stimulate development of this mode of operation, we ask that the NSF supercomputer office install a 56 kb/sec tail circuit at the Socorro facility.

## 4.3 Mass Storage and Archive System

The mass storage/archive is a difficult area to address because there does not yet appear to be a good hardware solution. What we desire resembles the common file systems employed at some supercomputing centers. However, most of these systems in their present form are not optimized for mean file transfer sizes as large as we require. Most are set up to back up files automatically, in an environment where resource sharing between many users is a primary consideration. In contrast, we typically support a few concurrent large users who require a convenient way to manually load and unload data blocks of order 1 to 4 Gbytes in size. Of the systems now in use, that at the National Center for Atmospheric Research seems closest to meeting our needs.

For budgeting purposes we have considered a modification of the common file system used at the supercomputing centers. We hope, however, that an alternative and more suitable system will become available in the near future. There are options, for example, based around NFS (Network File System), which use IBM cartridge tape drives. We propose not to acquire the mass storage/archive facility until the second year of the plan (1990), when the alternative

possibilities will be clearer.

#### 4.4 Operations Personnel

We will provide operators to manage disk resources, to stage data into and out of the computers, and to help users to overcome logistical difficulties. Three new operations personnel will be required. The existing number of maintenance staff at the VLA (including those budgeted for the VLBA) will be adequate for the proposed facilities, if older equipment is phased out as planned.

#### 4.5 Existing Facilities

As the speed and sensitivity of the VLA grew, we took steps to increase computer resources for our user community using the existing operating budgets. Appendix A chronicles the development of array telescope computing at the NRAO. Our present computer capacity comprises a complex and heterogeneous mix of computers. We still depend on some systems acquired over 10 years ago. These are increasingly costly to maintain, and, within the limited manpower and capital budgets available, we plan to phase them out. The original on-line computers will be replaced by a new system at the end of 1987. The DEC-10 system, used for calibration of VLA data, for all monitor data processing and for various general purposes, is being phased out. New calibration and editing software has been developed in AIPS to replace the DEC-10 calibration functions and to provide the additional calibration procedures needed for the VLBA. The new on-line system will also have enough capacity to do more of the basic calibration and validation checks as the data are collected. Software to handle the VLA monitor data outside the DEC-10 is being developed. Other general purpose programs, such as those for observing file preparation, for scheduling and for archiving are being rewritten. Our goal is to remove the DEC-10 in 1989. The schedule is determined primarily by the software effort needed, and has been assumed in estimating the budget shown in Section 8.

The "Pipeline" system of networked PDP-11 computers, array processors, and shared disks is used for batch processing of large multi-channel projects. Since it is intimately coupled to the DEC-10 it must be removed at the same time. If new hardware is not available in 1989 this extra load will cripple the other existing systems.

As the new hardware described in this plan is acquired the workload on the VAXes will be redistributed, and eventually they will be removed. The Convex C-1's will likely remain in operation for general-purpose computing and for some image processing until the mid-1990's.

## 5. COMPUTING AT USER INSTITUTES

Much effort has gone into making the NRAO Astronomical Image Processing System (AIPS) truly portable. Thanks to this prior effort, we can help our user community to run most of our code in many models of computer, large and small, with only minimal modifications. We have done this by using generic FORTRAN wherever possible and by isolating all system-dependent code into well-defined groups of routines. AIPS, when implemented in different computer architectures and under different operating systems, tries to present as nearly the same interface to the user as possible.

We have sought this level of hardware and operating system independence in AIPS for two main reasons: first, to ensure that our software can migrate quickly, without major recoding, whenever manufacturers' advances in hardware and in compiler technology make this attractive; and second, to serve the needs of VLA users at their home institutes, where available hardware and operating systems may differ substantially from our own. It is important for this plan that we continue to export software to our users' home institutes, whether or not they follow our hardware choices, for two reasons:

- We can support our users' science better by supporting their computing where they (and their students) can do it most conveniently and most thoughtfully—at their home institutions. There, they have more time to plan and to test reduction strategies than they do while visiting the NRAO.
- Successful software export from NRAO reduces the need for concentrated computing facilities at the NRAO itself.

The advantages of software portability are now widely recognized in the astronomical community, and other imaging processing packages, notably the National Optical Astronomy Observatory's IRAF, also seek to provide it. The AIPS package is now actively used on over 100 different computers worldwide [20]. As a result, about equal amounts of computing power are now devoted to VLA data processing inside and outside NRAO. AIPS has been run successfully on Cray and Fujitsu supercomputers and on Convex and Alliant mini-supercomputers, as well as on VAX, MicroVAX, Apollo, Charles River Data Systems, Data General, Jupiter, MassComp, Nord, Prime, Ridge, and Sun machines. This broad base of non-NRAO machines running AIPS is growing rapidly. The total machine power devoted to VLA data reduction at non-NRAO sites in the U.S. grew by 62% from 1985 to 1986 [20], while that in other countries grew by 91%. (There are now about equal total non-NRAO machine powers devoted to VLA data reduction in the U.S. and in all other countries.) In response to a recent survey of the active AIPS sites, many site representatives foresaw significant expansions of their AIPS processing power in the future. Forty-eight expect their sites to acquire workstations for AIPS processing, either as stand-alone systems or for use with more powerful computers. Eleven have plans to acquire "mini-supercomputers". Fourteen state that they expect to have access to supercomputers for AIPS processing.

We plan to foster continued growth of AIPS for VLA (and VLBA) data reduction outside NRAO, so that the data from about 40% to 50% of all VLA and VLBA observing time is

processed at user institutions. We will do this by maintaining the ability of AIPS to run on many types of computer, from workstations to supercomputers, by exporting it as widely as possible, and by importing code of general interest that is written at AIPS user sites. All implementations of AIPS will continue to provide the same functionality—differing only in performance, according to the computer architecture. The following NRAO activities are therefore an integral part of this plan:

- Distribution of AIPS installation tapes and documentation to all astronomical institutes that request them;
- Assistance with installation of AIPS at user sites when problems arise for which the NRAO staff have relevant experience;
- Consultation with active AIPS sites about operational problems associated with AIPS at their institutions;
- Co-operation with computer vendors to test and benchmark AIPS in promising classes of machine as they come to market. This enables us to advise user institutions on their procurements and sometimes to pre-certify that a machine is suitable for running AIPS, even if NRAO will not acquire one for its own use;
- Co-operation with computer vendors to arrange the loan of equipment for which machine-specific AIPS interfaces need to be developed;
- Promotion of a worldwide AIPS Users' Group through newsletters, electronic mail networks and workshops, so that users at institutions whose hardware differs from ours can use the NRAO as a communications channel for problem-solving and for code exchange;
- Distribution of non-NRAO coded AIPS tasks and machine-interface routines along with NRAO's own AIPS installation tapes.

These activities are already underway, to the extent that manpower limitations permit. But the recent 60% annual growth in the number of active AIPS sites has meant that NRAO staff can barely keep up with the distribution and maintenance of AIPS. Code reliability and development are suffering. It is essential that additional manpower be provided, so that we can continue to provide an acceptable level of service to the user sites, while still developing the AIPS package. Such manpower is budgeted in this proposal (see Section 7.3).

Finally, we emphasize that the benefits of exporting AIPS from NRAO are not confined to VLA (and eventually VLBA) data reduction. As much non-NRAO machine power is devoted to running AIPS for other radio astronomy data processing as is in use for VLA data reduction. Furthermore, although the package was designed primarily for processing data from Fourier synthesis radio telescopes, it has enough generality to be attractive for single-dish radio astronomy and for other, non-radio, applications. Use of AIPS for non-radio applications (including processing of IRAS data, optical CCD data, X-ray and UV imaging data, visualization of numerical fluid mechanical simulations and some medical imaging) amounts to about 10% of the usage for VLA data processing [20]. Strengthening our commitment to support the package outside NRAO thus has "spinoff benefit" beyond the VLA and the VLBA user communities.

## 6. SUPERCOMPUTER ACCESS

### 6.1 Purpose

Access to supercomputers will permit the data analysis required by some exceptionally computer-intensive projects that deserve high priority but whose requirements would otherwise enormously overburden the proposed facilities at the NRAO and its user institutions. We estimate that 10% of future VLA projects, and a few per cent of VLBA projects, will fall into this category. Supercomputer access will also provide an "escape valve" if new algorithms, or greater scientific interest in one of the most computer-intensive research areas, begin to saturate other systems at the NRAO and its user sites. It will also be important to have access to supercomputers for algorithm research, to ensure that algorithm development is not afflicted by "tunnel vision" about the available resources.

### 6.2 Supercomputer Experience and Overview

We have invested significant effort in using supercomputers for VLA reductions and VLBA simulations. NRAO staff have worked directly with a Cray 1-S at Los Alamos National Laboratory, and with Cray X-MP's at Digital Productions in Los Angeles and at the Pittsburgh Supercomputer Center. AIPS and stand-alone NRAO image processing software (notably deconvolution by the Maximum Entropy method) have been run on these and other supercomputer systems, including a Fujitsu VP50 supercomputer in Japan and Cray X-MP's at the Naval Research Laboratory and at the University of Toronto, Canada. Our experience at Digital Productions showed us that AIPS can be made to run in supercomputer centers, although the needs of interactive image processing are difficult to support. Nevertheless, our experience there has encouraged us to take the time needed to develop a robust, updatable implementation of AIPS that could be exported to other Cray sites that operate under COS. This work is currently manpower-limited, but we expect soon to install this improved AIPS implementation in the Cray X-MP at Pittsburgh, and elsewhere.

We therefore regard continued involvement with supercomputers, particularly those operated under the auspices of the NSF Supercomputer Access Program, as a key ingredient in this plan. We intend to supply supercomputer centers with image processing software that will allow their users to process VLA and VLBA data. We also seek close ties with individuals and groups who are studying image processing and display techniques for supercomputers; co-operation with such groups should accelerate development of NRAO-supplied software for supercomputers. (There are active efforts in these areas at the University of Illinois, the University of Minnesota, the University of Toronto, and elsewhere.) Our immediate goals *vis-à-vis* supercomputers are twofold:

- We will maintain AIPS in one supercomputer operating system (currently COS). We will also support AIPS within UNIX emulations running in supercomputers that have other operating systems. (There is little uniformity of supercomputer operating systems, so it is

had the same operating system version as at Digital Productions, the Crays in Pittsburgh, Illinois and Minneapolis have different operating systems, and installation of the same software will be far more time-consuming. The anticipated use of UNIX as a Cray operating system should lessen this problem.)

- We propose to use the NSF Supercomputer Access Program between 1989 and 1991 to develop and run NRAO reduction programs on at least one nationally-available supercomputer.

We believe that these activities, if undertaken with adequate manpower, could significantly improve our ability to handle the most computer-intensive part of the VLA and VLBA data processing load by 1991. However, experience shows that certain features, not routinely found at large supercomputer centers, are essential for our applications to use supercomputers effectively. We describe these explicitly in the next section.

### 6.3 Future Use of Supercomputers

We propose to use supercomputers outside NRAO for the most computer-intensive 10% of the data reduction problems. To support this effectively, NRAO must have enough funds and administrative leverage to obtain features in the host systems that will maximize throughput. For this approach to be successful, it will be important to minimize the wall clock time for a given reduction, as well as the CPU time. The essentials of our plan are as follows:

- The center administration must be ready to circumvent operating restrictions that inhibit the data flow in our applications, such as mandatory archiving to tape of large data sets and codes. They must be willing to schedule critical resources, including disk storage, to maximize data transfer rates for our applications.
- Experience has shown that constant attention to data reduction details and intermediate results is necessary, to avoid costly mistakes. Interactive display facilities are therefore essential.
- A three year commitment should be made to the chosen supercomputing center, so that the initial effort of developing an appropriate hardware and software environment there has time to bear fruit. Based on the NRAO experience with the Cray X-MP at Digital Productions in 1985-1986 in which 400 service units were used, we estimate that 1000 service units (su) of Cray time in 1989 will be necessary to install AIPS and the stand-alone software, to integrate the new hardware, and then to begin synthesis data reduction. We will request this allotment from the NSF's Supercomputer Access Program.
- Proposals that require processing time at the NSF centers should be refereed by the NRAO as part of the observing proposal evaluation.
- As most of the work to be done at these centers will involve large data cubes, for which remote display and analysis are unlikely to be practical by 1989, most groups will have to travel to the center to do their data reduction. When in full operation, the array telescope processing facility should be able to handle two groups at once to optimize throughput. Communication between co-resident user groups will ultimately be beneficial in minimizing mistakes and promoting transfer of expertise.
- One full-time programmer, knowledgeable in image processing and supercomputers, should reside at the center to support the image processing systems and to advise visiting users.

- The viability of this part of our proposal is contingent on the computing capacity of the supercomputer centers growing at a rate that at least tracks the growth of capability in the mini-supercomputer industry. We assume that policies are in effect to ensure this.

The software to be run in such a supercomputer center will be a combination of AIPS and several crucial stand-alone programs, as developed either at the NRAO or by involved scientists at the supercomputer centers. The number-crunching will be concentrated in a few important programs (wide-field imaging and deconvolution, mosaicing, self-calibration) for which highly optimized code, perhaps initially in stand-alone packages, will be appropriate. Our experience at Digital Productions showed, however, that many of the less computer-intensive AIPS functions are also needed to guide and monitor the progress of the image processing. The ability to edit and re-calibrate the data in the midst of processing is occasionally crucial. Code transport to other supercomputers and to other architectures will be helped if stand-alone programs are ultimately incorporated into AIPS.

## 7. SOFTWARE AND ALGORITHM DEVELOPMENT

The software tools needed for Fourier synthesis imaging fall into three categories: (1) data validation, editing and calibration; (2) image construction and refinement; and (3) image display and analysis. Figure 7-1 shows a simplified outline of the main stages of Fourier synthesis image processing.

The initial stages of the analysis, including data validation, editing and calibration, have hitherto been done mainly at the VLA site in the DEC-10 computer. The remaining functions involving images and their analysis have mainly been carried out either in the NRAO-supplied Astronomical Image Processing System (AIPS) at the NRAO or at the user's home institution, or in the VLA "Pipeline" computer.

### 7.1 The Astronomical Image Processing System (AIPS)

AIPS has been the principal tool for display and analysis of both two- and three-dimensional radio images (i.e., continuum "maps" and spectral line "cubes") from the VLA since early in 1981. It was initially intended that AIPS be confined to this role. In 1981, as it became evident that the DEC-10 computer and the Pipeline alone would be unable to handle the full VLA data stream, the scope of the AIPS effort was expanded to include manipulation and calibration of the cross-correlated data, and image formation (see the chronology in Appendix A). Since then, AIPS has provided the main route for self-calibration and imaging of VLA continuum data. It contains facilities for display and editing of data in the aperture plane; for image construction by Fourier transformation; for deconvolution of the point source response by "CLEAN" and by maximum entropy methods; for image combination, filtering, and parameter estimation; and for a wide variety of TV and graphical displays. A batch mode of operation is also available.

In 1983, the scope of the AIPS effort was expanded again in anticipation of the needs of the VLBA and of phasing out the older computers at the VLA site. By the end of 1987, AIPS will embrace all stages of radio interferometric calibration, both continuum and spectral line, as well as the geometric and delay calibrations required for VLBI. AIPS will then contain most of the calibration and editing functions previously performed only in the DEC-10, in addition to the post-calibration capabilities outlined above. At that point, AIPS will handle all data reduction steps that normally occur after correlation of radio interferometric data.

The AIPS software is not static. It continually evolves in response to new developments and to the needs of its users. AIPS must accommodate (1) new astronomical applications, (2) the benefits of optimizing code for new hardware, and (3) improved visualization of, and extraction of astronomical parameters from, images and cubes. Support of a large software package (now 600,000 lines of code) for a growing user community also requires increased emphasis on remote user support and on documentation.

Significant improvements to AIPS in all these areas are an indispensable part of this proposal. We now review briefly the main areas where development and additional manpower are needed.



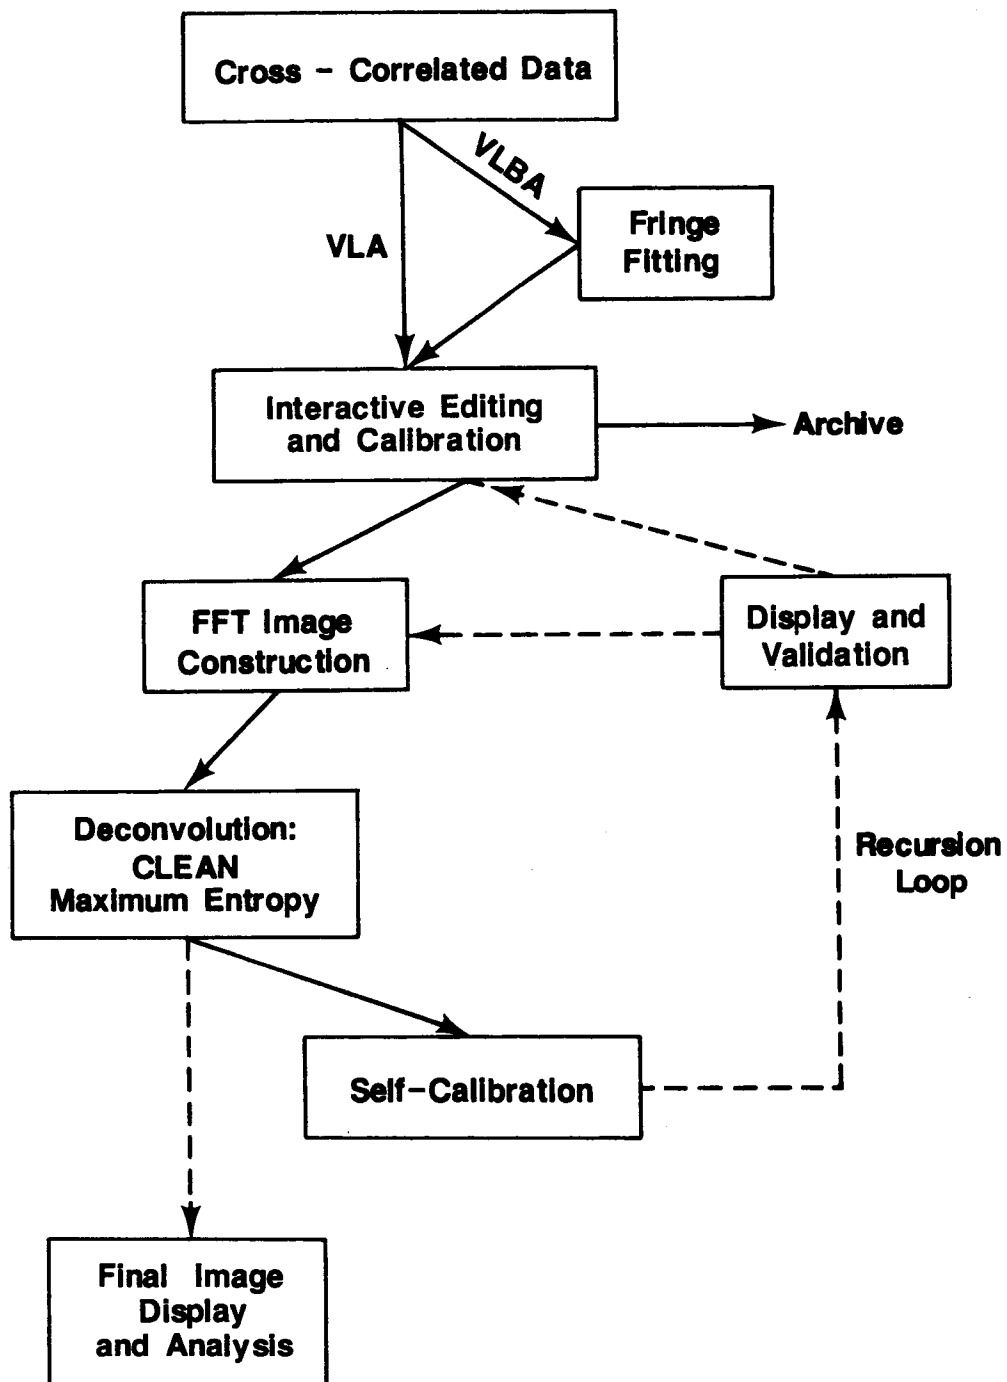


Figure 7-1. The main stages of Fourier synthesis image processing and analysis.

### 7.1.1 New applications software in AIPS

Wide-field imaging problems require corrections for the varying instantaneous geometry of the array as seen from different positions in the source. Although the theory is well-developed, the requisite algorithms have yet to be realized efficiently in AIPS. These problems will be particularly acute at the lower VLA frequencies, where sensitivity to low surface-brightness, steep-spectrum emission is the greatest. Many studies require the ability to mosaic, or "tessellate", images of adjacent fields into a single wide-field image. The techniques for handling this task in AIPS need considerable refinement.

The technique of self-calibration relies on constructing ever-improving models of the intensity distribution over the sky, and feeding these back to the calibration process (Figure 7-1). The present techniques for doing this in AIPS are based on CLEAN models, i.e. on linear combinations of finitely many point sources; they are cumbersome for wide fields of view and for complex images. More sophisticated approaches exist and AIPS should incorporate them.

The software demands of spectral line observations involve the need not only for more efficient algorithms to handle large data cubes (e.g., for repeated image deconvolution and self-calibration) but also for software unique to spectroscopy. A good example of the latter is an algorithm for self-calibration in the spectral domain using the spectral line itself as the calibration "source." Pioneering work in software suited to array spectroscopy is being done elsewhere. We hope to take advantage of this expertise to incorporate additional capabilities into AIPS, as well as to continue our own development work.

The use of precision geometric models will be critical to the success of the VLBA, whose long baselines and full time operation will make it a powerful instrument for radio astrometry and geodesy. It will be important to import software that is now in use within the astrometric and geodetic communities into AIPS, so that these procedures become part of the calibration and imaging toolkit of the VLBA, and can be maintained by NRAO staff.

### 7.1.2 Improved visualization

When handling large databases, simply looking at the data can become a major difficulty. In Figure 7-1, this function is described as data "validation," meaning the process of identifying discrepant data points and of deleting them from the database. While it is common practice to apply crude numerical filters to the data, the assumption behind these techniques is that one knows the nature of the errors in advance. Often this is not the case. More subtle effects are best identified in visual displays using, for example, color-coded images of the data sorted in various ways. Development of effective displays for aperture-plane data, and of algorithms to support interactive data editing, is much needed.

A second important aspect of visualization relates to the last stage of analysis, where the astronomer has produced images and seeks to understand them quantitatively and to extract their astrophysical content. The spectral line astronomer will need to page rapidly through the data base; the continuum astronomer may want to compare images made at different frequencies or in different polarizations, to re-grid the images or change the projection. In most applications, there is a need to fit models, to compute image statistics, and otherwise to parametrize the astronomical content of the final results. The expansion of AIPS to include

image formation and calibration has left this area underdeveloped, and we plan greater attention to it in future. Intelligent routines for finding and classifying emission regions in wide field images are also needed.

These functions—manipulating images to extract their astronomical content—are vital to successful use of the final images by the astronomer. They are best done in an uncongested, single-user workstation environment which can promote interactive analysis and scientific insight. Development of appropriate image display and analysis software for workstations in AIPS has high priority, and needs additional manpower.

## 7.2 Optimization

Vector architectures and computers with large memories offer such great potential for synthesis array computing that we must ensure that our code uses them effectively. Optimizing AIPS software for widely available hardware will be an important consideration. Optimization must also be a high priority for all code to be run in the supercomputer centers, whether within AIPS or as stand-alone programs.

Particular attention will be given to effective use of large memory resources, where these are available. The volume of data from a full synthesis observation at the VLA can range from 100 Mbytes to 1 Gigabyte. The computers now at the NRAO handle these data with a memory of 4 to 64 Mbytes by constantly paging the data in and out of memory, so that I/O contributes appreciably to the run time of a given task. The present AIPS software is optimized with this limitation in mind. However, the next-generation hardware will be able to approach the problem far more efficiently and all our imaging software must evolve to recognise this. For example, both run time and disk storage requirements could be reduced by eliminating precise sorting of the data before imaging, and the re-sorting that now occurs during each self-calibration cycle.

Greater efficiency in the processing of multichannel spectral line data is also needed, to provide the rapid cube-building capabilities that currently exist in the highly optimized but hardware-specific VLA "Pipeline" software.

It is difficult to predict the overall gains in performance that can be achieved by increased attention to optimization. They will, of course, vary with the hardware environment and with the degree to which present code has already been optimized. In the past, greatest attention has been given to the code segments that most affect throughput, but, even so, past experience suggests that improvements by factors of two may be achievable in some significant areas. The tradeoff between optimization and portability in all NRAO code will be closely examined.

## 7.3 Manpower Needs

Evolution of the AIPS software to address the priorities outlined in Sections 5, 7.1 and 7.2 requires substantial additions to the AIPS programming staff. This staff has not grown since 1982 despite the greatly increased scope and complexity of AIPS and an approximately 60% per year annual growth rate of the AIPS user community. Many tasks outlined below are now being undertaken on an *ad hoc* basis by AIPS programmers and other NRAO staff. Defining new positions for personnel with special responsibilities in the areas outlined below will ensure that they can be given the priority necessary for AIPS to play the role envisaged for it in this

plan. The widespread use of AIPS outside NRAO, and our belief that this must continue as an integral part of this plan, mean that investments in AIPS development will bring highly leveraged benefits to the astronomical community.

The complex computing environment and large user load that is envisaged at the proposed NRAO facilities by 1991 will also require additional personnel to support production image processing.

Additional positions are needed in the following areas:

- Graphics interfaces and workstation support (2 positions). Improved AIPS interfaces are needed for a variety of graphics hardware that will be supported both at the NRAO facility and at user institutes. The need for work on improved visualization methods was described in Section 7.1.2. The present AIPS effort is seriously manpower-limited in these areas. Local support of the workstations at the NRAO facility will also be needed, not only for image processing but also to exploit the workstations fully for other applications relevant to the array telescopes, e.g. engineering analysis and design.
- Scientific applications programming (2 positions). Initially, we will emphasize development of new code for wide-field imaging and for fast spectral line cube-building and analysis.
- Code vectorization and optimization (1 position). A dedicated effort in these areas, especially within AIPS, is essential both to maximize throughput in the NRAO facility and to ensure that NRAO code makes efficient use of architectures that are widely used elsewhere.
- Supercomputer support (1 position). As described in Section 6, access to supercomputers is an integral part of this proposal for extremely computer-intensive reductions and will require full-time support by one programmer. As noted in Section 6.3, this programmer should reside at a supercomputer center; their work will be closely co-ordinated with that of the team at the NRAO.
- Support of AIPS at non-NRAO user sites (1 position). This includes assistance with installation and maintenance of the AIPS code, as well as advice on peripheral interfaces and on "bug" handling.
- Documentation (2 positions). The AIPS package contains over 1.5 Mbytes of "help" text that provides on-line documentation, and there is a suite of manuals for users and for programmers (including users who wish to code their own applications). Increased effort is needed to ensure that documentation tracks the development of the software promptly and reliably. Production and maintenance of documentation describing the non-AIPS aspects of the new computer systems will also require greater effort than is now devoted to existing systems.
- AIPS code validation (1 position). The AIPS package is now sufficiently complex that it is no longer adequate to rely exclusively on exercising of the code by NRAO staff and by other users to reveal and report its deficiencies. A co-ordinated, dedicated effort is now necessary.
- Systems analysis (1 position). With the addition of a more sophisticated network of high performance computers, the present level of systems analysis staffing at the NRAO will be inadequate.
- Networking and Communications (1 position). The management of the internal network and of remote access to it from external networks and by modem will need additional support. NRAO support of computing at user institutes and supercomputer centers will

also depend heavily on good networking and communications.

- Management and supervision (2 positions). It will not be possible to expand either the AIPS staff or the computer operations staff on the scale envisaged here without also providing management and supervisory assistance, and developing improved code management procedures.

## 7.4 Algorithm Research

Image processing algorithms are so important to array telescopes that we must make special efforts to improve them. Despite the early successes, many problems remain and new avenues need to be explored. The largest uncertainty is the robustness of the procedures that are now used. Although deconvolution and self-calibration are vital to the success of array telescopes, research done so far has emphasized developing new methods, rather than the more difficult task of analyzing the stability of existing procedures. We plan special efforts aimed towards removing the art from this part of image processing. Examples of important problems are:

- Can we quantify the conditioning of the problem (i.e., whether the solution is well or poorly determined), so that the astronomer can be warned of the ill-advised use of a procedure?
- Mosaicing allows imaging of objects much larger than the primary power reception patterns of the elements of the array. Although this procedure is used increasingly often at the VLA to image objects of large angular size, neither the inherent limitations of the method nor the influence of (and implications for) array design have been adequately explored. This problem is acute at short centimeter and millimeter wavelengths, where the primary power patterns of array elements are restrictively small.
- At low frequencies, the wavefront distortions induced by the ionosphere are *non-isoplanatic*, i.e., radiation from different parts of the same object suffers different collimation errors. The self-calibration procedure must be augmented to allow for this. One attempt has been made to describe a suitable algorithm [11] but no implementation exists. Both this and the previous problem will afflict the VLA when observing at 75 MHz.
- Can we exploit techniques invented in other fields? For example, analogs of the closure phase of Jennison, which has been used in self-calibration, were discovered independently in two other disciplines: in X-ray crystallography as Hauptmann's structure invariant, and in optical imaging as Weigelt's triple-correlation. The latter technique is of some use for array telescopes when imaging weak objects in the presence of large, rapidly varying collimation errors [21].
- Radio images, though they often are viewed simply as pictures, are quantitative. (Typically, an image represents a statistical estimate of the brightness of some region of the radio sky, at a given wavelength.) There are two reasons why statistically sound error estimates do not routinely accompany these presentations of the data. First, the data processing operations are too complicated for any deductive error-estimation formula to apply. Second, because there usually is such a large volume of data, the basic data processing itself—let alone the error estimation—is extremely burdensome. Two techniques that could be used routinely to assign error estimates are the grouped jackknife and the bootstrap method of Efron [22]. These techniques (especially the latter one) are extremely computer-intensive,

however—even in much simpler settings than Fourier synthesis—and further research is needed.

Systematic exploration of effective algorithms is a crucial step in developing powerful and efficient code to handle the most difficult computing problems associated with Fourier synthesis telescopes. In view of this, we propose to add one new research scientist and two support programmers to the existing algorithm research effort at the NRAO. Because this research will benefit other fields of astronomy and image processing, funding for this part of the proposal will be sought separately, from the NSF's Science and Technology Centers program. It is therefore not included in the budgetary estimates given in Tables 5 and 6.

## 8. BUDGET

The tables below summarize cost and manpower estimates for all aspects of array telescope computing support described in this proposal. These are: construction and operation of a computing facility at the NRAO; direct NRAO support of computing at users' home institutes; support of computing by NRAO users at supercomputer centers; and software and algorithm development and optimization.

Table 3: Required Capital Equipment Investment

| Item                       | 1989        | 1990        | 1991        |
|----------------------------|-------------|-------------|-------------|
| Mini-supercomputers        | \$3,000,000 | \$1,900,000 | \$1,900,000 |
| Workstations – display     | \$300,000   | \$400,000   | \$300,000   |
| Archive/mass store         | —           | \$1,000,000 | \$700,000   |
| Hard copy equipment        | \$50,000    | \$125,000   | \$125,000   |
| Network and communications | \$100,000   | \$50,000    | \$50,000    |
| Total Required*            | \$3,450,000 | \$3,475,000 | \$3,075,000 |

\*including anticipated discounts

Table 4: Proposed Capital Equipment Funding Schedule

|                            | 1989        | 1990        | 1991        |
|----------------------------|-------------|-------------|-------------|
| Total Capital Needed       | \$3,450,000 | \$3,475,000 | \$3,075,000 |
| Less VLBA Equipment Budget | —           | \$1,500,000 | \$1,500,000 |
| Total New Capital Funds    | \$3,450,000 | \$1,975,000 | \$1,575,000 |

Table 5: Operating Costs

|   | 1989               | 1990               | 1991               | 1992*              |
|---|--------------------|--------------------|--------------------|--------------------|
| <b>(a) Personnel:</b>                   |                    |                    |                    |                    |
| Total Personnel Costs                   | \$1,800,000        | \$2,110,000        | \$2,680,000        | \$3,050,000        |
| Less Current NRAO personnel             | \$1,340,000        | \$1,410,000        | \$1,480,000        | \$1,550,000        |
| Less VLBA Budgeted                      | \$50,000           | \$100,000          | \$310,000          | \$590,000          |
| <b>New Personnel Funds Needed</b>       | <b>\$410,000</b>   | <b>\$600,000</b>   | <b>\$890,000</b>   | <b>\$910,000</b>   |
| <b>(b) Maintenance/Operations:</b>      |                    |                    |                    |                    |
| Operation of Old Equipment              | \$1,145,000        | \$955,000          | \$940,000          | \$875,000          |
| Operation of New Equipment              | \$365,000          | \$925,000          | \$1,415,000        | \$1,705,000        |
| <b>Total Operation of Equipment</b>     | <b>\$1,510,000</b> | <b>\$1,880,000</b> | <b>\$2,355,000</b> | <b>\$2,580,000</b> |
| Less Presently Budgeted                 | \$1,160,000        | \$1,235,000        | \$1,310,000        | \$1,385,000        |
| Less VLBA Budgeted                      | \$60,000           | \$120,000          | \$360,000          | \$665,000          |
| <b>New M/O Funds Needed</b>             | <b>\$290,000</b>   | <b>\$525,000</b>   | <b>\$685,000</b>   | <b>\$530,000</b>   |
| <b>(c) Total of (a) and (b):</b>        |                    |                    |                    |                    |
| <b>Total New Operating Funds Needed</b> | <b>\$700,000</b>   | <b>\$1,125,000</b> | <b>\$1,575,000</b> | <b>\$1,440,000</b> |

\* and beyond



Table 6: 1992 Manpower – Operations and Production Software

| Activity                 | Total Needed | VLBA funded | Current NRAO | Addition Proposed |
|--------------------------|--------------|-------------|--------------|-------------------|
| On-line computers        | 5            | 2           | 3            | 0                 |
| Systems Analysis         | 4            | 1           | 2            | 1                 |
| Operations               | 9            |             | 6            | 3                 |
| Applications Programming | 9            | 3           | 4            | 2                 |
| Graphics, Workstations   | 3            |             | 1            | 2                 |
| Optimization             | 1            |             |              | 1                 |
| Code Validation          | 1            |             |              | 1                 |
| User Support             | 5            | 1           | 3            | 1                 |
| Supercomputer Access     | 1            |             |              | 1                 |
| Documentation            | 3            |             | 1            | 2                 |
| Communications           | 2            |             | 1            | 1                 |
| Maintenance              | 8            | 3           | 5            | 0                 |
| Management, Supervision  | 6            | 1           | 3            | 2                 |
| <b>Total</b>             | <b>57</b>    | <b>11</b>   | <b>29</b>    | <b>17</b>         |

Table 7: Manpower – Algorithm Research\*

| Activity     | Total Needed | Current NRAO | Addition Proposed |
|--------------|--------------|--------------|-------------------|
| Scientists   | 2            | 1            | 1                 |
| Programmers  | 2            |              | 2                 |
| <b>Total</b> | <b>4</b>     | <b>1</b>     | <b>3</b>          |

\* To be funded separately, through Science and Technology Centers Program

## Bibliography

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, pp. 1293-1322 (1983); "The Very Long Baseline Array", by K. I. Kellermann and A. R. Thompson, *Science*, Vol. 229, pp. 123-130 (1985).
2. "Adaptive Calibration of Radio Interferometer Data", by F. R. Schwab, *Proc. SPIE*, Vol. 231, pp. 18-25 (1980).
3. "Image Formation by Self-Calibration in Radio Astronomy", by T. J. Pearson and A. C. S. Readhead, *Ann. Rev. Astron. Astrophys.*, Vol. 22, pp. 97-130 (1984).
4. "Superluminal Expansion of Quasar 3C273", by T. J. Pearson, S. C. Unwin, M. H. Cohen, R. P. Linfield, A. C. S. Readhead, G. A. Seielstad, R. S. Simon and R. C. Walker, *Nature*, Vol. 290, No. 5805, pp. 365-368 (1981).
5. "New High Resolution, High Dynamic Range VLA Images of the M87 Jet", by F. N. Owen, T. J. Cornwell and P. E. Hardee, (1987, in preparation).
6. "A National Computing Environment for Academic Research", by K. K. Curtis, National Science Foundation, July 1983.
7. "The Physical Process Which Shapes Cassiopeia A", by R. Braun, S. F. Gull and R. A. Perley, *Nature*, Vol. 327, pp. 395-398 (1987).
8. "The Jet and Filaments in Cygnus A", by R. A. Perley, J. W. Dreher and J. J. Cowan, *Astrophys. J.*, Vol. 285, pp. L35-L38 (1984).
9. "The Faraday Rotation of Cygnus A", by J. W. Dreher, C. L. Carilli and R. A. Perley, *Astrophys. J.*, (1987, in press).
10. "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.
11. "Relaxing the Isoplanatism Assumption in Self-Calibration; Applications to Low Frequency Radio Interferometry", by F. R. Schwab, *Astron. J.*, Vol. 89, No. 7, pp. 1076-1081 (1984).
12. "A Search for Gravitational Lensing", by J. N. Hewitt, Ph. D. Thesis, MIT (1986).
13. "OH/IR Stars Within 50 pc From the Galactic Center", by A. Winnberg, B. Baud, H. E. Matthews, H. J. Habing and F. M. Olmon, *Astrophys. J.*, Vol. 291, pp. L45-L49 (1985).
14. "The Distance to the Center of the Galaxy" by M. J. Reid, M. H. Schneps, J. M. Moran and C. R. Gwinn, *I.A.U. Symposium 125*, pp. 554-555 (1987).
15. "Data Display: Searching For New Avenues in Image Analysis", by A. Rots, in *Synthesis Imaging*, Proc. NRAO Workshop No. 13, ed. R. A. Perley, F. R. Schwab and A. H. Bridle, pp. 231-252 (1986).
16. *Serendipitous Discoveries in Radio Astronomy*, ed. K. I. Kellermann and B. Sheets, Proc. NRAO Workshop No. 7 (1983).
17. "Large, Highly Organized Radio Structures Near the Galactic Center", by F. Yusef-Zadeh, M. Morris and D. Chance, *Nature*, Vol. 310, pp. 557-561 (1984).
18. "Astronomical Requirements for Future VLA Processing", by R. D. Ekers, E. B. Fomalont and F. N. Owen, VLA Scientific Memorandum No. 150, September 1983.
19. "A Computer Plan for the VLA", by R. Duquet, G. Hunt and R. Burns, VLA Computer Memorandum No. 168, September 1983.
20. "The 1986 AIPS Site Survey", by A. Bridle, D. Wells and N. Wiener, AIPS Memorandum No. 52, June 1987.

21. "Radio Interferometric Imaging of Weak Objects in Conditions of Poor Phase Stability: The Relationship Between Speckle Masking and Phase Closure", by T. J. Cornwell, *Astron. Astrophys.*, Vol. 180, pp. 269–274 (1987).
22. "Computer-Intensive Methods in Statistics", by P. Diaconis and B. Efron, *Scientific American*, Vol. 248, No. 5, pp. 116–130 (May 1983).

# APPENDIX A

## CHRONOLOGY OF NRAO ARRAY TELESCOPE COMPUTING

---

|      |           |  |
|------|-----------|--|
| 1969 | January   | Proposal Vol. III specifies computing for 27-antenna VLA; \$3.2 M budgeted.  |
| 1972 | August    | Congress approves construction of VLA; \$4.5 M budgeted for computing.   |
| 1974 |           | CLEAN deconvolution algorithm published.   |
| 1974 | June      | DEC awarded subcontract for initial continuum off-line computer, a DEC-10/KI.  |
| 1974 | July      | Modular Computer Systems awarded subcontract for on-line computer; Initially a network of 4, later 7, ModComp II minicomputers.  |
| 1974 | October   | SAIL analysis and imaging software for DEC-10 begun.   |
| 1975 | December  | First 2D maximum entropy deconvolution algorithm published.  |
| 1976 | February  | Contract with ERIM to design optical processor for VLA image generation.   |
| 1977 | May       | First VLA image of an extended source.   |
| 1977 | July      | EXPORT facility at VLA allows first off-site processing.   |
| 1977 | December  | Optical processor abandoned for pipelined network of PDP-11's and special purpose hardware — the VLA "Pipeline".   |
| 1979 | February  | AIPS post-processing project begun in Charlottesville; code developed in Modular Computing Systems computer with FPS120B AP.   |
| 1979 | July      | Self-calibration algorithm available in Charlottesville ModComp.   |
| 1979 | December  | Acquire VAX 11/780 with AP in Charlottesville.   |
| 1980 | December  | VAX 11/780 with AP delivered to VLA site.  |
| 1980 | December  | Final VLA computer construction budget increased to \$6.5M.  |
| 1981 | January   | First use of AIPS at VLA site.   |
| 1981 | January   | Decision to include general image formation in AIPS.   |
| 1981 | March     | DEC-10/KI upgraded to DEC-10/KL.   |
| 1981 | April     | Acquire second VAX 11/780 with AP at VLA site.   |
| 1982 | March     | NRAO Computer Advisory Committee meets at VLA.<br><i>"Need a long range plan based on astronomical requirements ... requires a major new infusion of capital from NSF."</i>  |
| 1982 | March     | First UNIX implementation of AIPS exported from U. Texas.  |
| 1982 | November  | First Maximum Entropy deconvolution, H2MEM, available in AIPS.   |
| 1983 | January   | Decision to include primary calibration in AIPS.   |
| 1983 | September | NRAO study of VLA computing finds Class VI supercomputer capacity needed.  |
| 1983 | September | First wide-field imaging algorithm, MX, available in AIPS.   |
| 1983 | October   | DEC abandons DEC-10 line of computers.   |
| 1983 | October   | Scientific Review Committee assesses NRAO's computing plans, concludes:<br><i>"We are convinced of the need for computing capacity at least in ... the small supercomputer range ... the VLBA will add to this need".</i><br>Committee urges the NRAO to make such a plan. |

---