

A I P S L E T T E R

Volume XXVII, Number 1: June 30, 2007

A newsletter for users of the NRAO
Astronomical Image Processing System

Written by a cast of *AIPS*

Edited by

Eric W. Greisen

National Radio Astronomy Observatory

P.O. Box 0, Socorro, NM 87801-0387

505-835-7236; Fax 505-835-7027

daip@nrao.edu

<http://www.aoc.nrao.edu/aips>

General developments in *AIPS*

FILLM

31DEC06 contains a revision of FILLM which is essential to support the new data form to be produced by the VLA beginning in July 2007. VLA users will have to upgrade their copy of *AIPS* to 31DEC07 or 31DEC06 by that time. See below for more information about FILLM, including important patches to the 31DEC06 version of FILLM. Changes which are not in the 31DEC06 version may also be important to current observers.

Current and future releases

We have formal *AIPS* releases on an annual basis. While all architectures can do a full installation from the source files, Linux, Solaris, and MacIntosh OS/X (PPC and Intel) systems may install binary versions of recent releases. The next release is called 31DEC07 and remains under active development. You may fetch and install a copy of this version at any time using *anonymous* ftp for source-only copies and *rsync* for binary copies. This *AIPS Letter* is intended to advise you of improvements to date in this new release. Having fetched 31DEC07, you may update your installation whenever you want by running the so-called "Midnight Job" (MNJ) which uses *cvs* and transaction files to copy and compile the code selectively based on the code changes and compilations we have done. The MNJ will also update sites that have done a binary installation using *cvs* and *rsync*. There is a guide to the install script and an *AIPS* Manager FAQ page on the *AIPS* web site.

The MNJ serves up *AIPS* incrementally using the Unix tool *cvs* running with anonymous ftp. The binary MNJ also uses the tool *rsync* as does the binary installation. Linux sites will almost certainly have *cvs* installed; other sites may have installed it along with other GNU tools. Secondary MNJs will still be possible using *ssh* or *rcp* or NFS as with previous releases. We have found that *cvs* works very well, although it has one quirk. If a site modifies a file locally, but in an *AIPS*-standard directory, *cvs* will detect the modification and attempt to reconcile the local version with the NRAO-supplied version. This usually produces a file that will not compile or run as intended.

AIPS is now copyright © 1995 through 2007 by Associated Universities, Inc., NRAO's parent corporation, but may be made freely available under the terms of the Free Software Foundation's General Public License (GPL). This means that User Agreements are no longer required, that *AIPS* may be obtained via anonymous ftp without contacting NRAO, and that the software may be redistributed (and/or modified), under certain conditions. The full text of the GPL can be found in the 15JUL95 *AIPS Letter*, in each copy of *AIPS* releases, and on the web at <http://www.aoc.nrao.edu/aips/COPYING>.

NRAO is 50 years old

The NRAO is celebrating its 50th year of existence. All sites, including the high site in Chile, were connected by television conferencing for brief remarks and a ceremonial cutting of cakes at each site. A symposium was held June 17–21 in Charlottesville, Virginia to celebrate the event entitled *Frontiers of Astrophysics: A Celebration of NRAO's 50th Anniversary: Astrophysics Across the Electromagnetic Spectrum*. To obtain more information, see <http://www.nrao.edu/50/>. It is hoped that the PowerPoint talks will appear on the web and that the whole symposium, which was said to be very good, will appear in a book.

Among the many letters received by the NRAO congratulating us on this anniversary was one which we thought would be of interest to *AIPS* users. It reads:

“Fifty years is a long time for the Italian radio astronomical community, which started 43 years ago with the completion of the Northern Cross in Bologna. It is probably fair to say that early interaction between Italians and US radio astronomers was at most sporadic at that time, with very little collaboration.

The construction of the VLA changed all this, and from the 1970s on, contacts have been intense. A whole generation of Italian radio astronomers could make use of advanced instrumentation and became strong admirers of the “American way” of doing science. Curiously enough perhaps it was not the hardware the most influential of all but a seemingly minor software development, which became known in the whole world as AIPS. The possibility to bring the software back home and have it run on extremely powerful computers (VAX) really changed the way in which scientific research was done. NRAO has made this revolution possible.

In the early eighties one still had to go to the VLA site to do the reduction of your own data: down there it took no more than a few minutes to get 100 CLEAN components, and this breathtaking speed could only be reached with the super-modern VAXes at the site. Many Italian astronomers have beautiful memories of nights spent at the VLA site, working in the AIPS-cage (after having struggled to reserve time on the VAX — only the feeding of sharks must have been a more ferocious and bloody show) and consuming gallons of (according to Italian taste¹) awful coffee.

The next 50 years of NRAO will no doubt be quite different in nature, with personal contacts reduced because of high-speed electronic connections. Anyway, the future of radio astronomy is bright, owing to intense international collaborations and the new generation radio telescopes. Obviously the role of NRAO in the planned new generation instruments will be fundamental. The international radio astronomical community can only rejoice.”

Signed for The Director and the Staff of the Istituto di Radioastronomia INAF by Dr.ssa Luigina Feretti.

Improvements of interest in 31DEC07

We expect to continue publishing the *AIPS Letter* approximately every six months along with the annual releases. There have been quite a few changes in 31DEC07 in the last six months. Many have been the usual bug fixes, but there have also been a number of new verbs, tasks, and procedures as well as significant improvements in existing tasks. New tasks include ODSUB to apply a model to visibility data with spectral-dependent corrections, VLANT to correct phases for corrections in the antenna locations at the VLA, CUBIT (written by Judith Irwin) to fit models to full spectral cubes, and UVDI1 to subtract the average of one *uv* data set from another. New verbs include VLA, EVLA, VLBA and HSA which read an antenna table and place suitable antenna numbers on the *POPS* stack.

31DEC04 through 31DEC07 use a new numbering scheme for magnetic tape logical unit numbers that is incompatible with previous versions. Thus all tape tasks and the server TPMON must be from one of these four releases. Other than this, 31DEC07 is compatible in all major ways with the with the 15OCT98 and later releases. There are significant incompatibilities with older versions.

¹Editor's note: not just Italian...

UV data calibration and handling

It has come to our attention that users now trust *AIPS* altogether too much. These users have assumed that all *AIPS* tasks apply a flag table if it is present. But that is not the case. The simple rules are: if FLAGVER is not an adverb to a task, flag tables are not applied in that task. If DOCALIB is not an adverb to a task, then SN and CL tables are not applied in that task. If DOBAND is not an adverb to a task, then bandpass tables are not applied in that task. We have begun a review of all *uv*-data tasks to determine which tasks would benefit by having flagging and calibration options. A number of changes have been made to the calibration system to enable the changes to the old, non-calibration tasks. Note that the package of data reading-flagging-calibration routines have limitations, primarily to restrict the output data to one sub-array and one frequency ID. Thus the newly modified tasks, while acquiring new capabilities, may have some of their previous capabilities curtailed.

The most sweeping change to the calibration routines is to allow the data to be in sort orders other than 'T*' where possible. For non-time sorted data, flagging is limited to 6000 flags applying to the full data set. Time sorted data work when no more than 6000 flags apply to a single time. If the data are not in time order, then time-dependent calibrations, DOCAL true and DOBAND greater than 1, are not allowed. Time independent bandpasses and polarization calibration are allowed in all sort orders. The tasks that apply flagging with their own flag routines, UVCOP, AVSPC, CVEL and SDVEL, were also modified to work with greater restrictions on data that is not in time order. The labeling and selection of polarizations in the calibration package was generalized to allow linear polarization (with many limitations), to allow a few additional STOKES values, such as 'CROS' and 'RLLR', and to allow specification that the I Stokes value is to be "formal" I. This last means that Stokes I is flagged unless *both* RR and LL or XX and YY are present. The new STOKES values 'F' and 'FQU' request this. See the revised STOKES help file for details.

Tasks FUDGE, UVLSF and UVLSD have been rearranged to offer the full suite of calibration adverbs. FUDGE is a "paraform" task used as a model for the creation of many "simple" one-data-set-in, one-data-set-out tasks and provides a pattern for changing many of the older, non-calibration tasks into the more complete form. The other two tasks really do not need the calibration adverbs, but flagging, particularly on a channel-dependent basis, will improve their performance.

UVMTH

Recently it was discovered that the VLA had developed a "correlator bias." This appears as a baseline, spectral channel, polarization, and IF dependent offset in the visibilities. When imaged, an object appears at the phase reference position (that used during observing) which cannot be Cleaned like true sources. In the case of the VLA, it appears that this bias has been present from September 2006 until late in June 2007. It is known to affect spectral-line observations at 25-MHz bandwidth, producing an "object" of some 100's of micro-Jansky. It is believed not to affect continuum data at this bandwidth. We have been unable to determine whether spectral-line observations at narrower bandwidths are similarly affected. The effect is seen only when there are adequate integration time and bandwidth to reduce the noise below these levels and when there is no strong source at the field center masking the offending object. It is thought that the VLA bias is only mildly time dependent, allowing the bias to be measured over significant periods during an observing run and applied to other times during that run. Measurements on consecutive days did show significant changes in the bias. UVMTH is a newly-upgraded task that can help VLA observers deal with this problem. It may also be useful with GMRT data, for which there is also evidence of correlator bias.

UVMTH averages the visibilities in a *uv* data set on each baseline and correlator and adds them to, subtracts them from, multiplies them by, or divides them into the visibilities in another data set. It was written in 1989 and then mostly forgotten. With the reports of correlator bias, this task was revised to apply all of the data selection, calibration, and flagging options to the first input *uv* data set while averaging each baseline and correlator over all TIMERANG. The result is applied to the second data set which is read, modified, and written with no application of data selection (other than sub-array and frequency ID), flagging, or calibration. For the VLA correlator bias problem (and a similar one sometimes seen in GMRT data), the user should specify the same data set for both inputs. SOURCES should select those pointings with no significant source at the center of the field and FLAGVER should select flags for at least these sources, but all calibration should be turned off. The default subtraction operation should be selected. The new task UVDI1 was written to do

a similar operation. It requires a single-source data set for the averaging and does not apply flagging or calibration to either data set. It can time-average the output difference data, however. UVAVG has a similar operation which subtracts a full average of a data set from itself with optional averaging. At present, this SUBT operation functions correctly only on single-source data, and there is some thought to removing the option when calibration adverbs are added to UVAVG.

FILLM

FILLM is the old task that converts the VLA on-line data format into an *AIPS* format. It applies to data from the now old ModComp controlled system and from the new control system (through IDCFAF) as long as the old correlator is in use. At the end of 2006, it was revised to handle the data scaling difference between the ModComp output and the new IDCFAF output. Two significant bug fixes were patched into the 31DEC06 release as well. One caused the loss of the first data record from the next input disk file. At best this meant simply a loss of a single integration time. But, if the mode changed at the file boundary, that change might have gone unrecognized. If the number of polarizations or IFs changed, then errors would be recognized quickly when writing the CL table. A lingering fear that data might have gotten mixed without recognition of a problem remains. The second “error” was a poor decision to use an embedded blank in station names for the VLA. This did not confuse *AIPS*, but other software packages were not so lucky. The names are now VLA: *xxx* and EVLA: *xxx*, where *xxx* is the station name, such as W1 or N72.

Changes which did not make it to patch include the addition of support for Z mode observing. In that mode, true zero-lag data are retained, offering the possibility of new tasks to re-appodize the spectra. The master pad is now correctly recognized as a valid station. EVLA antennas do not have a “front-end” T_{sys} , so the back-end one is used for them at all times. FILLM was changed to write the data with the “normal” *AIPS* axis order: Stokes, frequency, and IF. Previously, it wrote whatever order was convenient to the observing mode, usually frequency, Stokes, and IF, which caused trouble when comparing freshly loaded data sets to data that had had some *AIPS* processing down to them. Support was added for the ten “official” EVLA band codes and separation frequencies. The new telescopes are able to, and do, observe at frequencies which were not understood previously. The default handling of TIMERANG was changed to that which users now expect and the mischievous adverb DOALL was removed. The bad exponent message was curtailed, and a summary of bad exponents was added at the end. It was discovered that FILLM would load two data files into one output, without distinguishing them, even when the two were taken with different configurations. Code sensitive to any change in antenna position greater than 0.02 ns was added to prevent this. Note that this test cannot be made more sensitive for ModComp-recorded data. The antenna positions recorded by the ModComps are changed continuously for earth tides and polar wander. The routine that attempted to correct antenna locations for the shift that must occur when Pie Town is used at Q band was found to be defective. It was generalized so that it will work without reference to band.

VLANT

When an antenna at the VLA is moved to a new station, it is pressed into service after a brief check of its pointing parameters. At a later date, a significant period of time is used to observe a wide variety of calibration sources in order to determine improved estimates of the true locations of the antennas. These data are reduced by the data analysts using a procedure that invokes *AIPS* tasks CALIB and LOCIT and suitable changes are then made to the antenna database at the telescope. These changes, as increments, are also added to a hand-maintained text file available to users. The web site contains an access tool to these text files and a complicated prescription for how the user should add up corrections and apply them to his/her data. The changes were to be made, one antenna at a time, via CLCOR operation ANTP.

We suspect that users seldom did this; the corrections, while useful, were usually small enough that they were not worth the trouble for all but high-frequency data. To rectify this, a new task called VLANT was written. The analysts’ text files for all available previous years were cleaned up to remove various formatting issues and a few obvious errors and then included with *AIPS*. VLANT fetches the latest year’s text file from the web, if necessary, in order to have all of the analysts’ recent changes. (Note that this implies that *AIPS* must be running on a machine that has web access, at least for data that require the latest text file.) Then VLANT performs the complicated prescription for all antennas and does the modification of the AN table and

a new CL table for all antennas at once. Now there is no excuse for not applying the best known antenna locations to a users' data. Note, that antenna position information continues to be updated until the antenna is moved from its pad. This can be many months in some cases.

OOSUB

OOSUB is a version of UVSUB written using the object-based classes of such tasks as IMAGR. It offers the frequency-dependent options found in IMAGR to correct for primary-beam and spectral-index effects. It also has a special option not available anywhere else. This function allows the user to omit from the OOSUB all CCs that are inside the highest frequency beam, outside the highest frequency beam, inside the lowest frequency beam, or outside the lowest frequency beam. The special option will allow the subtraction of partial models to enable better imaging of the inner portions of the data and will probably have other uses when we get clever. It will be interesting to experiment with modifications of PEELR to, for example, subtract a model in a frequency-dependent manner, but add it back without the frequency dependencies.

Other *uv*-editing matters

TVFLG and **SPFLG** were changed to use modern rules for FLAGVER, rather than their curious rules. Now, if a flag table is used on the input, it is copied to a new output flag table and the newly-created flags are appended. A fatal error when looping over channels and IFs to apply an existing clip pattern was corrected. The error arose only when one of the channels had no data.

WIPER was changed to label sub-windows as well as those displaying all of the data. It now will display up to 2 antenna pairs contributing to the sample under the cursor, with a plus sign added if there are more than 2. It will now accept any reasonable STOKES specification, displaying all that are requested and flagging things appropriately. The handling of hour angles, elevations, azimuths, and parallactic angles was corrected.

FLAGR was changed to look for averaging intervals containing only one time (at the ends of scans) which can be reasonably included in the previous interval within that scan. Otherwise, they are flagged since they cannot have an rms. The data buffer sizes were increased to allow for four polarizations with large numbers of antennas and IFs. Other changes to improve reliability were also made.

EDITR was made more forgiving about format errors when reading user-provided numbers. This forgiveness also affects EDITA, SNPLT, SCMAP, and SCIMG.

Other *uv*-display matters

LISTR and **RLDIF** were changed to average all requested spectral channels in the MATX display and to loop appropriately over IF. **RLDIF** was corrected to use vector averaging between scans as well as within scans. **LISTR**'s **GAIN** display required a correction to its format.

PRTUV was given the option to print real and imaginary parts of the visibility rather than amplitude and phase. The alternate display of random parameters was corrected and generalized.

VPLT was given a new option '**IFRA**' to plot the ratio of two IFs.

UVPLT had all restrictions on STOKES removed; all requested polarizations will be plotted together. It had corrections for VLBI to the computation of hour angle, azimuth, elevation, and parallactic angle. The display when using such axes was changed to show the usual information as well as the REFANT.

DTSUM confused antenna number with position in the antenna file. The displays of which baselines are in the data were revised to do the total summary, followed optionally by a scan-by-scan summary. More user and internal controls over the displays were added.

ANBPL computed elevation incorrectly for single-source files since the source position information was never read. This oversight was corrected.

SNPLT subtracted 6 hours erroneously from all hour angles.

Other *uv*-related matters

- SPLIT** could lose the value of *DOCALIB*, *DOPOL*, *DOBAND*, and *DOBL* if a source in the list had no data found. All sources later in the list would be written out, but with no calibration applied and no warnings. The code was corrected and a patch issued.
- QUACK** was given new options to limit its function by Stokes and IF.
- INDXR** was changed to work with single-source files as well as its usual multi-source files. *CALIB* will use this information to good advantage in self-calibration. *INDXR* was changed to support the new EVLA band definitions.
- CALIB** extended index table times by 5 seconds in each direction. This is disastrous when using a short averaging time within scans, so corrected the round-off protection to 0.1 sec.
- LOCIT** was given the option of using the difference in the phases of the 2 IFs to solve for antenna positions. This means that baseline errors are multiplied by the small frequency difference of the 2 IFs rather than the full observing frequency. An optional slope in phase with time was also added. All of this was needed to find the VLA master pad!
- VBGLU** was improved to reformat flag tables for the new IF nomenclature. *VBGLU* was changed to refuse to glue multi-frequency ID data sets. This is better than making a mess of the job. Gluing of the *STATE* column in PC tables was corrected. A multi-dimensioned character column is confusing to FITS tables, or at least to *AIPS*.
- CVEL** had a programming error cause *CALINI* to issue numerous warnings erroneously and another error that could cause an abort (with some compilers) before the extension files were copied.
- BLCAL** is now recommended with some VLA data. It was corrected to work with normal multi-source data sets. It did bad things when there was non-zero delays in the CL table due to a misunderstanding about units. The task was also changed to use dynamic memory instead of rather small buffers. Fixed scaling of calibrator source models to be properly IF dependent. (This change was also made a wide variety of other tasks.)
- VLBAUTIL** was changed to try a variety of download methods to fetch files from the web; *wget* is not universally available. It was also changed to split frequency IDs before calling *MSORT*. That task could take a semi-infinite time to run when two frequency IDs observed at the same time were appended. Such data usually need to be separated, then sorted and *VBGLU*ed.
- UVFIX** was given the option of suppressing differential aberration in its computations. An error of the reference channel in frequency was made for uncompressed data when the source position was changed.
- CLCOR** had a second-order error in phase corrections for a change in source position which has been corrected.
- SMOOTH** was left out of 17 tasks that otherwise apply calibration. The oversight was corrected.
- FITLD** was changed to support bandpass tables within the IDI convention.
- MULTI** used parameters to the precession routine other than those used normally in *AIPS*. The apparent coordinates in the output source table were somewhat affected.

Miscellaneous

New verbs

Four new, related verbs have been written. They are named *VLA*, *EVLA*, *VLBA*, and *HSA* and place a temporary numeric array on the *POPS* stack for processing by other verbs. The verbs look in the specified antenna file to make a list of the antenna numbers that match (or do not match) the specified array. The temporary numeric array is a new concept in *POPS* and required a number of modifications to the language

and its fundamental operators. Verbs such as MAX and PRINT can now handle the output of such verbs. There are, however, a number of limitations to the use of these verbs, the worst being that one cannot use two of them in the same command. If you want to set an adverb array to, for example, the list of EVLA antennas, you must include an equals sign; thus ANTENNAS = EVLA.

Other matters

- TABED** was changed to default the output file name parameters individually to those of the input. An inadvertent OUTDISK with otherwise blank OUTNAME sent the output table to very unexpected places.
- DFQID** failed to renumber frequency IDs in the tables.
- DQUAL** failed to renumber the sources in tables.
- FXTIM** failed to correct the TIME and TIME RANGE in tables.
- TAPLT** was corrected to recognize blanked values. The scaling was changed for binned plots to ignore empty bins and plot them at the bottom anyway. The Y function string displayed was corrected.
- KNTR** was changed to allow ZINC less than zero to plot the planes in the opposite order. Error reporting, especially for crazy polarization lines, was improved.

Analysis

CUBIT

Judith Irwin of the University of Toronto wrote CUBIT quite some time ago to fit spectral-line cubes with models of density as a function of radius in, and height above, a galactic plane and of velocity as a function of radius. She has continued to develop the code and help file over the years and has now kindly made it available to all AIPS users. We have modified the code to modernize its structure and to replace modest fixed arrays with dynamic memory, allowing cubes of any size to be fit. This task is not easy to use — there are a great many parameters to fit and it takes patience to find good minima in the very large Chi-squared space. The results are rather more general than those found fitting solely moment images. If you use this task to analyze your data cube and make reference to the results in a paper, you should reference the original paper describing this task: Irwin, Judith A. 1994, “Arcs and bridges in the interacting Galaxies NGC 5775/NGC 5774,” ApJ, 429, 618–633.

Other analysis matters

- FLATN** was changed to allow it to make images of the expected noise or weight (one over noise squared) from mosaic observations. Repaired bug which caused the NOISE adverb in mosaicing to be ignored, leaving all pointings with the same basic weight. FLATN will now use the header parameter ACTNOISE, if present, as the default for NOISE.
- VTESS** was changed to test whether the incoming images for a mosaic are on the same geometry as required. Changed UTESS, LTESS, and STESS to apply these same tests. Corrected a bug in STESS and added a note to its help file pointing out that FLATN now does its job without all the problems with image geometry.
- SAD** was enhanced to be able to use the optional noise image when selecting the islands to be fit and during the decision process in which fits are accepted or rejected. Previously, the noise image was only used to modify the uncertainties assigned to the results of the fit.
- SMODEL** was applied without the $w \times z$ term which is important at large angles. The extended source portions of SMODEL were also ignored by low-level OOP routines.
- Models** of calibration sources must be scaled by the source total flux, which is different for each IF in the data. CALIB did this correctly, but most other model application tasks ignored this nicety.
- UVCON** was corrected for problems with its computation of elevation and its creation of antenna files.

Imaging

IMAGR

The TV menu was rearranged to put the choices that stop things well away from the choices that are frequently used. This should reduce the premature terminations. Added two new TV options - to turn off Cleaning of a facet now and to turn it back on. Changed the beam histogram computation to avoid edges and corners which can be pathological due to accumulated round-off errors. Moved the forcing of D03DIM to true for OVERLAP=2 mode to before the images are created. It is too late after that. Fixed the positioning of the inscribed circle which may be drawn optionally to guide the eye in setting boxes. Cleaned up the handling of the on-line help which had stopped working for some options and were too demanding on the programmer for the ones containing field numbers.

SETFC had an error in the formula for phase error which was corrected in the help file and code. The output is unchanged because the default phase error was changed accordingly. Fixed errors handling a radius of zero. Changed the channel to avoid channel 1 which is often fully flagged. Changed the round up levels to recommend larger images rather sooner.

General and programming matters

ATPS does much of its high-performance computing with a cpu-based emulation of an array processor. General mathematical routines function on data located in a large memory accessed via pointers. Such routines are usually highly optimized by compilers. Previously, the size of this memory was determined at compile time and set by the local *ATPS* Manager. A fundamental conflict arose in choosing this size between those people at a site with older, small-memory computers and those with large problems to solve. Furthermore, since the size of the memory was set independent of the problem, the programmer had to write algorithms which could adapt to almost any size of memory. Such algorithms tend to be limited and/or much more difficult to write.

In 31DEC07, the pseudo-AP has been changed to use dynamically allocated memory. The programmer specifies the amount of memory needed when the AP is allocated by QINIT. Pointers in regular code remain normal integers since we assume that we will not allocate more than 8 Gbytes. However, the actual subscripts to APCORE are now required to be LONGINTs which are the sum of an offset called PSAPOF and the calling routine's pointers. Computers which use 8 bytes for address pointers in C require us to use INTEGER*8 "LONGINTs." The pre-processor automatically handles the conversion of LONGINT in the *ATPS* Fortran code into the form needed by the local computer.

We had hoped that this change would improve performance by allocating more memory in many cases to avoid the use of disk inside algorithms. The standard Y2K tests, however, showed no change in performance, suggesting that 20-Mbyte APs are already adequate for these tests. Linux disk I/O does a great deal of in-core caching as well, which would reduce the differences between "disk" and in-memory algorithms. The main purpose of this change was to allow for the creation of new algorithms which would be easier to create and more efficient by never invoking disk files.

MNJ reports now contain a list of files noted by the cvs update, with a character showing whether the file was new, updated, different at the remote site, etc.

install.pl was changed to cause cvs to ignore more of the files it should ignore in its reporting. The options of the rsync calls were altered to improve reliability. Calls to chgrp were added to change the group ownership of files to that requested by the installer.

Servers for messages and Tektronix plot emulation are programs that run inside an xterm or other suitable window program. The binary releases require the location of supplied libraries when these programs start in the spawned xterm. The files UNIXSERVERS and XASERVERS were changed to have shell- and architecture-dependent case statements to try to set this information where it can be used. The situation with these servers seems to have improved with this addition, but mysterious troubles continue to be reported.

Patch Distribution for 31DEC06

Important bug fixes and selected improvements in 31DEC06 can be downloaded via the Web beginning at:

<http://www.aoc.nrao.edu/aips/patch.html>


Alternatively one can use *anonymous ftp* to the NRAO server <ftp.aoc.nrao.edu>. Documentation about patches to a release is placed on this site at [pub/software/aips/release-name](#) and the code is placed in suitable subdirectories below this. As bugs in 31DEC07 are found, they are simply corrected since 31DEC07 remains under development. Corrections and additions are made with a midnight job rather than with manual patches. Since we now have many binary installations, the patch system has changed. We now actually patch the master version of 31DEC06, which means that a MNJ run on 31DEC06 after the patch will fetch the corrected code and/or binaries rather than failing. Also, installations of 31DEC06 after the patch date will contain the corrected code.

The 31DEC06 release had a number of important patches:

1. CALIB handled scan times badly when averaging within a scan *2007-01-02*
2. UPDCONTROL in the MNJ used obsolete syntax for sort *2007-01-10*
3. UVFIX did not contain the latest leap second *2007-02-27*
4. CLCOR operation SUND did not work *2007-02-27*
5. VBGLU did not handle the PC table STATE column correctly *2007-02-27*
6. CVEL had a bad call sequence which could cause an abort *2007-02-27*
7. SNSMO had a bad call sequence which could cause bad smoothing of rates including flagging them *2007-04-26*
8. UVFIX had a frequency error for uncompressed data only and CLCOR had a minor conceptual error both affecting phases after a position shift *2007-04-26*
9. BLCAL and UVFND set the integration time to one day, causing bad amplitude calibration when there were rates and delays *2007-04-26*
10. aips.1 man page was lost *2007-05-04*
11. FILLM skipped a record at ends of file which could lose a data sample and possibly cause confusion if the mode changed *2007-05-24*
12. SPLIT lost the calibration flags when a source was not found so that later data did not have calibration applied *2007-06-10*
13. SNPLT lost data from phase plots of PC tables due to failure to check for wraps and got hour angles wrong by 6 hours *2007-06-14*
14. FILLM used a blank in the middle of some station names, confusing other software packages *2007-06-16*

AIPS Distribution

We are now able to log apparent MNJ accesses and downloads of the tar balls. We count these by unique IP address. Since dial-up connections may be assigned different IP addresses at different times, this will be a bit of an over-estimate of actual sites/computers. However, a single IP address is often used to provide *AIPS* to a number of computers, so these numbers are probably an under-estimate of the number of computers running current versions of *AIPS*. In 2007, there have been a total of 750 IP addresses so far that have accessed the NRAO cvs master. Each of these has at least installed *AIPS* and 186 appear to have run the MNJ on 31DEC07 at least occasionally. During 2007 more than 180 IP addresses have downloaded the frozen form of 31DEC06, 117 in binary form, while more than 517 IP addresses have downloaded 31DEC07, 346 in binary form.



June 30, 2007



AIPSLETTER
National Radio Astronomy Observatory
Post Office Box O
Socorro, NM 87801-0387
USA

FIRST CLASS

To:

Library
Natl Radio Astronomy Obs.
520 Edgemont Road
NRAO

DATED MATERIAL-DO NOT DELAY

A I P S L E T T E R

Volume XXVII, Number 2: December 31, 2007

A newsletter for users of the NRAO
Astronomical Image Processing System

Written by a cast of *AIPS*

Edited by

Eric W. Greisen

National Radio Astronomy Observatory

P.O. Box O, Socorro, NM 87801-0387

505-835-7236; Fax 505-835-7027

daip@nrao.edu

<http://www.aoc.nrao.edu/aips>

General developments in *AIPS*

Current and future releases

We have formal *AIPS* releases on an annual basis. While we offer a full binary installation method for both the frozen and development versions for MacIntosh OS/X (PPC and Intel chips), Solaris, and Linux systems, all architectures can do a full installation from the source files. The current release is called 31DEC07 and is now “frozen.” If you took a development copy of this version at some earlier date, you should use the “Midnight Job” (MNJ) to bring it up to date. You need to run a MNJ only once in 2008 to convert your copy of 31DEC07 into the frozen version. When patches to 2007 are announced, you may apply them with the MNJ. This *AIPS*Letter is intended to advise you of corrections and improvements in this release.

We have begun a new version, called 31DEC08, which is now under development by the *AIPS* Group. You may fetch and install a complete copy of this version at any time. Having fetched 31DEC08, you may update your installation whenever you want by running the MNJ which uses cvs, rsync, and transaction files to copy and compile the code selectively based on the code changes and compilations we have done. We expect users to take their source-only or binary version of 31DEC08 *AIPS* over the Internet (via *anonymous* ftp). Both versions require you to copy the installation procedure `install.pl` via ftp; the source-only version also requires you to ftp the 90-Mbyte `31DEC08.tar.gz` compressed tar file.

From `mnj.aoc.nrao.edu`, the MNJ will serve up *AIPS* incrementally — or as a whole — using the Unix tool `cvs` running with `anonymous` ftp. Binary MNJs also use the `rsync` tool. Linux sites will almost certainly have `cvs` installed; other sites may have installed it along with other GNU tools. Secondary MNJs will still be possible using `ssh` or `rcp` or NFS as with previous releases. We have found that `cvs` works very well, although it has one quirk. If a site modifies a file locally but in an *AIPS*-standard directory, `cvs` will detect the modification and attempt to reconcile the local version with the NRAO-supplied version. This usually produces a file that will not compile or run as intended.

AIPS is now copyright © 1995 through 2008 by Associated Universities, Inc., NRAO’s parent corporation, but may be made freely available under the terms of the Free Software Foundation’s General Public License (GPL). This means that User Agreements are no longer required, that *AIPS* may be obtained via `anonymous` ftp without contacting NRAO, and that the software may be redistributed (and/or modified), under certain conditions. The full text of the GPL can be found in the 15JUL95 *AIPS*Letter and is included with every distribution in file `$AIPS_ROOT/release-name/COPYING`.

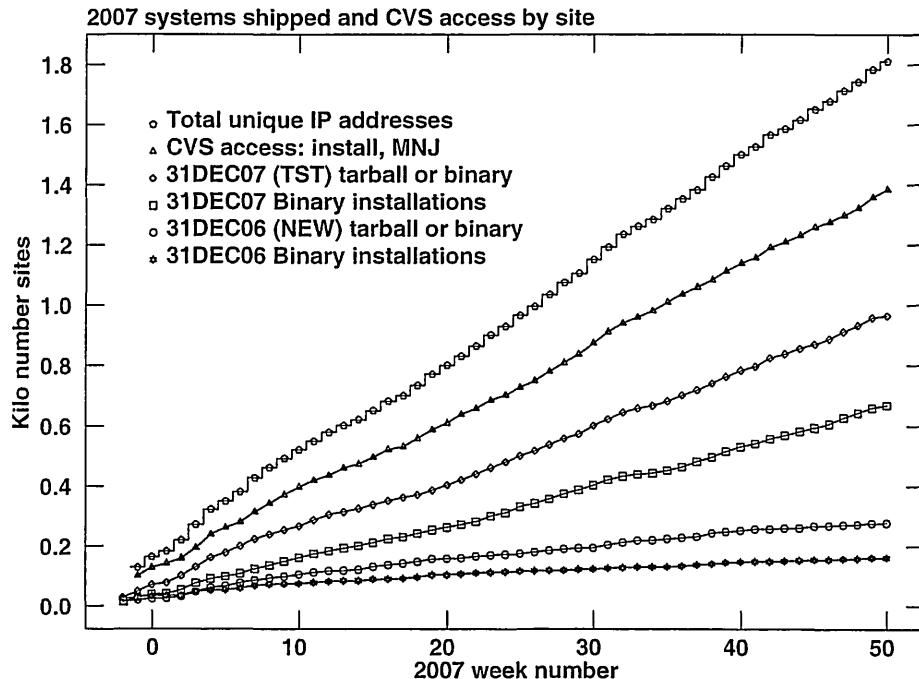
Installing a new version

If compiling locally, new releases must be installed from the tar ball for that release. If using the binary installation, a full new installation must also be done with `rsync`. The `cvs` system requires this. When installing a new *AIPS* release in a system that already has a previous release, we recommend that `install.pl` be used and that the previous release be left in place, at least until the installation has been seen to work. If you do this, then you will not have to re-edit the disk, printer, and tape lists and can simply skip all those pages in the `install.pl` menus. The old `$HOME/.AIPSRC` file may be left in place, but it will need to be edited. The lines giving the `DOWNLOADED` and `UNPACKED` parameters should be deleted and the `CCOMOPT` line should be changed to point to the current release rather than the previous one — the `-I` parameter really should be `-I$INC` but it gets its full path name instead. This forces a re-edit with each release. If you have made special versions of `UPDCONFIG` and `do_daily.host`, you should preserve them under new names and restore them after the install. If you have an odd set of *AIPS* versions, the `$AIPS_ROOT/AIPSPATH.*SH` files may need to be edited after the install to set the desired versions.

For Linux, Solaris Ultra, and Macintosh systems, a binary installation could be available from CDrom, supported by `install.pl`. Alternatively, the frozen version may be installed with the binary installation method now present in `install.pl`. The `ftp` site for downloading files directly has been eliminated.

AIPS Distribution

We are now able to log apparent MNJ accesses, downloads of the tar balls and `rsync` accesses. We count these by unique IP address. Since dial-up and some university connections may be assigned different IP addresses at different times, this will be a bit of an over-estimate of actual sites. However, a single IP address is often used to provide *AIPS* to a number of computers, so these numbers are at the same time an under-estimate of the number of computers running current versions of *AIPS*. In 2007, a total of 277 different IP addresses downloaded the frozen form of 31DEC06 and 965 IP addresses downloaded 31DEC07 in tarball or binary form. Fully 1385 IP addresses accessed the NRAO cvs master. Each of these has at least installed 31DEC07 and 304 appear to have run the MNJ on 31DEC07 at least occasionally. The total number of unique IP addresses in these three lists was 1811. 161 sites accessed 31DEC06 in binary form, while 669 sites used the binary form of 31DEC07. The attached figure shows the cumulative number of unique sites, cvs access sites, tar-ball/binary download sites and binary access sites known to us as a function of week in 2007. These numbers represent substantial increases over those for 2006.



Since the registration system, always under-utilized, has now been abandoned, we are left with analysis by IP address. The table below lists the IP addresses for 2007 by the final qualifier for shipments of 31DEC07, 31DEC06, and access to the cvs site. The numbers in the cvs column include those sites that install or run a midnight job for these releases. The comments come from what appears to be a semi-official list of Internet codes. Sorting is on the “unique” column, which counts unique IP addresses over the other three columns:

Code	31DEC06	31DEC07	cvs site	unique	Comments
net	19	110	353	403	Network
edu	35	220	300	373	US Educational
jp	29	58	69	88	Japan
uk	16	48	46	69	United Kingdom
de	7	25	37	53	Germany
com	11	29	31	49	US Commercial
in	12	32	10	38	India
es	5	23	30	34	Spain
nl	5	22	25	33	Netherlands
au	8	19	23	33	Australia
it	6	23	17	33	Italy
ca	2	23	22	31	Canada
org	3	13	19	24	Non-Profit Organization
pl	3	12	15	20	Poland
fr	7	11	11	17	France
za	4	12	7	17	South Africa
gov	5	10	11	15	US Government
mx	7	10	7	13	Mexico
ru	6	9	5	13	Russian Federation
br	1	7	9	13	Brazil
mil	0	11	7	12	US Military
ie	3	6	6	8	Ireland
se	1	5	3	7	Sweden
gr	1	1	4	6	Greece
tw	2	4	4	5	Taiwan
pt	0	4	5	5	Portugal
cn	2	2	2	4	China
be	1	2	2	3	Belgium
cl	0	2	2	3	Chile
pe	0	2	3	3	Peru
hu	0	1	2	2	Hungary
ua	0	2	0	2	Ukraine
arpa	1	0	0	1	Old style Arpanet
ar	1	0	1	1	Argentina
mu	1	0	0	1	Mauritius
il	0	1	1	1	Israel
no	0	1	0	1	Norway
fi	0	1	0	1	Finland
ch	0	1	0	1	Switzerland
eg	0	1	1	1	Egypt
lt	0	1	0	1	Lithuania
hn	0	0	1	0	Honduras
None	0	2	6	5	
Unknown	73	199	288	368	
Total	277	965	1385	1811	

Preview of coming attractions

The 31DEC08 release already contains a few minor changes that we decided were a bit risky or not needed in 31DEC07. There are counters of I/O transactions by ZMIO and ZFIO that can be displayed when tasks end by issuing the AIPS verb SETDEBUG to a value greater than zero. This will be of use when we re-examine imaging algorithms to try to reduce I/O to improve performance. The enormous data sets which the EVLA will produce require us to examine this issue closely. The task IMEAN has had the DOINVERS option added, which led to changes to IMSTAT as well. The *CookBook* is being reviewed thoroughly and updates, which will be available to all versions, will appear soon.

Improvements of interest to users in 31DEC07

We expect to continue publishing the *AIPSLetter* every six months along with the annual releases. There have been a number of changes in 31DEC07. In the last six months, we have developed new tasks TAPPE to append one version of a table to another version, TYSMO to edit and smooth system temperature tables, TYAPL to remove and/or apply system temperature corrections to uv data, FLOPM to reverse the spectral order of a data set with or without special fixes for a VLA issue, and HAFX to determine if there is a time error in the (u, v, w) values in a data set. A stand-alone program REBYTE was created to convert the byte order of whole *AIPS* directories. This should ease users' transitions from Solaris and Mac PPC to Linux and Mac Intel computers. Described in the June 30, 2007 *AIPSLetter* were changes in 31DEC07 including new tasks OOSUB to apply a model to visibility data with spectral-dependent corrections, VLANT to correct phases for corrections in the antenna locations at the VLA, CUBIT (written by Judith Irwin) to fit models to full spectral cubes, and UVDI1 to subtract the average of one uv data set from another. New verbs include VLA, EVLA, VLBA and HSA which read an antenna table and place suitable antenna numbers on the *POPS* stack.

31DEC06 contains a revision of FILLM required to support the new data form which the VLA began producing on June 27, 2007. VLA users must upgrade their copy of *AIPS* to read these data. Also, 31DEC07 contains a significant change to FITAB. This task is used to fill the archive with partially processed data from the VLBA and more fully processed (pipelined) data from the VLA. Versions of *AIPS* prior to October 15, 2007 are unable to read these data.

31DEC04 and later releases use a new numbering scheme for magnetic tape logical unit numbers that is incompatible with previous versions. Thus all tape tasks and the server TPMON must be from 31DEC04 or later. Other than this, 31DEC07 is compatible in all major ways with the with the 15OCT98 and later releases. There are significant incompatibilities with older versions.

UV data input/output

FILLM

FILLM, the task that translates VLA on-line data into *AIPS*, was changed quite a bit during the first half of 2007; see the June 30 *AIPSLetter*. The most significant change in the last half of 2007 was the long overdue change in the recorded times to the center of the sample interval. Previously, the times differed from the center time by one-half of the on-line integration time (traditionally 10 seconds, but no longer limited to that). This affects usage of UVFIX which required a time offset to be applied, but now only for data filled by versions of FILLM prior to July 5, 2007. A message is written by FILLM warning you of this change.

Other changes in FILLM include correcting the help file to tell the truth about CPARM(2)=1, which only masks the T_{sys} fluctuating error condition, correcting the total bandwidth recorded in the source table, and correcting a mismatch in the default nominal sensitivity when it is recorded as zero. The new capability of FILLM to load correlation coefficients (rather than data corrected by system temperature) was enhanced by writing a header keyword to record the presence (or absence) of correlation coefficients and by writing correct weights for correlation coefficients. The weights for correlation coefficients are constant with scaling only for bandwidth and integration time.

FITAB

FITAB writes visibility data in the form of FITS binary tables rather than the old, and now mildly deprecated, random groups form. The fact that the details about the visibility data including numbers of samples, channels, IFs, polarizations, and the like are not available to the reading program until after it has had to swallow history, antenna, calibration, source, and other tables made construction of reading programs for this form difficult. When FITAB was first written, it was decided to take a short-cut and actually provide the needed information in the main header in the form of special *AIPS* HISTORY cards. Unfortunately, when a non-*AIPS* software package sees these data, it should save these history cards and re-emit them with output data and/or images. Then when *AIPS* re-encounters the data, these history cards are used even though they no longer correctly represented the data. Data reading programs were changed to reduce their use of HISTORY *AIPS* cards in some circumstances and so to reduce the likelihood of this error arising.

But, to avoid the error entirely, versions of FITAB after October 15, 2007 no longer take this short-cut. Corresponding versions of UVLDD, FITLD, and PRTPP were created to handle this more difficult case. VLBA data first appear in the archive not in raw correlator form but in a mildly processed form written by FITAB. The raw data arrive in the archive substantially later. Thus, VLBA users will need to update to the full frozen version of 31DEC07. Additionally, the *AIPS* VLA pipeline is being run on many VLA observations with its results available in the archive. The visibility data also use FITAB, so any users of these processed VLA observations from the archive will also need the last version of *AIPS*.

Calibration application

VLA u, v, w error

Deep wide-field 21-cm observations of a particular direction made at two widely-spaced epochs have revealed an apparent error in the VLA on-line computation of (u, v, w) used by *AIPS* in imaging and other operations. The error appears to have crept in sometime in 2004 and to have remained until the ModComps were replaced with modern computers and new software at the end of June 2007. The most probable nature of the error is an offset in time of 10 seconds between the recorded (u, v, w) and the correct ones. The new *AIPS* task HAFX is available to you, if you wish, to study this error in your data. However, simply running UVFIX on your data set prior to imaging will correct any error. Be careful about the offset in time needed by UVFIX for data loaded by versions of FILLM prior to 31DEC07 July 5 (see above).

SPLIT and SPLAT errors

Several errors in the frequencies written by SPLIT and SPLAT when averaging channels in multi-channel data sets were discovered. A fairly recent error caused the frequency increments in the FQ table to be wrong when written by SPLAT and also to be wrong out of SPLIT when the initial increments were not identical. An error of long-standing was made by SPLIT when averaging channels with $BCHAN > 1$, while still writing a multi-channel output file. These errors can cause smearing when doing bandwidth synthesis imaging as well as an image scaling error. The two tasks also did not change the reference frequency when averaging all spectral channels (but not IFs) although that will cause a simple scaling error in imaging. There was also an error in the number of words copied which certainly caused some loss in accuracy of compressed data and, depending on compilers, could have caused rather more serious consequences.

TY tables

Amplitudes in interferometers with digital correlators are found by measuring a correlation coefficient and a system temperature and multiplying the two. Errors in the latter directly cause errors in the visibilities. During the VLA-EVLA transition, we have found “peculiarities” in the measured “nominal sensitivities” used to scale the correlation coefficients to visibilities in Jy. There are occasional, really bad, values and more noise than one would hope. There is an interactive editing task called SNEDT which will allow one to modify a TY table. However, for 31DEC07, we have written two new tasks. The first is TYSMO which copies a TY table to a new one while median-window filtering out the bad samples and replacing them and/or

the good samples with time-smoothed values. The second is TYAPL which will remove one TY table from a visibility data set and replace it with another, presumably smoothed and edited, TY table. This, combined with FILLM's new ability to write correlation coefficients rather than scaled visibilities, will allow new tools for the user to evaluate this crucial step in calibration.

Miscellaneous

3C147 21-cm calibrator model was installed in the usual place (\$AIPSTARS). We now provide complete models for 3C48, 3C138, and 3C286 (at all VLA bands except for P and 4) and some models for 3C147 (K, L, Q, U).

3C286 X-band model had a mysterious error causing all Clean components to be shifted by one pixel.

UVAVG was changed to set time intervals for TB sorted data from the first sample of the interval. For BT data, it still uses its old method of using intervals which are integers times YINC within each day. The full set of calibration and flagging options was added and the 'SUBT' operation was removed. The averaging buffer is now dynamic and so can handle any reasonable size data set.

CORER was changed to support calibration and flagging of the input data set.

TI2HA was changed to support calibration and flagging of the input data set, including allowing the input to be multi-source. Only one source may be output. STUFFR was changed to allow DOCALIB and FLAGVER, although both will use the highest version gain and flag tables in each input file if invoked.

AVSPC was given the full range of calibration and flagging functions.

DIFRL was given the full range of calibration and flagging functions.

NX tables are now created and populated by many of the above tasks as well as UVCOP, FUDGE, UVMTH, UVLSF and FLGIT as they proceed. This avoids the need to run INDXR under many circumstances.

BL tables were seldom used previously and so were overlooked when tables were edited for data selection while being copied. The VLA-EVLA transition has encouraged the use of BLCAL and thus led to this and other corrections affecting BL tables.

Polarization calibration application was found to have a serious error when the application task called the data read initialization for a second time. Fortunately, this error affected only equatorial telescopes (*i.e.*, WSRT) and so did not affect a large number of *AIPS* users.

Data editing

OUTFGVER is a new adverb which has appeared in most, if not all, visibility-data editing tasks. This allows the users to specify whether the new flags will be written to an existing flag table or to a new table which will also include any flags applied on input. If OUTFGVER if ≤ 0 or $> FG_{\max}$, then a new flag table will be created with version $FG_{\max} + 1$ and the input flag table (FLAGVER) if any will be copied to it along with any new flags. If OUTFGVER specifies an existing table, then FLAGVER is not copied to it even when $FLAGVER \neq OUTFGVER$. Tasks TVFLG, SPFLG, DEFLG, FLAGR, EDITR, EDITA, UVMLN, WETHR, and WIPER were changed in this fashion. Other tasks which had more extensive changes are described below.

IBLED was changed to use OUTFGVER on single- and multi-source data. It no longer will apply flags and copy single-source data sets. The options PREV BASELINE and FLAG ALL TIME were added and the FLAG INTERACTIV option was made to work correctly. It now switches IFs with no question when possible. (EDITR remains the interactive editing task of choice in most circumstances.)

CLIP is the new name for CLIPM, which is the task users should have been using. It has the full calibration and flagging input options include STOKES which was previously omitted. It handles single- and well as multi-source files writing OUTFGVER. The old CLIP task was removed.

UVFLG was rewritten to support flagging by source elevation correctly. It now writes OUTFGVER with no input FLAGVER. Flagging in place was dropped, making adverbs BDROP, EDROP, and the previous APARM(1) through APARM(3) obsolete and thereby changing the meaning of the APARMS. Improved the writing of the history file including adding more information and removing mis-information.

FLGIT was given the adverb OUTFGVER, but will still write out a new edited data set when OUTFGVER ≤ 0 .

EDITR and EDITA now both support OUTFGVER as do SCIMG and SCMAP which also use the edit class. NEXT ANTENNA and NEXT BASELINE were changed to get new antennas/baselines rather than recycle old ones. ENTER IF was made friendlier in crowded display modes.

CORER was made to support OUTFGVER and to offer the choice of printing or not and of write a flag table or writing out the edited data. It also acquired all the calibration and input flagging options.

Other *uv*-data changes

FLOPM is a new task written to handle a temporary error made in the VLA-EVLA transition in which sidebands were confused. This 'VLAE' mode, reverses the order of the spectral channels and the sign of the phase, but leaves everything else untouched. The default mode, which was requested by another user, reverses the order of the spectral channels and fixes all of the header and FQ table accordingly while not messing with the sign of the phase.

UVLSF was also changed in an attempt to assist with another transition difficulty, in this case with spectral aliasing on EVLA-EVLA baselines affecting narrow bandwidths. It will now fit up to fourth-order orthogonal polynomial baselines and has the option to add the continuum back to the spectral data. A long-standing error in flagging on rms was corrected.

DBCON was fixed for errors when converting one of the data sets to true Stokes when that data set is compressed.

OOSUB and UVSUB use CHANNEL to mean to subtract or divide the data of that channel but to copy all channels, leaving the rest unchanged. Clarified this in the code and documentation and fixed OOSUB to handle compressed data and copy tables.

MATCH was corrected to take header information for antenna, source, and frequency tables from the input file rather than the master file.

UVHOL was changed to handle polarization states other than the usual RR, LL, RL, LR. Options to average data to all the reference antennas and/or to average in time within a pointing were also added. Adverbs BCOUNT and UVRANGE were dropped and the default of BDROP was changed to zero.

UVDIF was changed to allow the two files to have different states of compression. It is still pretty fussy about which data sets it is willing to compare.

Averaging of visibility data was found to ignore weights in several OOP routines. Corrected this, which affects tasks CL2HF, UBAVG, SCIMG, and SCMAP.

Imaging and analysis

FLATN was corrected; it was writing out noise and weight images *squared* rather than with the correct scale.

COPIXEL was given the ERROR output adverb so that coordinates which are not on, or are illegal for, the image can return a suitable indicator without stopping a script.

Plotting

VPLOT has a new default SYMBOL, uses adverb FACTOR to scale the plotted symbols and optionally connect the symbols with lines, to offer hour angle, elevation, azimuth, and parallactic angle axes, to plot ratios of true STOKES types, to allow averaging for all axis types, and to do time averaging over normal rather than integer-step intervals (same correction as UVAVG above).

TVCPS has an option to read a full image from disk, rather than plot only the subimage which fits on the TV. It was fixed to make a correctly-colored plot when multiple images are involved.

Slice plots default to the min/max of the slice itself, but then used the full image max/min for scaling the plot internally. If the slice flux range is small compared to the image flux range, this causes no end of trouble. Changed all slice plots to ignore the image max/min, fixed them to pass the scaling information in the image catalog header, and fixed the TV slice plotting verbs to use the current size of the TV rather than its full size.

Miscellaneous

TAPPE is a new task to append two tables, then sort the result and eliminate duplicate rows. TABED with OPTYPE 'COPY' does a selective append operation, but does not do the sort and elimination steps.

REBYTE is a stand-alone program to convert a disk directory full of *AIPS* data files to the opposite byte order. Almost all of the conversion is done with full knowledge of the file formats. Only PL file headers and TPUT/TGET and SAVE/GET files are converted with the assumption that four printable characters in a row must be an actual character string. For floating-point variables, this has a 2% chance of being wrong. The RUN WRTPROCS procedures WRTDISK and READISK offer another approach, through FITS-disk files, to this conversion. They make no attempt to preserve plot and slice files nor the TPUT/TGET and SAVE/GET files. Note, however, that the SAVE/GET files may be preserved fully via SG2RUN with either approach.

Batch jobs in *AIPS* use particularly high *AIPS* numbers. Failures to move to base 36 which affected mostly batch were repaired and several more adverbs are now tested against local system limits rather than generic limits found in the help files. AP-using tasks are now allowed in batch queue 1.

gfortran is the GNU Fortran 90 compiler. It has not worked on *AIPS* in earlier tests, but version 4.2.1 was found to produce working load modules on an AMD computer. They run faster than those produced by g77 3.4.4, eliminating the slight advantage that the Intel compiler had on AMDs.

Linux kernels beginning with 2.6.19 (RedHat 2.6.9-55) and continuing until 2.6.23.1 (RedHat 2.6.9-61) had what appeared to be a race condition in the automounter. If an open occurred to mount a file system while it was being un-mounted due to a timeout, the file system could appear to be empty, leading to erroneous "file does not exist" *AIPS* error completions.

Patch Distribution for 31DEC06

As before, important bug fixes and selected improvements in 31DEC06 and 31DEC07 can be downloaded via the Web beginning at:

<http://www.aoc.nrao.edu/aips/patch.html>


Alternatively one can use *anonymous* ftp to the NRAO server [ftp.aoc.nrao.edu](ftp://ftp.aoc.nrao.edu). Documentation about patches to a release is placed on this site at [pub/software/aips/release-name](#) and the code is placed in suitable subdirectories below this. As bugs in 31DEC08 are found, they are simply corrected since 31DEC08 remains under development. Corrections and additions are made with a midnight job rather than with manual patches.

The patch system has changed because we now have binary installations. We now actually patch the master copy of the frozen version. This means that a MNJ run on 31DEC06 after the patches listed below will fetch

the corrected code and/or binaries rather than failing. Similarly, patches announced for 31DEC07 during the next year will be available via MNJ as well as ftp. Installations of 31DEC06 and 31DEC07 after the patch date will contain the corrected code.

The 31DEC06 release is no longer available for installation. It had a number of important patches. They are

1. CALIB handled scan times badly when averaging within a scan *2007-01-02*
2. UPDCONTROL in the MNJ used obsolete syntax for sort *2007-01-10*
3. UVFIX did not contain the latest leap second *2007-02-27*
4. CLCOR operation SUND did not work *2007-02-27*
5. VBGLU did not handle the PC table STATE column correctly *2007-02-27*
6. CVEL had a bad call sequence which could cause an abort *2007-02-27*
7. SNSMO had a bad call sequence which could cause bad smoothing of rates including flagging them *2007-04-26*
8. UVFIX had a frequency error for uncompressed data only and CLCOR had a minor conceptual error both affecting phases after a position shift *2007-04-26*
9. BLCAL and UVFND set the integration time to one day, causing bad amplitude calibration when there were rates and delays *2007-04-26*
10. aips.1 man page was lost *2007-05-04*
11. FILLM skipped a record at ends of file which could lose a data sample and, possibly cause confusion if the mode changed *2007-05-24*
12. SPLIT lost the calibration flags when a source was not found so that later data did not have calibration applied *2007-06-10*
13. SNPLT lost data from phase plots of PC tables due to failure to check for wraps and got hour angles wrong by 6 hours *2007-06-14*
14. FILLM used a blank in the middle of some station names, confusing other software packages *2007-06-16*
15. BLAVG scaled by the sum of the weights once too often, giving bad output amplitudes very dependent on whether the weights were calibrated *2007-09-05*
16. FILLM used 2 different defaults for a nominal sensitivity recorded as 0.0. This led to bad weights. *2007-09-28*
17. SPLAT set the wrong IF frequencies whenever channels were averaged to make a multi-channel output. SPLIT had a similar error, but only when the channel increments were different in different IFs. *2007-10-15*
18. SPLIT set the wrong frequency reference pixel when averaging channels and BCHAN was not 1 and messed up the FQ offsets when the IF increments were not the same. *2007-10-23*
19. SPLIT and SPLAT copied and computed on too many correlators, leading to buffer overruns and less than optimal scaling for compressed data. *2007-11-06*



December 31, 2007



AIPSLETTER

National Radio Astronomy Observatory
Post Office Box O
Socorro, NM 87801-0387
USA

FIRST CLASS

To:

Library
Natl Radio Astronomy Obs.
520 Edgemont Road
NRAO

DATED