A I P S L E T T E R

Volume XXVIII, Number 1: June 30, 2008

A newsletter for users of the NRAO
*A*stronomical *I*mage *P*rocessing *S*ystem

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.aips.nrao.edu/

General developments in *AIPS*

FILLM

31DEC06 contains a revision of FILLM which is essential to support the new data form that has been produced by the VLA since late June 2007. However, a bug affecting the scaling of cross-hand data in the new form was not corrected in 31DEC06. Therefore, VLA users will have to upgrade their copy of *AIPS* to 31DEC08 or 31DEC07 to read such data correctly. There are additional changes only in the 31DEC08 version of FILLM, which may 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 last release is called 31DEC07; 31DEC08 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 *AIPSLetter* is intended to advise you of improvements to date in 31DEC08. Having fetched 31DEC08, you may update your installation whenever you want by running the so-called "Midnight Job" (MNJ) which copies and compiles the code selectively based on the changes and compilations we have done. The MNJ will also update sites that have done a binary installation. 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.

Patch Distribution for 31DEC07

Important bug fixes and selected improvements in 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. 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. Since we now have many binary installations, the patch system has changed. We now actually patch the master version of 31DEC07, which means that a MNJ run on 31DEC07 after the patch will fetch the corrected code and/or binaries rather than failing. Also, installations of 31DEC07 after the patch date will contain the corrected code.

The 31DEC07 release has had a number of important patches:

1. REBYTE did not handle tables with long rows (IM and possibly BP) correctly *2008-01-09*
2. FITLD did not translate WX (weather) tables correctly *2008-01-18*
3. DFT model division did not set weights correctly *2008-03-05*
4. FILLM did not scale and weight cross-hand data for some baselines correctly *2008-03-05*
5. VISDFT did not do multi-scale model division and subtraction correctly *2008-04-29*
6. FILLM did not set the CORRCOEF keyword correctly for recent data *2008-06-19*

$\mathcal{AIPS}$ Distribution

We are now able to log apparent MNJ accesses and downloads of the tar balls. We count these by unique IP address. Since some systems assign the same computer 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 $\mathcal{AIPS}$ to a number of computers, so these numbers are probably an under-estimate of the number of computers running current versions of $\mathcal{AIPS}$. In 2008, there have been a total of 981 IP addresses so far that have accessed the NRAO cvs master. Each of these has at least installed $\mathcal{AIPS}$ and **214** appear to have run the MNJ on 31DEC08 at least occasionally. During 2008 more than 155 IP addresses have downloaded the frozen form of 31DEC07, while more than 611 IP addresses have downloaded 31DEC08. The binary version was accessed for installation or MNJs by 218 sites in 31DEC07 and 529 sites in 31DEC08. The attached figure shows the cumulative number of unique sites, cvs access sites, and binary and tar-ball download sites known to us as a function of week — so far — in 2008.

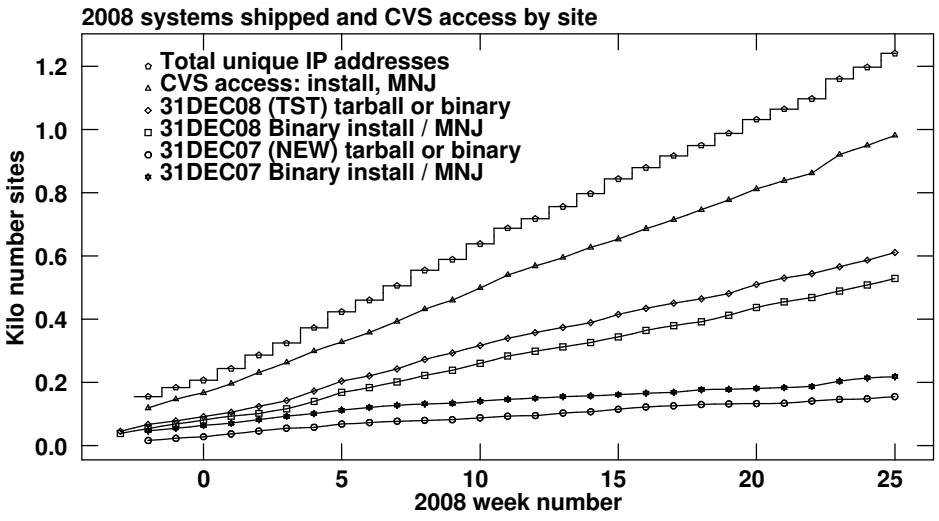

Improvements of interest in 31DEC08

We expect to continue publishing the *AIPSLetter* approximately every six months along with the annual releases. There have been quite a few changes in 31DEC08 in the last six months. A significant effort has been made to upgrade the capabilities of the "TV" display (XAS) and to make imaging and model computation more efficient through reduction in disk I/O. Although both of these should improve performance significantly for some users, neither causes any significant change in user inputs. A new verb SETMAXAP allows the user to guide AIPS in the matter of the amount of memory which it is safe to use within individual tasks. The model computation changes make this guidance significant. There are two new tasks to support on-line flagging tables which presently are written only for the VLA. These are PRTOF to print such tables and OFLAG to use these tables to generate entries in flag (FG) tables. A new task FIXAL and procedure FXALIAS have been written to deal with a temporary aliasing problem on EVLA-EVLA baselines. A new diagnosis task TIORD was written to check *uv* data sets to make sure that they are in strict time order, reporting any failures.

31DEC08 contains major changes to the display software. Older versions may use the 31DEC08 display (XAS), but 31DEC08 code may not use older versions of XAS. 31DEC04 through 31DEC08 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 five releases. Other than these issues, 31DEC08 is compatible in all major ways with the with the 15OCT98 and later releases. There are significant incompatibilities with older versions.

TV display

The AIPS "television" display dæmon XAS was significantly modified to allow greater numbers of image memories and wider dynamic ranges. This upgrade was done by the addition of new operation codes so that older versions should be able to use the new display program without modification. The number of TV memories, each the size of the display screen, was changed from 4 to 16. This change will allow for much larger TVROAMs up to 4 x 4 planes or even 16 x 1 or 1 x 16 planes and for larger TVMOVIEs using a combination of more spectral channels and larger sub-images of each channel. The AIPS TV allows the user to do very complex combinations of images; the larger number of memories will allow "layering" up to 16 simultaneous images in the display.

The range of data values in each memory has been changed from 0–255 to 0-8191. The data go through look-up tables called LUTs which previously had 256 input values and output values in the range 0-255. The new range is 8192 inputs and output values in the range 0-2046. These outputs are summed over those image memories which are "on" and enter the output function memory look-up tables (OFMs) which now have 32752 possible input values (2047×16). The output data range is still 0-255, which is all the display screen can actually handle. This extended dynamic range should allow for greater display flexibility after the image has been loaded to the TV memory and should also allow mathematical combinations of images, such as TVHUEINT, of greater accuracy.

Other than the ability to ask for TVCHAN up to 16, there are few changes visible at the user level. Tasks and verbs that use two graphics or grey-scale channels now use new adverbs GR2CHAN and TV2CHAN, respectively, to specify the second display. The verbs GRON, GROFF, TVON and TVOFF that used to take a decimal-coded immediate argument, *e.g.*, TVON 12 meant turn on channels 1 and 2, now take a binary-coded immediate argument. Thus, in the new system, TVON 12 turns on channels 3 and 4 ($12 = 4 + 8 = 2^{3-1} + 2^{4-1}$). Other verbs and tasks which allowed TVCHAN and GRCHAN to have decimal-coded multiple values now no longer support that option. Task TVCPS now displays whatever graphics channels are visible rather than requiring the user to specify which of the visible and invisible ones it should use. AIPS INPUTS and GO now know the number of TV memories and the size of the TV screen locally and use those limits when checking adverb values.

Imaging and model computation

When $\mathcal{AIPS}$ computed the visibilities from a source model, it used to have to re-read the *uv* data for every facet and, for frequency-dependent models, every spectral channel. The latter occur in tasks `IMAGR` and `OOSUB` which allow the user to correct models for the primary beam pattern and for known images of spectral index. All tasks involved with model computation deal with multiple facets, including calibration tasks such as `CALIB`, `FRING`, and so forth, as well as some modes in `IMAGR`. The basic routines that compute models in both DFT and gridded forms have been changed so that they can allocate a large "pseudo-AP" memory and compute the models for as many channels and facets as possible for each read through the *uv* data. This should greatly improve performance in large bandwidth-synthesis imaging problems involving wide fields and/or wide bandwidths. We will continue to look for more ways to reduce the disk traffic.

A new verb has been added to $\mathcal{AIPS}$, named `SETMAXAP`, to allow the user to specify the maximum amount of memory for an $\mathcal{AIPS}$ task to use. This allows users on small-memory machines, or machines doing many simultaneous operations, to limit the algorithms described above to reasonable memory sizes. If they are not limited, severe paging problems could cause tasks to take nearly infinite times to complete. Most $\mathcal{AIPS}$ tasks have algorithms that adapt to available memory, so a limited allocation will still work and should be faster than a blindly page-faulting version. We have to leave this parameter in the users' hands since we cannot find an operating system service to provide this information. Standard memory allocation will quite happily allow memory sizes in excess of available physical memory and will only fail if they exceed available swap space.

During the testing of the new gridded and DFT subtractions, a number of disturbing things were noticed. It became clear that the order in which facets were subtracted from the *uv* data mattered, especially in the less than totally accurate (but much faster) gridded subtraction. Steps were taken in the code to retain the original u, v, w values rather than to rotate them for one facet and then rotate them back before starting the next. In single-precision, numerical error could accumulate in this operation and subtle changes in cell position can change how the gridded algorithm does its interpolation for those few points that are almost exactly on cells. A better-known error is seen in imaging with large numbers of samples, either many visibilities or many spectral channels in bandwidth synthesis. The images, including the beams, frequently show large excursions in the corners. This is the result of a random-walk accumulation of numerical error in gridding followed by a very large correction (in the corners) for the Fourier transform of the gridding function.

We have experimented with making the pseudo-AP operate fully in double precision. We did not add options to send in double precision values, but just took advantage of the improved accuracy internally. Indeed, the problems described above were greatly reduced by this. There are two costs in switching the pseudo-AP to double precision. The first is that the number of data words available is halved, which reduces the number of facets and channels which can be handled simultaneously. This will add to the real time in large problems (where the errors are more important). The second is that the cpu and real times increased by a noticeable amount, probably due to double-precision arithmetic being a bit slower and due to cache limitations when twice the memory is required for an operation. It was concluded that one may make a few algorithmic changes to avoid the worst of the errors (*e.g.*, avoiding image corners, caching u, v, w values for re-use) and that the scientific results of the computations will not be enhanced by the greater accuracy. It is important to pay attention, however, to details. If a single-precision number is used to count a very large number of small numbers, then it will reach a maximum beyond which the small numbers will not contribute even if there are still very many of them. Double precision, or smarter algorithms, are the solutions in such cases.

UV data calibration and handling

FILLM

A significant error, mentioned above in the patches listing, was discovered in the task that translates the current VLA data format into $\mathcal{AIPS}$ *uv* files. `FILLM` reverses the direction of some baselines in the data to make them consistent when numbered with the actual antenna numbers; they are written by the VLA on-line system in a consistent fashion based on "dcs" address instead. Unfortunately, the nominal sensitivity (T_{sys}) correction factors were not swapped. During the ModComp era (prior to June 27, 2007), this affected only

Solar data and the weights of cross-hand data. The latter is likely to be of almost no significance, but Solar data should have encountered troubles calibrating polarization. The new, post-ModComp system writes the data as correlation coefficients, forcing FILLM to scale the data as well as the weights by the nominal sensitivity. This means that some baselines will have the cross-hand data incorrectly scaled for data taken after June 27, 2007 using versions of FILLM prior to March 6, 2008. Parallel-hand data were correctly scaled and weighted. Solar data since June 27, 2007 should now be correct.

The 31DEC08 version of FILLM also had a number of improvements. The option to average data on input was extended to include correct averaging of the output tables (TY, PO, OF, WX, and OT), to check integration time as well as elapsed time to terminate an average, and to restart an integration if the first time sample(s) had no valid data. A few more holes, in which FILLM could miss changes of mode or the number of spectral channels at file boundaries, were plugged. The task now writes an on-line flagging table in a new, more general and extensive, format. The task was changed to use an improvement in the on-line format which actually tells the reader which receiver was used. For older data, FILLM will split two IFs from the same receiver if they are at frequencies on opposite sides of one of the "official" band boundaries. Messages will now be issued when the task omits (or includes) pointing- and tilt-mode data. The writing of the header keyword CORRCOEF was corrected. It indicates whether the data are correlation coefficients (+1) or visibilities (−1) and is used by the TYAPL task.

Two new tasks were written to do useful things with the on-line flagging table written by FILLM. They are PRTOF which displays the table, interpreting the flag bit patterns into meaningful words, and OFLAG which may be used to apply selectively the information in the OF table. Note that both of these read only the new OF table. FILLM used to write an OF table once in a while, containing almost nothing of any use.

EVLA-EVLA spectral aliasing

At the current time, the VLA is being operated using some antennas with the old VLA electronics and some antennas with the upgraded EVLA electronics, all of which feed the old correlator. Due to the absence of some prohibitively expensive filters, data from below baseband is aliased into the observing band on EVLA-EVLA baselines only. For narrow-band, spectral-line observations, this causes a serious error in the band shape on those baselines only. Extensive study suggests that the form of the data on these baselines is

$$V(n) = A_c e^{2\pi i\phi} + f(n)\, A_c e^{-2\pi i\phi} + V_l(n)$$

where A_c is the amplitude of the continuum, ϕ is the uncalibrated phase of the continuum, V_l is any spectral-line signal as a function of channel number n, and $f(n)$ is the strength of the aliasing. It appears that, to first order anyway, $f(n)$ is independent of time, direction, IF, polarization, and antenna and is a real function with no phase term.

We have written a new task FIXAL which fits observations of calibrator sources to determine $f(n)$ and then fits that function to line-free channels in the main data set to determine A_c and ϕ to correct the data for the aliasing. A procedure FXALIAS was written to assist in the operation. It runs BPASS using only VLA-VLA and VLA-EVLA baselines, applies the bandpass to all data with SPLAT, separates the bandpass calibrators with UVCOP, and then runs FIXAL. Note that this operation must be done on totally uncalibrated data — if any phase correction has been applied, the above formula will have been rendered incorrect.

Users should note two things. This problem is temporary. When the new WIDAR correlator is used, the problem will disappear. However, the problem will remain with us until then and will remain in the VLA data archive forever. *At present, the new task and procedure should be regarded as experimental.* They appear to work most of the time and to remove most of the problem. There are niggling bits left and there seem to be isolated cases in which they do not work well.

Other *uv*-editing matters

EDITA and EDITR were changed to handle flagged table rows without dying, to "restore area" with the same complex logic used in "flag area", to know which antennas have data and ignore those that don't, to handle phase plot ranges better, to keep track of source name/number even when only one source is included, and to apply the FG table correctly to TY tables before plotting them.

SNFLG now counts the data to determine the required dynamic memory correctly and benefits from the changes to the EDIT class described for EDITA.

TVFLG now supports channel averaging with NCHAVG and CHINC adverbs and allows auto-correlations to have phase (which they do in cross-hands).

WIPER now has interactive options FLAG BASELIN and UNFLAG BASEL to eliminate/restore all points from a user-entered antenna pair. Up to 10 antenna pairs are remembered for each plotted point. Phase plots handle wraps better.

Other *uv*-display matters

LISTR now supports two gain conversions EFST (effective system temperature) and SEFD (system equivalent flux density) and honors the FREQID specification on gain listings.

VPLOT can now average spectral channels under control of adverb AVGCHAN and can plot channels and IFs separately or together under control of the CROWDED adverb. It can now draw connected lines in color when requested and selects better phase ranges for plotting when possible.

UVPLT was improved to plot phases with the least wrap problem, to implement the fixed scale within a fixed range option, to bin phases in a vector fashion rather than scalar, and to scale u, v, w by frequency.

Other *uv*-related matters

FITLD was corrected for an error that caused weather tables to be garbled and for an error in the logic that decided which antennas to include in the next record being written in the CL table. Code was added for the new software "DifX" correlator to read a new keyword and, based on it, to avoid one of the VLBA digital corrections, while making all of the others.

SPLIT and SPLAT now write a new index table automatically. The index table can help CALIB do a better job of averaging within scans even in single-source files and running INDXR is a nuisance.

DELZN was changed to use the antenna name rather than number since the latter can vary between data sets.

CLCOR was changed to offer a moving-source correction option using either fixed rates or an INFILE and to use the new format of the output from DELZN.

Weights in model division should be multiplied by the model amplitude squared. The gridded method was corrected some time ago, but the DFT method was only corrected in March.

TIORD is a modest task intended simply to report all points in a data set at which the data are not in strict time order. INDXR is very fussy about this, but quits at the first such point without supplying any useful information.

Analysis

RMSD was changed to offer the histogram-fitting method of rms determination, to allow control of the number of iterations in the robust method, and to allow circular as well as rectangular apertures for the computations.

CONVL was changed to offer the option of doing a cross-correlation of two images and to find and fit the maximum in the result.

IMEAN and IMSTAT can now do their thing inside or outside the specified window.

MFPRT was upgraded to offer a variety of new output options including full user control of which columns are displayed and more mnemonic column labels.

IRING was changed to offer more plot options including choice and size of symbol, error bars, connecting of points, and a rescale and relabel option for the x axis. This allows conversion to, for example, kpc rather than arcsec. It now also offers a text-file output option (for use with `PLOTR` or other plot programs) and uses a plot type fully understood by `EXTLIST` and `PLGET`.

ISPEC had displays of total flux, sum of plotted points, and number of non-blanked points added. Previously, one had to add up the numbers by hand.

XGAUS had a number of bugs corrected, one of which caused it to go catatonic when a retry was requested. Additional descriptions were added to the help file to clarify what it is attempting to do.

IMFIT and `JMFIT` reported the major and minor axis sizes in arc seconds after conversion from pixels, but reported the pixel-fit position angle rather than the one the corresponds to CCW from North in coordinate space. `SAD` reported this position angle correctly.

Miscellaneous matters

TABED had the operations of delete, clip, and unflag added.

PEELR now supports the `SOLMODE` option.

I/O count is now displayed in megabytes at the end of each task. If one enters `SETDEBUG 1`, then all following tasks will display separately the total count and size of `ZMIO` and `ZFIO` reads and writes from which the total is computed.

CookBook was reviewed thoroughly and numerous upgrades and corrections were made about January 1.

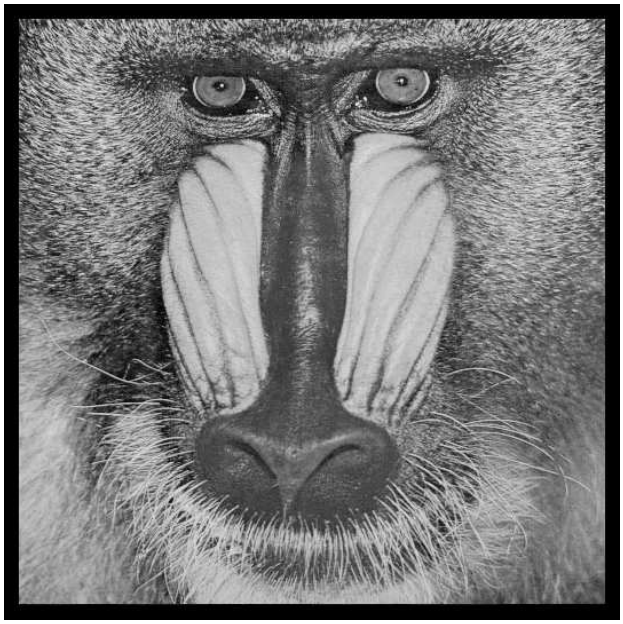

June 30, 2008

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