

AIPS MEMO NO. 53

THE 1987 AIPS WORKSHOP

Alan H. Bridle

National Radio Astronomy Observatory
Edgemont Road, Charlottesville, Virginia

8 October 1987

This memorandum is intended to capture the main points that were brought out at the AIPS Workshop held in Green Bank on Tuesday, September 15 and Wednesday, September 16, 1987. It is gleaned from my notes of the meeting, so it may not be fully comprehensive. I hope that even so it may be a useful reminder of the main issues that were raised at this workshop, whose agenda is reprinted in Appendix A. Those who were unable to come may want to contribute to ongoing discussion of these topics via BANANAS.

There were 34 "official participants" (Appendix B); several Green Bank staff and visiting observers also came to some sessions. The outside user groups that were represented were: CSIRO Radiophysics (1); Fujitsu (2); Lab. di Radioastronomia CNR Bologna (1); Max-Planck-Institut fur Radioastronomie Bonn (1); E.S.O. (1); Jodrell Bank (2); D.R.A.O. (1), H.I.A. Ottawa (1); University of Toronto (3); University of New Mexico (1); Haystack Observatory (1); Naval Research Laboratory (2); NASA Goddard Space Flight Center (1); Space Telescope (1); COBE Project (3). There were twelve participants from NRAO -- ten from Charlottesville and two from the VLA.

The workshop comprised seven sessions of talks and discussion. An arrival reception and an "informal discussion" session in the Residence Hall lounge also provided good opportunities for the participants to get to know one another and to exchange ideas and concerns about AIPS. A few stalwarts continued discussions on the second evening until about 2 a.m. on the 16th !

SESSION 1

Richard White summarized his use of AIPS for reducing IRAS data, and outlined the proposed role for AIPS in analyzing data from the COBE (COsmic Background Explorer) Project. Both applications rely heavily on the ability of AIPS to display, and re-grid, wide-field data into different co-ordinate systems. The edge filtering and Fourier domain editing capabilities of AIPS are also important for these applications. Richard pointed out several areas of AIPS that will be critical to the

COBE project and for which further development may be needed. These include support of further co-ordinate systems, support for interchange of data with other archiving and image display packages using FITS disk format, and routines for point and diffuse source parameter extraction and classification. Richard also noted that software written for analysis of LANDSAT images has some useful tools (e.g. for deriving pixel spectra from data cubes) that should be incorporated into AIPS.

Bill Cotton described the status of continuum and polarization calibration software in the 15OCT87 release of AIPS. He described the use of multi-source u,v data files with a much-expanded Tables structure to avoid having to copy the data set whenever the data are recalibrated or edited. He noted that the new data file and table format are fully supported in FITS format so that all calibration information can be written to tape. The new system permits some error recovery if users keep intermediate versions of calibration Tables on disk. Bill briefly described the new interactive TV-based editing routines.

Points arising in discussion were the need for an improved terminal interface to handle multi-page inputs for complex tasks such as CALIB, the likelihood that workstation interfaces will provide good tools for doing this, and the need for the polarization calibration code to incorporate ionospheric Faraday rotation correction. Bill noted that AIPS should soon provide the ability to make Faraday rotation correction (though it does not at present), but that the programs for calculating them, and the ionospheric data on which they are based, are currently maintained by VLA staff. The programs will have to be ported into AIPS, and new arrangements made for distributing the ionospheric data to AIPS user sites.

Phil Diamond described his progress on spectral line calibration software for AIPS. The software is station-based and is written for a general interferometer (i.e., it is not NRAO-specific). Task BPASS can make amplitude bandpass correction spectra from autocorrelation data. It can also use the same least-squares algorithms as in self-calibration to determine antenna-based amplitude and phase correction spectra from observations of calibration sources. Use of this algorithm has already shown that allegedly "calibrated" VLA data require significant amplitude corrections across the band for some antennas. The bandpass corrections will be applied as multi-source data sets are split into multiple single-source data sets for imaging. Three options will be available: (1) averaging all bandpass data for each antenna to make constant corrections; (2) using the bandpass data nearest in time to each scan to make variable corrections; (3) interpolating bandpass spectra in time for some antennas, to deal with rapidly changing responses. The third option will be computer-intensive and should rarely be used. Phil also described plans for supporting VLBI needs such as aligning spectra in velocity after stations have used fixed frequency offsets for the observations, and for template autocorrelation methods for maser observations.

Gareth Hunt reviewed how the new VLA on-line system (expected to assume full control of the VLA by the end of 1987) will affect sites that

intend to use the new AIPS calibration package. The new VLA archive tapes will only be written at 6250 BPI, with 26624 byte records and an industry standard format. Derived quantities (such as true center frequencies and antenna locations), autocorrelation data, and delay and phase quantities of interest to astrometry, will all be written to the new tapes. He emphasized that the new on-line system will be able to record 512 spectral line channels, rather than 36, with the full array. 3.33-second averaging times will also be routinely available for continuum observing. Projects that use the full multichannel capabilities of the on-line system (either for spectral line observing or for wide-field continuum imaging) may therefore generate up to 18 times more data than before, and will require correspondingly more disk storage space. Gareth noted that although the VLA staff hope eventually to be able to convert archive tapes routinely to UVFITS format, present manpower and computer resource limitations at the VLA mean that only archive tape copies can be provided in the near future. (The original archive tapes remain the property of NRAO and the data on them can be made available to other users 18 months after the observations). The AIPS FILLR program that reads archive tapes will however be maintained by VLA staff, to ensure that any changes in the VLA on-line format are tracked promptly within AIPS.

SESSION 2

Jun Maekawa described the AIPS installation on the Fujitsu VP-50 supercomputer at Nobeyama Observatory. This is the most powerful CPU in which AIPS is now actively in use anywhere in the world. AIPS is run in batch mode with job control instructions generated on the M380S computer. The 15JAN87 version of AIPS has been running there since April 1987. Major modifications are needed to adapt AIPS to the FACOM operating system, which resembles IBM's MVS. Particular difficulties are: the 40 kb minimum file size; the need to manage task activation from AIPS within a batch job queue structure; and modifications required to run Imagen displays. Installation of a new AIPS version takes up to three months.

Yoshihiko Mizumoto described the tuning of AIPS for the VP-50. The Q-routines were vectorized using compiler directives that were introduced into the Convex C-1 source code by INCLUDE statements. This gained a factor of about 2 in CPU time and 3 in real time. Use of virtual I/O to the main memory for temporary files, especially when the DFT was used, typically gained a factor of 2 in real time. The Nobeyama Observatory has prepared a full report on the performance of AIPS in the VP-50 ("AIPS on the vector processor FACOM VP-50", by M. Ishiguro, A. Kurokawa, J. Maekawa and Y. Mizumoto) describing its CPU and real-time speeds, and the accuracy of resulting images relative to those computed from the same data in NRAO's Convex C-1's. There is an apparent loss of precision for large images; the reasons for this are unclear. They may include the presence of Q-routines from different epochs of AIPS in the Nobeyama code, and the effects of tape transfers of the data using only 16-bit FITS format. He emphasized that AIPS in vector-register machines should do more to leave behind the FPS120B history and that the large memories now available on some machines should be used for scratch file

storage to accelerate real time performance of I/O-bound tasks (such as UVSRT). In discussion, Don Wells pointed out that the benchmark images for the DDT test should ideally be obtained by computing to 64-bit accuracy in two different supercomputer architectures, finding agreement between them, and then writing the "master images" to the DDT tape to 32-bit precision.

Dave Shone described the AIPS installation of the 4-CE Alliant FX/8 at Jodrell Bank. This machine is connected via an Ethernet LAN to Apollo DN560 workstations and to several VAXes. He also described the experience of using both OLAF and AIPS for reductions of MERLIN and VLBI data. OLAF has been used more for editing, calibration and difference mapping, but less for image processing. Users now employ both OLAF and AIPS according to taste, passing data between them by direct file transfer under control of AIPS. OLAF has achieved large speed-ups in the Alliant (e.g. a factor of 8 for a 2000-iteration, 256 x 256 Hogbom CLEAN) by use of memory for in-place storage of images during CLEAN. Dave stressed that the entire factor of 8 is relevant only to OLAF and its own earlier performance -- but he believed that similar strategies might improve current AIPS CLEAN throughput in large-memory vector machines by a factor of about two. If many tasks try to use 38 Mbytes of virtual memory simultaneously, the Alliant performance degrades rapidly, however. Dynamic memory allocation will soon be attempted. He described the current strategy for use of the Alliant's multiple processors: tasks such as MX, ASCAL, VM and APCLN are run in concurrent-vector mode, while tasks such as AIPS and TKPL run on single processors. The system dynamically allocates the clustering of the computing elements. The system manager also varies the tuning to favor keeping the cluster together overnight (when the number of users is small) but to favor splitting it into single CE's in daytime (when there are more users).

Dave also reviewed strategies for using networked workstations as AIPS displays. Current practice at Jodrell, based on an implementation developed for OLAF, uses a server running in the workstation to listen for calls from the Alliant. Jodrell will soon beta-test an Alliant package that will instead implement remote procedure calls; the Alliant will not know where the TV's are, but will use a calling structure resembling a Y-routine that itself calls up remote TV's.

In discussion, Richard Simon noted that difference mapping can now be done in AIPS through procedures that combine APCLN, UVSUB and ASCAL, but that a more efficient interface is desirable. Eric Greisen observed that software for "virtual TV-by-wire" was developed in AIPS to provide remote TV's while NRAO was using a Cray X-MP at Digital Productions. Work on this was suspended in favor of other priorities when funding for the Digital Productions project ended, but will continue when NRAO acquires SUN workstations for use with its Convex C-1's.

Mark Calabretta described the current status of AIPS at CSIRO, where there are about 50 locally-coded AIPS tasks and about 25 AIPS tasks have been significantly modified. (This is the largest local coding effort in AIPS at any user site known to the NRAO -- the NRAO-supported

AIPS package contains about 200 tasks). Mark noted that CSIRO protects its local code at update time by modifying COMRPL to distinguish the CSIRO code and to use it to replace the NRAO code where both exist.

CSIRO has also implemented code to allocate image planes and track-ball resources dynamically between two Gould-DeAnza TV monitors. This code also allows conversion between graphics planes and image planes. They have also coded a WYSIWYG task to add captions to the TV display and to send hard copy of grey-scale, labels and captions to a Laserwriter printer. They also have facilities for storing, recalling and manipulating TV look-up tables. A MIXPL task will translate AIPS .PL extension files into device-dependent files specified by an adverb. This task can support Zeta, HP or Tektronix pen plotters (including color), Macintosh PC's as terminals, and various graphics terminals. Mark reported that a popular option is to use MacDraw and MacGraph software to embellish a plot and then route this to the laser printer for hard copy.

AIPS is run at CSIRO in two VAX750's with a common disk farm and Mark described tools developed for this. A PROCESSOR logical for each machine includes different ranges of POPS number for the two machines, and the start-up is modified to assign POPS numbers in the appropriate ranges. Scratch files and batch must be handled specially, as neither processor knows what the other is doing. Scratch file sequence numbers are linked to the processor number; cataloging and destruction of scratch files are modified so that processors operate only on files that they have created. As QMNGR will have difficulties with the AIPSB coming from the other machine, duplicate system files are used to make DA00 point to different directories for each processor, so that QMNGR uses a different BQ file for each.

Mark also described some system management tools that have proven particularly useful. As TIMDEST is not used, they have tools to archive files to tape by renaming and de-cataloging them, then writing them to tape with VMS BACKUP and deleting them. When files are restored from tape, they are recataloged and updated to the current AIPS format if necessary. An automatic procedure deletes extension files whose image file no longer exists. There is a verb to mount and dismount a disk from within AIPS; this makes it easier to use a demountable disk.

In discussion, it was pointed out that the Zeta plotting KONTR task in AIPS will not survive the forthcoming code overhaul and that its CSIRO substitute could be imported to NRAO to replace it. Mark stated that some of the CSIRO code had already been sent to NRAO for redistribution to any AIPS sites that are interested in it, and that the rest would follow. Don Wells noted that there are now BAKTP and BAKLD tasks in AIPS for producing machine specific tape backups under both VMS and UNIX.

SESSION 3

Colin Lonsdale described the installation of AIPS in a Ridge

computer at Haystack under the Ridge Operating System (ROS), a dialect of UNIX. The principal difficulties he experienced were: 1) the compiler was sensitive to data alignment errors that were not caught by VAX or Convex compilers; 2) the compiler often failed to compile complex arithmetic expressions correctly; 3) the performance of some device drivers was unsatisfactory; 4) the system's multi-tasking abilities were poor; 5) it uses a proprietary bus; 6) problems arose where AIPS uses a feature that is not part of "official" UNIX. On the bright side, the system has a bit-mapped display that does TEK graphics especially fast. Colin has written Lexidata Y-routines to overcome some driver limitations. His future plans call for use of the Ridge as a file/print/tape server to a network of SUN 3/260 workstations, and he encouraged the AIPS group to support such networked systems.

Richard Simon echoed Colin's negative evaluation of Ridge systems based on his experience at the Naval Research Laboratory, then described enhancements to AIPS that are regularly used in the VAX environment there. They run AIPS batch in their system batch queues using a SUBMIT verb to replace GO; the simplicity of this encourages use of batch even by inexperienced users. They use a SHELL verb to exit from AIPS to the operating system, e.g. to access DEC editors. Gripes are automatically E-mailed to NRAO. AIPS is run from private log-ins to enforce disk quotas.

Richard argued for an UPDATE task to be written to work around the format changes that will come with code overhaul. He stated that most archiving at N.R.L. had been done using VMS "BACKUP" and that obsoleting that archive would be a "disaster". This prompted considerable discussion, which made it clear that no other AIPS managers present shared Richard's level of concern. Eric Greisen stressed the exceptionally large amount of work that would be needed to produce an UPDATE task for this transition. He also emphasized that old system-dependent archives were already likely to be invalid because UPDATE tasks could only be guaranteed to work incrementally (i.e. from one format change to the next). A strategy that emerged from the discussion was: 1) site managers should encourage their users to archive in FITS format, and to convert existing archives to FITS format before the code overhaul is complete (by the 15APR88 release at the earliest, and it may well be later); 2) sites that have large machine-specific archives may need to retain an old executable file for FITTP to support writing FITS tapes from data in obsolete formats.

Richard also supported priority for TV-by-network, for a full graphics package such as PGPLOT or MONGO, for ASCII output of AIPS Tables (so that these could be read into PC spreadsheet and data base programs, for example), for a more modern terminal interface, and for standards for system-dependent AIPS verbs. Eric expressed reservations about use of external PC software to modify tables if they are to be read back into AIPS -- but the consensus was that this was desirable and that the user should simply take responsibility for reading them back in exactly the format in which they had been delivered out of AIPS. Don Wells noted that DIF format should also be considered as a standard export format, although some otherwise excellent PC software (e.g. REFLEX) does not use its

capabilities properly.

Carla Fanti described experience with AIPS in Bologna, where a VAX11/750 plus AP is used with an Aydin TV to run AIPS in a multi-user environment. The most useful enhancements to AIPS from her perspective would be: 1) ability to handle observations taken in linear polarization as well as circular; 2) ability to begin self-calibration with multiple Gaussian source models; 3) more flexible tools for CLEAN component editing, especially for defining windows in which to exclude negative components; 4) ability to plot model visibilities for more than 1000 components. (The importance of CLEAN component editing followed by CCMRG to compress the component representation was emphasized in discussion). Carla noted that the AIPS version running in Bologna often lags the current release significantly (e.g. by a year) and that it would be valuable to access CHANGE.DOC remotely to become more aware of changes that had occurred in Charlottesville. It was noted that this should now be particularly easy for sites with SPAN connections. Carla pointed out that about 5% to 10% of AIPS use at Bologna was for CCD reductions, and that AIPS was often used for such optical image processing despite the availability of MIDAS.

Gustaaf van Moorsel described the use of AIPS at E.S.O., where the competition with MIDAS is more severe ! He emphasized the difficulties of supporting AIPS in a "hostile" environment in which most users do not wish to give it priority. The assumptions made by AIPS about resource availability often conflict with rules set by system managers -- e.g., at E.S.O., disk space must be registered for in advance. He commented that installation with a DeAnza display was now straightforward; joy-stick control was favored at E.S.O. and he had obtained useful routines for this from Dwingeloo. Gustaaf emphasized the need to run AIPS from user directories. In discussion, many of the other AIPS managers present argued that this should be the norm, rather than the exception. This raised the issue of system security in the presence of password-less accounts such as AIPS. Gareth Hunt and Bob Burns noted that NRAO must soon confront this issue, so that AIPS may have to adopt a standard for running from user accounts. Eric stated that this had been done for years in Charlottesville with few problems, and asked that site managers who had experienced difficulties with running from user accounts, or possible security loopholes associated with AIPS, write to him giving details as soon as possible. It was also suggested that AIPS site managers who feel that they have particularly good solutions to this problem disseminate them via BANANAS. Richard Prestage (Starlink) and Chris Flatters (Bonn) stated that they would do so. Geoff Croes noted that a reason for "hostility" to AIPS at other sites is that the buffered I/O loads VAXes unacceptably (as perceived by non-AIPS users).

Chris Flatters described the use of AIPS in a Convex XP-2 (two C-1's) at MPIfR Bonn. Their TV's (Raster Technologies, IVAS) have yet to be interfaced. They need to have the Cal Tech VLBI package in the Convex for use with AIPS (it was stated that John Benson now has parts of this package running on the Charlottesville Convex). Their main wishes are: 1) interactive editing by baseline (this was already implemented, and will

appear in the 15OCT87 AIPS); 2) ability to handle different weighting functions to put error bars on VBPLT output; 3) ability to plot and edit calibration tables (an ASCII interface to external software, as discussed above, would be sufficient).

Jeff Bowser described the proposed use of AIPS for COBE data reduction, emphasizing the need for mosaicing.

Susan Neff stated the difficulties of managing an AIPS installation in a VAX cluster with shared disk and tape resources, as well as many "hostile" non-AIPS users. She emphasized the need for tools to support AIPS use from individual accounts and in network/cluster environments.

SESSION 4

Laura Carriere described how AIPS has been migrated from a "hostile" VAX8600+AP environment into a dedicated SUN workstation at the University of Toronto. The Toronto system has a SUN3/180 with 8 Mbytes of memory, the Weitek Floating Point Accelerator, a 19-inch color monitor, 1/4-inch cartridge and 1/2-inch tape drives, 1.1 Gbytes of disk for AIPS data storage, and an Imagen printer. The total cost was about \$100,000. It is linked by Ethernet to several VAXes, a Cray X-MP, and external networks. It runs AIPS, PGPLOT, TeX, NeWS and Focusview (a package developed at U. of T.'s Scarborough College to display images on the SUN monitor under IRAF). It also runs Cray satellite station software to use it as a window into the Cray X-MP. DDT tests have shown that the CPU times in the SUN3+FPA are about one-third longer than in a naked VAX8600, and about three times longer than in an empty VAX780+AP. The system is dedicated to AIPS use, however, and this clearly benefitted users whose jobs are not too large for it. They plan to overflow larger jobs to the Cray X-MP.

Brian Glendenning described AIPS Y-routines for the SUN color monitor that are being developed at the University of Toronto. AIPS does not handle the windowing interactions, but merely passes data to the SUN main memory through a simple op-code protocol that cannot get out of sequence. The image is retained in this memory, and "Sunview" handles the windowing interactions. Zoom and look-up tables are applied as the data are written to the screen memory. Only 6 bits are used for grey-scale display, the other two being reserved for graphics and to allow text in other windows to show through. Two cursors are implemented, so that one of them can be moved out of the AIPS window to control other functions without changing the AIPS transfer function, etc. The track-ball buttons are implemented as icons; a resizing "button" can be used to reduce the display area to the minimum needed to display the non-zero pixels. Eric Greisen noted that the design of Brian's routine is similar to that in the (unreleased) AIPS TV-by-wire routines. Brian also mentioned that Patricia Monger (U. of T.) has also written an X-windows program to display images on the SUN.

Brian also observed that on SUN3's the FPA performance is limited by bus and memory speed rather than by the Weitek chip, and that assembly optimization is not expected to improve performance much.

Takako Amano showed illustrations of SUN color monitor output produced by the task SUNPL that she has coded at the Herzberg Institute of Astrophysics in Ottawa for use on a SUN3/180 linked to a Gould 32/67 machine. SUNPL is not a Y-routine but will take the output of any AIPS .PL file, including grey scale displays, and write it in a window on the 256-color, 1152 by 900 pixel SUN monitor. Pull-down menus and keyboard function keys are then used to adjust zoom, scroll, color tables and transfer functions. A variety of high-quality displays was illustrated with 35mm slides. Loading a 256 by 256 image takes about 23 seconds; larger images take proportionally longer.

An open discussion of workstations and graphics issues followed. There was clear interest in workstations as stand-alone AIPS systems because they can be afforded by individual investigators. It was noted that I/O limitations become apparent when there is more than 10 MIPS capability in the CPU, and that a new generation of faster machines will use disk striping to overcome this. Three stages of evolution can be envisaged. The first, running AIPS in stand-alone workstations such as the SUN, is already with us and is clearly of interest for jobs of the right class. The second, using the workstation as a TV device for a more powerful computer, is almost ready (witness the Toronto SUN Y-routines). The third will be to use the workstations as controllers and to make data files equally available to them and to a more powerful partner, e.g. a minisupercomputer or a Cray. This is being actively considered but implementation is still in the future. Bob Hanisch questioned whether 8-bit displays will be considered adequate. Discussion centered on the anticipated 16- or 32-bit displays. A target for the near future will be for groups that have imaging workstations to attempt to emulate the traditional AIPS TV functions to evaluate how well the workstations can perform. Eric stressed that the IIS is still the most competent display, and others should be judged by how many of its functions they can emulate (DeAnza's and IVAS do not support TVHUINT, for example). There is a need to be able to turn several planes on at once at the same pixel. It was not yet clear whether SUN workstations would be able to blink the entire display as fast as desired, though it should be close. The consensus was that it would be good for several groups to work simultaneously on workstation display implementations, particularly emphasizing emulation of the more sophisticated displays that are used on the IIS, and to communicate their successes or failures to the community via BANANAS. It was also noted that the Center for Astrophysics is developing a SUN3 display implementation for IRAF.

SESSION 5

Eric Greisen gave details of the planned AIPS code overhaul project whereby the coding standard will become close to ANSI FORTRAN-77. He stressed that this overhaul was being proposed for much stronger

reasons than mere aesthetics. He saw the most compelling reason to be the speed-up of integer computations, especially in the Convex, and gave examples in which Neil Killeen had found improvements of up to a factor of 10. Although the reasons for this were not completely clear, it appears that the degree of optimization that can be achieved by several compilers is increased by eliminating I*2 variables. The AIPS group had also found that many compiler bugs are related to I*2. Eric stressed that all opening and closing of files in AIPS would remain in Z-routines, and that Hollerith would still be needed for data structures. The overhaul will be used as an opportunity to update numerous features of the file formats and headers that have become unduly restrictive.

The code conversion will begin automatically with a list-driven program that will be provided to user sites with later releases. Much of the code will have to be hand-corrected, however, and the conversion and its debugging is expected to lead to cancellation of one AIPS release. All file formats except those of MA and UV files will change. All data sets will have to be backed up to FITS to work around the conversion, as an UPDATE program is judged to be about 6 months' work and will not be attempted with present manpower constraints. The earliest possible release containing the new code and formats will be 15APR88. The start of the project has been linked to reaching several targets in the calibration and editing project, and these have now been met. The remaining target to be met before the code overhaul begins is to run the DDT test successfully in the Cray X-MP at the Pittsburgh Supercomputer Center.

In discussion, Richard Simon asked whether this would be "the last format change". He encouraged the AIPS group to anticipate future needs as much as possible to maximize the likelihood that it might indeed be the last one. He reiterated that sites would need to keep an old AIPS shell and FITTP to compensate for the lack of an UPDATE program. Alan Bridle asked whether other site managers saw major difficulties in not having an UPDATE program to handle this transition, given that at least a year's notice of the change had been given in the AIPS Letter and that they could work around it by keeping an old FITTP alive. Gareth Hunt said that some VLA staff might experience difficulties, but that it was not obvious that the problem was severe enough to justify an UPDATE task. No other site managers expressed concern. Eric noted that Convex have decided that their initial adoption of the VAX floating point format was not optimal, and will support IEEE floating point. Version 6.1 of their operating system will have software support for an existing hardware switch to enable IEEE floating point. The code overhaul will anticipate this.

Kerry Hilldrup described how he is rationalizing the Z-routine structure to isolate the truly system-specific functions into lower leaves of the Z-routine tree. He noted that he has sought to preserve the names and calling sequences of routines that are called by the machine-independent code, and has isolated device specific code into separate routines with names that suggest which devices they involve. He has introduced new system parameters or constants into ZDCHIN to describe (for example) floating point formats, file control block sizes, offsets and contents, scratch volumes, short loop and vector register machine

indicators, and 2's complement versus 1's complement machine indicators. ZDCHI2 will extract and calculate system parameters and will make them available to routines written in C.

Geoff Croes urged the AIPS group to consider the advantages of pre-processors. Referring to Eric's statement that aesthetics was not the main reason for embarking on the code overhaul, Geoff noted that aesthetics is nevertheless important, as badly formatted code is hard to test and to maintain. Among the advantages of a pre-processor for AIPS code, he saw 1) a pre-processor can ensure portability, by producing FORTRAN-66 or FORTRAN-77, 2) a pre-processor can be made to generate target-specific code, 3) code produced by a pre-processor can be easier to debug than much hand-written code now in AIPS, because the output code is highly structured, 4) a pre-processor can give programmers early access to features that have not yet appeared in the FORTRAN standard, and this may encourage users to write programs, 5) a pre-processor can be used to enforce coding standards. Geoff related how these advantages had been important in increasing productivity, and in harnessing student labor, at DRAO. The disadvantages of pre-processors were some decrease in compilation speed, and the introduction of incrementally new language features that cause resistance from programmers with more entrenched habits.

Geoff also described how an automatic document processor had been used to generate programmer-oriented documentation at DRAO. He urged the combination of declaration and documentation, of providing subroutine help libraries, of keeping a maintenance history, and of writing user documentation before writing the code. Eric noted in discussion that AIPS has a document processor but that it had fallen into disuse. Alan Bridle stated that a planned expansion of the AIPS group has asked for more manpower to support documentation, and that Geoff's suggestions and the existing document processor should be reviewed if this occurs.

SESSION 6

Don Wells reviewed the impact of new technological developments on AIPS. The spread of network connections to most of the powerful machines that are now in use for AIPS had encouraged the group to turn on the BANANAS E-mail "exploder" for use as an AIPS users' bulletin board. Several other ways were envisaged to exploit networking to improve user site support.

Don has turned on an account at NRAO called AIPSSERV that uses E-mail as a file transport mechanism. It reads mail, executes commands in the mail, and returns files to the sender in a form that will survive BitNet transfer. He suggests sending "HELP" as a one-line mail message to AIPSSERV@NRAO, to get started.

The AIPS group has been using "midnight jobs" to update TST and NEW in VAX1, VAX3, nrao1 and cholla with daily changes to the AIPS code made in CVAX. This has become routine, and will be used to update the

code in the Cray X-MP at Pittsburgh and possibly also the Cray at the University of Toronto. It would be feasible in principle to use this technique to update code to any AIPS site that has E-mail connections to NRAO, though it would be neither desirable nor practical to do this on a daily basis to other sites (weekly updates would be more practical). Any sites wishing such service would also be obliged to maintain both TST and NEW, and to synchronize their versions with NRAO's at release time.

Don also suggested that network connections might be used to send Gripes-by-mail to Charlottesville (and to return receipt messages automatically), or to implement remote code check-out and put-back for distributed AIPS development. The latter facility will be needed for NRAO's own code development at supercomputer centers, but later discussion suggested that it might not be appropriate for more general use. Finally, Don stated that an E-mail exploder for FITS users, "INFO-FITS", would soon be available. This should be used to discuss such matters as floating-point standards, extensions, blocking and Tables conventions, standards for non-linear co-ordinates, and other FITS issues that the AIPS community should track closely.

Don observed that performance of minisupercomputers was doubling about every 18 months and that concurrency and vector registers were becoming ubiquitous. He noted that developments in VLIW ("very long instruction word") architectures and fiber-optic LANs should soon lead to interesting systems. The rapid development of workstations was leading to the "TV in a window" as a software abstraction that would permit great flexibility -- workstation users would be able to ask "what kind of a TV do I want to have today?". Don noted that NeWS supports pixel data and half-tone rendition and could have some display ability even on a single bit plane. Its Postscript-based technology would give it a strong advantage, especially as there is now a VLSI chip that executes Postscript. He expects Postscript to dominate the laser printer market. There would also soon be announcements of so-called "graphics supercomputers" that will integrate concurrency and vector processing in CPUs of Convex C-1 power with large, fast image displays and striped disks with entry prices just under \$100,000. Cheap "knob boxes" that could be used with workstations as control panels could also be of interest. They could be a powerful new approach to applications previously served only by manpower-intensive development of special hardware (such as NRAO's image storage units). Finally, new devices were becoming available for stereo viewing of three-dimensional data.

Don also reviewed how the already much-discussed area of TV-by-wire and remote AIPS displays would be affected by new networking techniques. Transparent file sharing between UNIX machines was already supported by Network File System (NFS). It would be important to use the IEEE floating point system in all machines, and full binary compatibility would be desirable if possible. Don predicted that byte-swapped architectures would soon be regarded as unacceptable. The goal for AIPS development should be that local AIPS operates as a server, that remote AIPSeS connect to it across a net and are serviced through network Z-routines. All AIPSeS would then be able to operate remotely. Virtual

TV's should be implemented by passing Y-routine arguments across networks in machine-independent form; this might be a good way to support multiple TV's on each AIPS.

Bob Burns added that cassette tape technology was also rapidly evolving in interesting directions. 10 Gbyte VHS cassette tapes should soon be on the market, and more tape cartridges were sold last year than conventional tapes. The technology was cheap but the disadvantage at present appeared to be search time. The 10-Gbyte cassettes would take about 5 minutes to fast forward from end to end, and about 5 or 6 hours to read. This would improve by about a factor of three with the new DAT tapes. Bob observed that optical disk technology was not evolving so rapidly and that NCAR had abandoned optical disks as their mass storage medium and was moving to IBM cassettes. In discussion, Richard Simon noted that there was now a 150 Mbyte cassette drive available for a PC for only about \$900.

Alan Bridle asked for discussion of how NRAO's move away from AP's would affect sites that now rely on them, especially with regard to new algorithm development in AIPS. Bill Cotton stated that no new AP micro-code would appear, but that new Q-routine based tasks could be expected. Sites that now have AP's would eventually find that maintenance costs will encourage them to consider buying vector register machines. Bob Burns noted that FPS sales are falling and the company is "retrenching".

SESSION 7

Alan Bridle reviewed NRAO's new "Array Telescope Computing Plan", which had just been sent to the National Science Foundation. It addresses fiscal years 1989, 1990 and 1991, and replaces the earlier proposal for a dedicated supercomputer. It proposes a four-pronged attack on the combined computing load of the VLA and VLBA.

First, an expanded computing facility at NRAO -- comprising four loosely-coupled second-generation minisupercomputers, each with about 3/4 of the power of a single processor Cray X-MP. Manufacturers have told NRAO that such machines will be available on the time scale addressed by the proposal. Each would have 1 Gbyte of memory, 20 Gbytes of disk, and would be networked to a set of workstations for use as image displays. One workstation on each machine would be a "graphics supercomputer" of the class described by Don Wells. This facility would provide about 14 times NRAO's present computing capacity and would be able to process between 40 and 50% of all proposals using the VLA and VLBA.

Second, increased explicit NRAO support for computing at user institutes. About half of the present computing power devoted to VLA data processing is outside NRAO (according to the last AIPS Site Survey) and the observatory seeks to preserve this, both to distribute the load and to support users' computing where they can do it with the most care and thought -- at home. The plan calls for increased manpower to support

exporting AIPS, for more assistance to user sites, for co-operation with vendors to benchmark new products, and for increased support for Users' Group activities (especially networking).

Third, access to supercomputers at national and regional centers to support the 10% of all VLA and VLBA projects that require the most powerful CPU's available to become tractable. The NRAO would maintain AIPS in one supercomputer operating system (currently COS), and support it under UNIX emulations in others. It would use the NSF access program to develop and run code for the most computer-intensive tasks at national supercomputer centers. AIPS would play a role in this, but stand-alone packages would also be used, at least initially, to optimize throughput. The plan called for one programmer to be dedicated to supercomputer applications.

Fourth, increased staffing for software and algorithm development. Areas given high priority include wide-field mapping (especially for low frequencies); improved interfaces between self-calibration and deconvolution; efficient cube-building and manipulation; astrometry and geodesy; improved visualization and parameter extraction (especially for cubes); efficient use of large memory resources; more systematic code validation and documentation; optimization of code for widely-available hardware; and improved code management. The plan calls for about seven positions to be added to the AIPS staff, in the areas of: graphics interfaces and workstations; scientific applications; code vectorization and optimization; supercomputer support; support of AIPS at non-NRAO sites; code validation; and documentation. Other staff would be added more generally, including positions to support algorithm research, systems analysis, networking and communications, and management and supervision.

There was a favorable reaction to the scope of this plan and to its balance between computing at NRAO itself and at other institutes. It was noted that the plan assumes that the computing capacity at supercomputer centers and at user institutes will continue to increase at about the same rate as the computing capacity at NRAO. Some users questioned the need to run stand-alone code in the supercomputers. Alan stated that the tradeoffs between optimization and portability in the supercomputer code would need careful study. Some big applications will need highly optimized code to get good throughput -- users with these applications may be willing to sacrifice some convenience and interactivity for computing speed. Alan stated that commonly used codes should eventually become part of AIPS, and that many AIPS utilities will be needed by users at supercomputer centers, but that the new plan does not prejudge the eventual balance between AIPS and non-AIPS code. Bob Hanisch commented that the plan was "a bargain" relative to NASA software costs, and that NASA funding should be sought if the NSF cannot support all of it.

The rest of the session was devoted to round-table discussion of priorities for future AIPS development. Kerry Hilldrup asked for discussion of the AIPS update cycle -- should the interval between

releases be increased ? The site managers present had few strong feelings about this. Eric argued for retaining the present cycle 1) because the update forces a recompile-and-relink test for all the code every quarter, 2) because the overhead of producing an AIPSletter is roughly proportional to the number of changes made, and thus to the elapsed time, 3) because it takes two releases to propagate "new" code to the outside world. He noted that a longer update cycle would delay innovation, would make it harder for the AIPS programmers to support the code that is being run at user sites, and would make it harder for users to write code that could be integrated back into AIPS. The consensus was to retain the present quarterly update cycle.

Bob Hanisch outlined the need for an AIPS Site Manager's Manual that was intermediate between the installation guide and GOING AIPS. He particularly asked for more organized discussion of the directory structure, management tools and debugging tools, and for advice on how to get an installation going again if it hangs up on a problem. Eric pointed out that the installation guide represents the group's best shot at describing how to instal AIPS, and includes its response to the difficulties that are encountered when installations are done at NRAO. He stressed that advice about difficulties encountered outside the observatory, and about strategies for recovering from such difficulties, could only be compiled by pooling the experiences of AIPS site managers themselves (much as the Cookbook had been compiled by a group of users rather than by the programmers). It was agreed that site managers who have accumulated installation tips and tricks to deal with recurrent problems should send these to Don Wells for compilation into a draft manual. The draft could then be reviewed and commented on by site managers, using AIPSSERV and BANANAS. Laura Carriere volunteered to assist with editing the results into the hoped-for manual.

The need for an "Advanced Users' Guide", such as a printed manual of selected EXPLAIN files, was also emphasized. Richard Simon advocated a tool for a keyword or context search through the existing HELP catalog. Bob Hanisch suggested that the SDAS "make help" facility might be used as a way to encourage programmers to document their code in a more disciplined fashion.

In the concluding discussion, every participant was asked to state their "one wish" for priority development in AIPS. Many themes recurred, and some participants made several suggestions. The following topics were mentioned:

- o AIPS Manager's training and trouble-shooting manual (8 requests)
- o Improved graphics capabilities (5 generic suggestions, but Eric asked that users now make explicit requests, perhaps via BANANAS, so that the AIPS community can set priorities for particular capabilities or packages)
- o Fuller explanation of what the installation is doing, on the lines begun in the UNIX installation guide (2 requests, especially for improvements to the VMS version).

- o More readable code and system documentation (e.g. help library for subroutines) (2 requests)
- o Eliminate I*2 (2 requests)
- o A COPY command between Convexes and VAXes
- o A regular AIPS port that tests for data alignment problems
- o Transparent support for "Cray-by-wire" (!!)
- o Simpler support for difference mapping, OLAF-style
- o More convenient editing of CLEAN components
- o Multiple Gaussian component input to self-calibration
- o "Midnight job to my site"
- o Installation option so that CTRL-Z does not stop spawned tasks
- o Up-to-date "What's out there" user site directory on AIPSSERV
- o Ability to process 8000-channel spectral line data
- o Image file creation date in AIPS CA file
- o "Spherical cube" co-ordinate system
- o Ability to use edited UV image files as UV data files
- o Clean up internal use of I/O buffers
- o System-dependent allocation of buffer files
- o Support a supercomputer system somewhere warm (an explicit suggestion was made)

Finally, the format and venue of the AIPS workshops was discussed briefly. The two-year cycle was endorsed. Those who expressed a preference between Green Bank and Charlottesville favored Green Bank as it made for a more interactive meeting. It was suggested that future workshops include a scientific session "so that computing demands can be related to scientific output". It was pointed out that they should be scheduled to avoid VLBI network runs by a bigger margin than this one.

The meeting concluded with a tour of the Green Bank telescopes led by Mark Clark, Harry Payne, Bob Vance and Ron Maddalena. Some discussion of the uses of AIPS for single-dish data reduction was stimulated by this.

ACKNOWLEDGEMENTS

Holding the workshop at Green Bank contributed greatly to its success, and I thank the Green Bank staff, especially Becky Warner and Richard Fleming, for their work on local arrangements. I also thank Bob Havlen and George Seielstad for supporting the workshop with NRAO and AUI resources, and George Seielstad for inviting us back in 1989.

APPENDIX A - AGENDA

THIRD AIPS WORKSHOP
National Radio Astronomy Observatory
Green Bank, West Virginia

14, 15, 16 Sept 1987

Monday September 14

20:00 Welcoming Reception (Residence Hall Lounge)

Tuesday September 15

07:30--09:00 Breakfast (Residence Hall Cafeteria)

Welcoming remarks (Auditorium)

09:00 G. Seielstad Welcome to Green Bank
09:05 A. Bridle Miscellaneous arrangements

SESSION 1. New applications (Chair: Don Wells)

09:10 R. White New IR Algorithms for AIPS
09:30 B. Cotton The new AIPS Calibration Package
10:10 P. Diamond Spectral Line Calibration in AIPS
10:20 D. King/G. Hunt The future of FILLR

10:30 Break (outside Auditorium)

SESSION 2. Experience at user sites (Chair: Ed Fomalont)

11:00 Y. Mizumoto AIPS on the Nobeyama Fujitsu VP-50
11:30 D. Shone AIPS and OLAF on the Jodrell Bank Alliant
12:00 M. Calabretta AIPS at CSIRO Radiophysics

12:30 Lunch (Cafeteria)

SESSION 3. More from user sites (Chair: Alan Bridle)

13:30 Review of AIPS experiences at other sites, including:

C. Lonsdale	Haystack, (Ridge)
R. Simon	N.R.L.
C. Fanti	Bologna
G. van Moorsel	E.S.O.
C. Flatters	MPIfR, Bonn (Convex C-1)
J. Bowser	STS Corp.

15:00 Break (outside Auditorium)

SESSION 4. Workstations and graphics (Chair: Eric Greisen)

15:30	L. Carriere	AIPS in the SUN at University of Toronto
16:00	B. Glendenning	AIPS graphics in a windowing environment with competing processes
16:15	T. Amano	SUNPL - a SUN graphics interface
16:30	Everyone	Discussion of workstations, TVs and graphics issues

Graphics standards
Better visualization of astronomical data

18:00 Dinner (Cafeteria)

19:30--whenever Informal discussions (Residence Hall Lounge)

Wednesday September 16

07:30--09:00 Breakfast (Cafeteria)

SESSION 5. The future - code and portability (Chair: Bill Cotton)

09:00 E. Greisen The AIPS Code Overhaul Project
09:30 K. Hilldrup New strategies for Z-routines
10:00 G. Croes AIPS Portability and Pre-Processors

10:30 Break (outside Auditorium)

SESSION 6. The future - hardware (Chair: Bob Burns)

11:00 D. Wells Advanced Technology Developments and AIPS
11:45 Everyone Discussion of future hardware paths

 Supercomputers
 AIPS in distributed systems (networks)
 Large memories

12:30 Lunch (Cafeteria)

SESSION 7. The future - priorities (Chair: Ed Fomalont)

13:30 A. Bridle AIPS in NRAO's long-range plans
13:50 Everyone Discussion of future priorities in AIPS

 Priorities for new coding effort
 Optimization versus Portability ?
 Role of an AIPS Users' Group
 Interfaces to other packages

15:00 Break (outside Auditorium)

 Telescope Tour - Mark Clark will meet
 those wishing to tour the telescopes,
 outside the Auditorium.

APPENDIX B - LIST OF PARTICIPANTS

Amano, Takako	Herzberg Inst. of Astrophysics, Ottawa	Canada
Bowser, Jeff	ST Systems Corp. (COBE Project)	USA
Bridle, Alan	NRAO, Charlottesville	USA
Burns, Bob	NRAO, Charlottesville	USA
Calabretta, Mark	CSIRO Radiophysics	Australia
Carriere, Laura	University of Toronto	Canada
Cotton, Bill	NRAO, Charlottesville	USA
Croes, Geoff	Herzberg Inst. of Astrophys., Penticton	Canada
Diamond, Phil	NRAO, Charlottesville	USA
Duric, Neb	University of New Mexico	USA
Fanti, Carla	Lab. di Radioastronomia CNR, Bologna	Italy
Flatters, Chris	M. Planck Inst. fur Radioastronomie	W.Germany
Fomalont, Ed	NRAO, Charlottesville	USA
Gaume, Ralph	Naval Research Laboratory	USA
Glendenning, Brian	University of Toronto	Canada
Greisen, Eric	NRAO, Charlottesville	USA
Hanisch, Bob	Space Telescope Sci. Inst.	USA
Hilldrup, Kerry	NRAO, Charlottesville	USA
Hunt, Gareth	NRAO, Socorro	USA
King, Dale	NRAO, Socorro	USA
Kronberg, Phil	University of Toronto	Canada
Lonsdale, Colin	Haystack Observatory	USA
Maekawa, Jun	Fujitsu/Nobeyama Observatory	Japan
Mizumoto, Yoshihiko	Fujitsu/Nobeyama Observatory	Japan
Napier, Leroy	NRAO, Charlottesville	USA
Neff, Susan	NASA Goddard Space Flight Center	USA
Prestage, Richard	NRAL, Jodrell Bank	UK
Shone, Dave	NRAL, Jodrell Bank	UK
Simon, Richard	Naval Research Laboratory	USA
Sinha, Ramesh	ST Systems Corp.	USA
van Moorsel, Gustaaf	European Southern Observatory	W.Germany
Wells, Don	NRAO, Charlottesville	USA
White, Richard	Greenbelt, MD (COBE Project)	USA
Wiener, Nancy	NRAO, Charlottesville	USA