CPG MEMO NO. 28

VLA COMPUTER MEMORANDUM NO. 170

Radio Imaging at NRAO

R. H. Miller

23 February 1984

0. Summary

The VLA is not realizing its scientific potential because of limitations set by the installed computing power available to reconstruct images of sources. Computing power adequate to the VLA's needs will catapult NRAO into a group that operates supercomputers. This statement is based on user assessments of the kinds of observing programs the VLA should be doing now, expressed in terms of the computer power required to process the images generated in those observing programs. Future synthesis arrays (VLBA, millimeter, combined VLBA-VLA array, satellite VLBI, optical interferometry) will have almost identical image reconstruction problems, and each will add to this load.

An image reconstruction facility within NRAO, serving needs from all NRAO observatories, appears to be a reasonable way to meet these needs. The facility would include major computational and graphics resources. It will be large and would represent a significant increase in the total size and funding level of NRAO.

Major reassessments of NRAO's perception of itself--its goals, its standing, and its role in astronomy and within the scientific community as a whole--are implied in such a move. The process must begin with NRAO's coming to grips with its perception of itself and to reconciling itself to the notion that it is now in the scientific big leagues by virtue of its past accomplishments and present capability. NRAO must begin an aggressive campaign to sell this viewpoint to the radio astronomical community, the astronomical community, and the scientific community as a whole. This should be fairly easy if NRAO plans boldly and capitalizes on its present strengths, accomplishments, scientific productivity, and capability.

The need is urgent. Scientific use of the VLA is crippled now by lack of adequate facilities. All future synthesis arrays similarly risk being unable to realize their full scientific potential if the problem is not straightened out for the VLA. The costs of doing nothing can be even greater than the costs of establishing a major imaging facility.

Indeed, NRAO must make some visible effort to provide good image construction and analysis facilities adequate to the needs of the VLA. NRAO cannot afford to settle for less than the maximum possible scientific utilization of its facilities. NRAO must be on record as having made every effort to make the full power of the VLA available to all qualified users.

1. The Problem

Scientific use of the VLA (and, presumably, of other synthesis arrays yet to be built) takes radio images as the principal observational material on which scientific inferences are based. A synthesis array like the VLA produces a huge volume of data in a form that may be considered to be an encoded image, contaminated by noise. The first step in doing science is to recover an accurate image from that noise-contaminated, encoded form, and to present it in a form that is easily recognized and interpreted by humans. The principal problem is to present the data to the observer in a manner that helps him recognize the scientific problems. Often, the observer presents the data in many different ways in an attempt to picture the source accurately. The observer should not have to waste effort trying to cut through awkward features of a display in order to get at the science underneath. Doing science is a somewhat fragile process, and any difficulties in getting to the true character of the source interfere with the scientist's ability to get down to the real problem at hand--what makes that source tick?

In the final analysis, the scientific power of the VLA is measured in terms of the quality (and quantity) of science produced. That power is severely limited if it is difficult to cast the observational results into a form comprehensible to the observer and to his fellow scientists. New, even daring, algorithms and display forms are also needed to discover unanticipated results.

An image, map, or picture, is the form most easily interpreted by humans. Radio sources often take unfamiliar forms, and visual presentation is the best way to inform people of those unfamiliar shapes. Few observations are made with the VLA that do not require image reconstruction. Further, it must be stressed that the images are quantitative. The images produced are as quantitative as a geologist's map with altitude contours and rock or soil types all presented in a single map. These images are the basic working tool of today's radio astronomer--his way of studying the Universe in a precise, quantitative way.

Like the geologist's map, a radio image contains a tremendous amount of information. The vast amount of data produced by the VLA in a short time, and the huge amount of information present in the images, define the problem of radio imaging. It is, quite simply, the problem of making all that data comprehensible to the observer--and, through the interpretations provided by the observer, to the scientific community in general and finally to the public at large--that we must face here. This is the process we are trying to facilitate.

There is a bottleneck in that chain between the VLA observations and the observer's ability to interpret them. The sheer volume of data to be handled, to be decoded and cleaned of noise, is beyond the capacity of the system that has been set up to handle it. In today's world, large volumes of data are usually handled by computers, and radio imaging at NRAO is no exception. The bottleneck is best described in terms of the computer power that would be required to break it.

2. A Market Survey

A very good start at describing the computer power needed to provide the necessary imaging power has been made by the VLA staff. Ekers, Fomalont, and Owen (VLA Scientific Memorandum No. 150, "Astronomical Requirements for Future VLA Processing," September 1983, hereafter EFO) described the astronomical needs and Duquet, Hunt, and Burns (VLA Computer Memorandum No. 168, "A Computer Plan for the VLA," September 1983, hereafter DHB) restated that astronomical requirement in terms of the computer power required to cope with it. They give a figure for the number of computer operations per second needed to cope with the steady-state data rate generated by VLA observations. That estimate defines the scale of the problem in computational terms, as 60 to 80 MFLOPS (million floating-point operations per second) minimally. It forms the basis of the remaining considerations of this note. We accept this estimate, and use it as >80 MFLOPS. (I detest the term, "MFLOPS," and use it only because it is standard. A flop means a failure to me, and megaflops implies millions of failures--hardly the message one wants to convey.) Furthermore, the estimated computer power requirement applies the present VLA operation (some options available to the VLA have not yet been implemented because of computer limitations); improvement over the next decade and a joint VLA-VLBA operation will increase the compututational requirement significantly.

Several comments are in order concerning these estimates. First, and most importantly, they reflect a user's perception of what the VLA should be doing and of current limitations to the scientific capability of the VLA. Second, they are current needs. They tell what the VLA ought to be doing now. It is unfortunate that the term, "future," slipped into the title and, quite possibly, into the thinking about the problem discussed in these two reports. Third, 80 MFLOPS is a time-averaged, steady processing rate, not a peak rate in a burst mode of operation. It will continue on a 24-hour-a-day, 7-day-a-week basis.

A fourth comment relates to a possible misconception that these reports do little to prevent. Table 4 of DHB relates the fraction of observational programs to the fraction of image-making resources consumed. That table shows that 7.5% of the most computationally intensive observing programs would consume 50% of the image-making resources. The point that is obscured in such a tabulation is that that 7.5% is likely to contain the most exciting and most important new science. It contains the observing programs that cannot now be done because of limited resources and includes qualitatively new kinds of observations that cannot even be attempted now. If the resources were available, perhaps 50% of the observing programs would need a large computer. Taken as an excuse to limit observing programs to those that can be handled with present resources, Table 4 could prevent the VLA from doing the most important science within its capability. It could cripple the VLA scientifically. The implications for future synthesis arrays are devastating.

We stress that this estimate of required computer power, needed to provide the necessary image-making power, is a user's estimate, drawn from the radio astronomical community's actual (and planned) usage of the VLA. It is not a number invented by the writer of this report. And it is very likely to be an underestimate because it is impossible to gauge the real needs of the radio astronomical community. Their usage of the VLA has been depressed because they realize that there is not enough image processing power to carry out the more demanding tasks. They do not even propose observations that would place heavy demands on image reconstruction capability. Complex sources are the most difficult to reconstruct. These are sources that are extended, that have complicated structure, and that have large dynamic range. The term, "dynamic range," as used here refers to the ratio of brightnesses within an image. One may want to "see" a faint source near a bright one. It is not unusual to seek structures 1000 times fainter than the brightest spot within a VLA map. The center of our own Galaxy is a good example of this kind of source.

The ultimate peril in all this is that the best radio astronomers, conscious of this limitation (whether real or perceived) will simply turn to other problems and that the most imaginative and most important problems in radio astronomy may simply not be done (at least not in the U.S.). One of the most beautiful and powerful instruments in all U.S. astronomy, the VLA, could remain crippled because it lacks the capability to carry certain important observational programs through to completion. Crippling will extend to other planned synthesis arrays like VLBA and millimeter array. They are unlikely to be adequately equipped if the VLA is not. The danger is real.

3. Immediate Consequences

The required computer power, obtained from careful translation of astronomical requirements (EFO) into computer power by DHB, catapults NRAO into the range of supercomputer users. 80 MFLOPS is near the top speed attainable on a Cray 1, taking full advantage of the special properties of the Cray. Probably no computer user anywhere has attained a sustained computation rate this high on a Cray. The situation is essentially the same with other supercomputers. (We single out Crays for the sake of a definite example. It is not at all clear that a Cray would solve the imaging problems faced by NRAO.)

The important point is that the required computer power, and the required image processing power, be made available. The precise form which that power takes (e.g., hardware, graphics, and software configurations) is unimportant as long as the power is there and usable.

It is premature to specify a "hardware" design, and I will not attempt to do so in this note. However, there are several matters that should be borne in mind as the system is specified that I would like to stress here. Not all of these have appeared in VLA reports, and the emphasis given those that have is often different from that given here.

Before going into more detail, it may be instructive to see how the present situation came about.

4. How NRAO Got Into This Bind

Image reconstruction was perceived to be a fairly simple and straightforward operation at the time the VLA was designed. Attention was quite properly concentrated on radio engineering aspects. The system was planned for a dynamic range of only about 100:1. The VLA is also an accurate spectrometer, a capability that was not stressed in its original design.

It is a testimony to the skill and wisdom of those who originally designed the VLA that it now routinely outperforms those original design specifications. The system is beautifully engineered. The VLA has proved flexible enough to meet the demands of the better receivers and vastly improved image reconstruction methods now in use. Imaging with dynamic range far in excess of the originally specified 100:1 is now routine. Dynamic ranges as large as 10,000:1 are obtainable if the observer is willing to put in the effort, and time on the image reconstructing computers is available.

The greatest contributor to this improved performance was simply that scientists learned how to use synthesis arrays. They learned how to CLEAN and to "self-calibrate" images. They discovered new modes of operation, such as snapshot and spectroscopy. These are improvements undreamed of at the time the VLA was designed. But the improvements demand a great deal of computation. NRAO is using these computationally intensive techniques without much more installed computer power than that envisioned at the time the VLA was designed and built. This quite naturally leads to a bind. The bind was somewhat unavoidable as well. During construction in the 1970's, the necessary computer power was not available, even though some within NRAO argued that it was needed. Telescope and electronics hardware had the top priority. Now the computational lack comes back in the form of limited imaging capability. Now is the time to remedy this lack.

Imaging techniques are so vital to doing science with a synthesis array that computers for image reconstruction are as important as other hardware components of the VLA. This is true of any synthesis array. The principal difference between computer power for imaging and other hardware components is that hardware that provides computer power can be shared with other synthesis arrays.

No reasonable alternative to computational methods is known that will yield the dynamic range routinely obtained in image reconstruction at the VLA. Optical methods have been used in other applications, but they do not include the special techniques used with radio-astronomical imaging, such as CLEAN and SELF-CAL (self-calibration). They are limited to low dynamic range. Even if the techniques could be carried over to other processing methods, the costs of doing so are likely to exceed the costs of an "adequate" computer image reconstruction facility by a wide margin. These devices also tend to be inflexible in the face of constantly changing needs. NRAO does not perceive itself as being able to afford to undertake the development work necessary to bring a different kind of system up to the quality of performance routinely achieved in its imaging work.

5. Present Situation

NRAO is now doing very well in terms of getting science out with modest image reconstruction equipment. Users are more or less satisfied. Attempts to convince the user community, the radio astronomical community, and others to support expanded imaging facilities can only be made more difficult by that general user satisfaction. For example, not all users will go along with the needs as set forth by EFO. Some users feel that the present image reconstruction facilities are adequate. However, the arguments of EFO are carefully considered and are directed toward realizing the scientific potential of the VLA as fully as possible within reasonable cost constraints. The limited resources presently available are not adequate for the imaging job the VLA presents today. Users who are satisfied with present imaging facilities are not those who are pressing the VLA to do new kinds of science.

Users seem comfortable with AIPS, NRAO's image processing package of computer programs. It does most of the things most of them want done. NRAO is to be complimented for this accomplishment. User satisfaction is a goal NRAO should strive to retain through any significant increase in imaging power.

On the other hand, NRAO is not doing very well in terms of using the presently installed computer power effectively. The array processors are actually used about 25% of the time within AIPS programs (this figure comes from DHB, who ran tests designed to make the number come out as high as possible). This poor efficiency results from hardware limitations on array processors that were commercially available at the time NRAO got into the business of using array processors. However, it indicates that both hardware and software designs will have to be re-thought to fit into an environment where 60 to 80 million floating-point operations are going on every second. Since the bulk of the major VLA imaging problems are well-known and few in number, efficient use of the computer hardware should not be difficult to achieve. It will require programming manpower, however.

VLA operations are hampered now by lack of image reconstruction power. Some observing proposals are rejected or restricted, others simply never get proposed in the first place. It is, however, difficult to get firm readings on just how severe a pinch this is and on how the user community copes with it. A question remains about how long this situation can or should be tolerated by a community anxious to do science with the VLA. NRAO must publicize this situation and lead the community in using the VLA for new science. Members of its staff know the present limitations better than much of the outside community. They are also more hampered by those limitations.

Users are not demanding as much as the VLA is capable of delivering. It is difficult to tell whether user satisfaction implies that the services now provided by NRAO are really good or if users are satisfied with something inadequate and simply don't know what current technology could provide. Perhaps the brute force sensitivity and resolution of the VLA have enabled good science to be obtained with modest computational power. Now that the obvious experiments have been done, new, more computationally intensive projects are appearing.

These projects are at the forefront and should be encouraged. On the other hand, NRAO may simply be catering to the demands of the "average" or "typical" user, rather than leading the user community and making them aware just what the VLA could be doing. If so, NRAO is not cultivating the demanding user. Conversely, the user community is not as demanding as it should be in insisting that the scientific capabilities of the VLA be fully realized. Whatever the reasons, the VLA is not as productive scientifically as it can and should be, and that is a pitiful state of affairs.

6. Some Design Considerations

A lot of thinking and design must go into an image reconstruction facility before it is seriously started.

a. I/O. I/O considerations are likely to dominate the final design. They are likely to be more important than the kind of processor selected. A processor that can run steadily at 80 MFLOPS will not do so unless data are available to it when it wants them. If it must wait for I/O, it is idle and the number of true operations per second decreases accordingly.

The basic consideration is the average number of arithmetic operations the central processor carries out per number moved in or out in I/O operations. In radio astronomical image reconstruction, that ratio tends to be around 100, much lower than the thousands in typical supercomputer programs. This is the origin of the I/O problem. It is the reason why I/O considerations are likely to dominate the design.

At 100 operations per number moved, the I/O processor must move 1 million words per second into staging memory and another 1 million words out, for a total I/O rate of 2 megawords/sec. With 32-bit words, that is 8 megabytes/second. Call it 10 megabytes for a round number.

The I/O processor that handles this will not be a small object. In Cray terms, a solid-state memory (SSD) at 24 million words (not addressable by the central processor) as a staging memory would help considerably, but one still has to move data into and out of it at 2 million words per second, which is now 16 megabytes/second because of the Cray's 64-bit word. If I/O is handled by VAXes, some 10 to 20 VAXes would be required just to handle this I/O rate.

Don't ignore the I/O problem!

b. Bigger Memory? I/O requirements are intimately tied in with memory size. The two cannot really be considered separately. However, some statements can be made about memory size-I/O rate tradeoffs.

First, it is not reasonable to consider a design that will not contain an entire map within the memory space directly addressable by the processor (by force of habit, I'll call that "core".) Depending on map sizes, that can range up to 4 million words for a 2K by 2K map. Staging memories are needed to keep the processor running (something like buffers, but likely to be considerably larger than your standard buffer), so main memory contains the map currently being processed, part of the last map processed being staged out, and part of the next map in sequence being staged in. The staging memory can probably be combined at some cost in programming development. We now have about 12 million words for our 2K by 2K map.

I/O rates could be reduced considerably if more arithmetic operations could be done on a map before it is moved back out. Ideally, a user's map should remain in core, without need to roll it out, until processing is complete. But this is probably prohibitively expensive of main memory.

However, a considerable step in this direction can be made if enough memory is available to do a CLEAN operation. This seems to require 4 times the map area (16 million words for our 2K by 2K map). With that much memory, the number of operations performed per number moved in or out rises to a thousand or more. That would allow reduction of I/O rates by a factor 10, to around 1 megaword/second rates. This lower rate is much easier to sustain with reasonable design. But staging areas would still be needed, and they now run to a total 48 million words, which you may as well call 50.

c. Map Sizes. The statement that the entire map must be core-contained clearly cannot continue to arbitrarily large maps. A facility should probably be designed to accommodate a reasonably large "standard" map size. 2K by 2K seems a reasonable target. However, 8K by 8K maps will become common, especially at frequencies below 1.7 GHz. A system too highly tuned to 2K by 2K would be too limiting. User demands for ever larger map sizes should be studied carefully from an information-theoretic point of view, both theoretically and empirically.

d. Overhead. The effectiveness with which the processors are kept operating is a crucial point in the overall design. As mentioned earlier, the array processors in the present AIPS configurations are probably not operating over 25% of the time, and this without a good measure of the speed being achieved while they are operating. In fact, if all the array processors now running at NRAO could be kept busy and running at full speed, the image processing load could be adequately handled. The factor that prevents this kind of operation seems to be I/O speed and directly addressable memory. These points have already been discussed in a-c above. We mention the point once more here to stress its importance in an overall view of the design.

e. Load Leveling. Not all VLA observations have the heavy computational requirements of DHB's "canonical task." This holds forth the promise that observations that impose a heavy computational load to reconstruct images could be processed at times the VLA (and other synthesis arrays) are generating data that do not produce as large a computational overhead. The fact that image reconstruction is intrinsically not a "real-time" operation at NRAO (yet!) encourages this hope.

Load leveling is, however, not a real possibility. The arguments underlying the specification of a "canonical task" in DHB have already taken it into account. Little more can be achieved in this direction according to our current understanding of the problem.

f. Reserve Capacity. NRAO traditionally has not planned computer installations with adequate reserve capacity. This is part of the reason for the present bind. Reserve capacity should be regarded as analogous to a safety factor in mechanical designs.

Many considerations enter into estimates of adequate reserve capacity. The primary considerations are (1) to allow room for growth of load, and (2) software costs incurred in trying to achieve near-design performance of installed hardware. Growth of load must include new kinds of applications, such as snapshot mode, CLEAN, SELF-CAL, and so on. It also includes any (currently unforeseen) applications that may arise. Most installations argue for increased computer power on the basis of extrapolations of actual load growth.

A factor 3 to 4 over the stated need seems to be a reasonably conservative reserve capacity allowance. That makes the 60 to 80 MFLOPS grow into 200 to 300 MFLOPS as a more realistic estimate.

g. Tradeoffs: Hardware-cost against Software-cost. Software development costs tend to run as large as hardware costs in many installations, even when production codes achieve perhaps 25% of the design speed of the installed hardware. Software development costs rise very rapidly as one tries to achieve a higher fraction of hardware utilization. Software costs could triple or quadruple to bring the hardware utilization up to 50%.

It is not cost effective to cut corners on hardware investment. Software costs are often seriously underestimated, and anything that can help to contain software costs is a good investment. It is safer to err on the side of greater hardware investment.

h. Graphics. Present graphics facilities at NRAO could not serve a faster or more complete image reconstruction facility. Individual frames of high quality can be produced, but they cannot be produced fast enough to meet even current demand.

As an example, galaxy images with velocities are basically three-dimensional. Intensity is represented as a function of three (or more) variables. A useful way to display such results is to use motion to help visualize the three-dimensional structure. Rotating the image has proven useful in other applications to help appreciate the intricacies of a three-dimensional structure. However, it is sufficiently difficult to chain a sequence of views into present-day NRAO graphics devices that this kind of display has been seldom, if ever, used. I/O rates required to support an adequate graphics system from an imaging computer are comparable to those estimated in Sec. 6a for data flow within the computer system itself.

A useful adjunct to other NRAO activities would be to have a small group actively working at making graphics equipment and software ever easier to use. As new graphics devices come on the market, NRAO should evaluate each to determine whether it could usefully help with NRAO imaging. That, by itself, can be a relatively expensive undertaking, depending on how thoroughly it is done.

7. How Will a User Cope?

The image reconstruction power suggested by EFO is nearly 100 times the present power. This kind of power could easily overwhelm a user, unless distributed among more than a few users.

Careful consideration must be given to helping the user cope with this system. Are 100 times as many users going to be on the system at once, each seeing something like the present speed and display capability? Or is one user going to see his maps come back 100 times faster? If the latter, can he even look at a map before the next one comes back? There are important psychological issues to match the image reconstruction capability to a user's capability to absorb and interpret images. The spirit must be that we want to help the user do science.

It is important that any new system be designed to keep the user perception of a friendly environment. Again, science is fragile, and the user must feel that he can concentrate on the science. He must not feel that most of his efforts are going toward coping with a difficult system.

Users may want to operate the system remotely ("dial in"), to use the system in conjunction with displays at their home institutions. This kind of operation would relieve some pressure on users, but at a cost of requiring high-speed data communications.

Many of the maps produced will be multi-dimensional. Users must be helped to understand what these maps contain. This will require substantially more powerful and higher resolution graphics capable (minimally) of displaying three dimensional data (RA, dec, velocity, for example). More and better display devices will be needed. Software to exploit the capabilities of these devices, and to make them easy and comfortable to use, will have to be developed. This will cost something, and that cost must be considered.

8. Implications for NRAO/VLA Operations

As mentioned above, a facility of this size catapults NRAO into the supercomputer class. Experience at other installations that have this kind of computer power indicates that substantially increased annual running costs and staff sizes may be expected. Compare NCAR (but be careful: their accounting procedures differ from NRAO's). It could entail doubling present computer staff and computational budget. This increase in scope of operations cannot easily be accommodated into NRAO's present organizational structure (financial and personnel).

It is very important that it be realized at all levels what the proposals for greatly improved image reconstructing power imply. For example, an improved CLEAN algorithm could be worth $1 million per year to NRAO. NRAO must change its perception of its scope of operation to go along with this increase of activity.

This is not a small operation. There is no way it can be kept small. It will have to include additional graphics capability along with increased computer power. A balanced facility is necessary, providing all the needs of users as they try to interpret their observations.

Planning will require a lot of effort. There are people on NRAO staff now who are quite capable of doing the planning, but they are busy with day-to-day activities and have little time for planning. Some arrangements for staff to plan a facility will be required.

Selling this to the radio astronomical community, and getting funding will require major efforts. If you are serious about this, groundwork must be started immediately.

On the other hand, doing nothing can cost even more than creating an imaging facility. The cost of doing nothing is that the VLA will continue to be crippled scientifically, and newer synthesis arrays will similarly be unable to realize anything like their full scientific potential. The cost of doing nothing is that NRAO can become branded as a batch of gadgeteers, interested in building new and challenging instruments, but with little interest in their scientific use once they have been built. It would be dangerous for NRAO to be so perceived by the scientific community. NRAO justifies its funding on the basis of the science produced, and this would run the risk that NRAO would be seen as not primarily concerned with doing science and seen as not serving its user community well. Synthesis mapping is here today and will be done more and more in the future with or without NRAO and the US astronomical community. A cost of doing nothing is to abandon synthesis mapping for others to take up instead.

Funding at a level that would not allow development and operation of a full radio imaging facility is a dangerous possibility that NRAO should try to prevent. The arguments on this point must be based on doing science with the VLA. Presumably, funding at a lower level would permit some enhancement of capability without going all the way to a full facility. The VLA (and future synthesis arrays: VLBA, millimeter array, combined VLA-VLBA, etc.) could be operated under whatever limitations this implies. Whether such limitations are tolerable hinges on the kinds of science that can and cannot be done at a given level of funding for imaging. These are contingency planning arguments, which should be prepared beforehand in hopes that they will never be needed.

The issue of urgency raises similar considerations. How long can the VLA be permitted to operate at this reduced level? How long will users tolerate this kind of operation? How do you convince users that they should not tolerate it? Would one seriously consider shutting the VLA down, or not proceeding with VLBA, if no additional imaging facilities become available?

9. Suggestions

After all this, let me offer a few suggestions.

a. NRAO should plan an image reconstruction facility, to serve VLBA and millimeter array in addition to VLA. VLA is the current proving ground, the area in which the need for greater power first became apparent because it is running now. It can also serve those other arrays, since their image construction problems are very similar.

The image reconstruction facility should be thought of in terms of imaging, and not in terms of computing. Even though there is a heavy computational component, emphasis must be kept on the science that is the end goal.

NRAO's need is great enough to justify a major facility for its own use. NRAO generates a large enough load to keep a very large image reconstruction facility busy now. The facility might also serve other astronomical users and other kinds of astronomical imaging requirements, but any shared facility would have to be much more powerful to handle the additional load brought in by those other users. New arrays now being planned or designed will increase the demand. NRAO cannot afford to wait while other groups decide what they want or need.

b. The need is urgent. How long can NRAO tolerate operation of a crippled VLA? Can NRAO risk bringing VLBA into operation without adequate imaging facilities to handle its output? It is dangerous to continue planning for VLBA while you cannot cope with the data produced by facilities that are running now (the VLA), and before you have plans for bringing that facility up to speed.

c. The place to begin is with groundwork. A facility this size and cost cannot be sprung on NSF without preparing the groundwork. Backing and support of the radio astronomical community, and of the astronomical community as a whole is needed. But before the radio astronomical community can be asked to support such an effort, NRAO must get its act together and decide whether it is ready to take on a program this size. NRAO has already done larger projects than this. It is a small program compared to the VLA or VLBA. But an imaging facility this size is breaking new ground. Nothing like it exists today, and nothing like it has been proposed.

d. NRAO should start immediately to reassess its perception of itself and of where it fits in the scientific community at large--its staff, users it serves, radio astronomers, astronomers in general, the scientific community as a whole, AUI, NSF, the congress and the public. NRAO's scientific achievements are impressive. The design, construction, and successful operation of an instrument like the VLA are significant accomplishments of which you may be proud. NRAO's self-image is too modest. NRAO should crow a bit. Staff contacts with the public and with public relations programs should be encouraged.

A new kind of facility like this, and one this size, cannot be sold by a timid organization. It cannot even be accommodated within a timid organization. NRAO's present self-image is too meek and timid.

NRAO should be bold and assertive in selling this project as well as in proclaiming its scientific achievements. You're as good as (if not better than) a lot of groups that trumpet themselves loudly.

e. An aggressive sales effort should be started to prepare all the groups mentioned in parts (c) and (d) for NRAO's entry into this more costly mode of operation, and to try to prepare them to support NRAO when it approaches the funding agencies. Scientific need must be emphasized in the sales pitch. Capitalize on the VLA's impact. Notice, for example, that the Field committee report has a VLA image on its cover--and think what that means.

f. Cornwell suggested forming a group to study problems of image construction. Such a group fits in well with the hardware facility considered here. It would be foolish to do either without the other. Forefront studies like this are now limited by lack of manpower as well as a shortage of computer power. A couple of problems this group should study are mentioned in sections 6c and 8 of this note.

10. Closing Remarks

This report started out to consider computing within NRAO, but it is in reality a discussion of map-making, an activity that has a very heavy computational component. Map-making is the main computational load, and it defines the computational requirements. Easy, thorough, and rapid map making is an essential part of the operation and scientific use of the VLA, and it will be an equally essential part for VLBA and the millimeter array. Making maps of the size and detail that are routine practice at the VLA makes every user a practitioner of the art of large scale scientific computing. Computation is the principal ingredient of making maps. This is not only a consequence of the fact that synthesis data are born as the Fourier transform of a map--the very essence of reconstructing an image from a noise-corrupted source, with imperfect "optics" (side-lobes) creates a heavy computational load. Computation is so much an ingredient of map making that it is taken for granted and often left unstated by users of the VLA. The VLA, and other synthesis arrays, would be nothing without computation.

The VLA exceeds its design goals in practically every way except the length of the track. With the VLBA-VLA combination, even this design will be exceeded. Improved hardware, such as better receivers, accounts for part of this, but the major contribution--that responsible for the fact that observers routinely produce maps with 1000:1 dynamic range, while the original design specifications for the VLA called for a dynamic range of 100:1--comes from new reduction methods that users have learned since the VLA was designed (CLEAN and SELF-CAL), all of which are computationally intensive. These procedures have completely revolutionized the process of map making, but they have done so at the cost of increasing the computational load by factors ten or more per map produced. There is every reason to expect that any future improvements in scientific utilization of the VLA will come either from new computationally intensive methods or from hardware developments that will increase the computational load.

It is better to focus on the scientific use of the VLA, and of the consequences of that scientific use on computational requirements, than to focus on the computational requirements themselves, in order to keep the emphasis properly on making the VLA easy to use as a research tool doing science. Emphasis on the computational requirements themselves runs the risk of creating a design that is elegant from a computer operator's or programmer's point of view, but difficult to use as part of the process of using the VLA as an astronomical research tool.

* * * * * * * *

This screed contains a potpourri of thoughts relating to the general problem of image reconstruction within NRAO. It consists largely of my impressions of the situation at the VLA, partly garnered through conversations with staff and users. It is meant to be provocative. If it stimulates discussion it will have served a major part of its purpose. Much of it consists of things you all know--there is very little that does not appear in some VLA or NRAO report already written, except for the stuff you all know so well that you never bother to include it in documents or reports. The reason for including it is to marshal some of the arguments that should be included in reports intended for audiences outside NRAO--points you may omit because they are so obvious. They may not be obvious to people at AUI, NSF, in the astronomical community at large, the congress, or the general public. Feel free to crib from it without reference or credit, or to use it for whatever purpose you please. (I think I'd rather not know <u>all</u> the uses to which it might be put!)