AIPSLETTER

Volume XXXII, Number 1: June 30, 2012

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.aips.nrao.edu/`

General developments in *AIPS*

Reduction of EVLA and ALMA data in *AIPS*

This *AIPSLetter* and those from 2010 and 2011 contain numerous improvements to *AIPS* that enable full calibration of EVLA data and most imaging operations as well. The one exception is the wide-band (bandwidth synthesis) deconvolution algorithm ("MSMFS") being developed in CASA by Urvashi Rao Venkata, for which there is no comparable function in *AIPS*. Calibrated *uv* data may be ported from *AIPS* in "UVFITS" format for use in that program. ALMA data may also be reduced in *AIPS*, although the package is not fully qualified to calibrate data from linearly-polarized feeds. See Appendix E of the *AIPS* Cookbook, available via the *AIPS* web site, for details.

AIPSLetter publication

We have discontinued paper copies of the *AIPSLetter* other than for libraries and NRAO staff. The *AIPSLetter* will be available in PostScript and pdf forms as always from the web site listed above. It will be announced in the NRAO e-News mailing and on the bananas list server.

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 (32- and 64-bit), Solaris, and MacIntosh OS/X (PPC and Intel) systems may install binary versions of recent releases. The last, frozen release is called `31DEC11` while `31DEC12` remains under active development. You may fetch and install a copy of these versions 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 `31DEC12`. Having fetched `31DEC12`, 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. Use a copy of the task and its help file in a private disk area instead.

Improvements of interest in `31DEC12`

We expect to continue publishing the *AIPSLetter* approximately every six months along with the annual releases. Henceforth, this publication will be primarily electronic. There have been several significant changes in `31DEC12` in the last six months. Some of these were in the nature of bug fixes which were applied to `31DEC11` before and after it was frozen. If you are running `31DEC11`, be sure that it is up to date to sometime this year (a MNJ after March 8 would be best). New tasks in `31DEC12` include `MORIF` to break up a data set into a greater number of spectral windows, `SNREF` to determine which reference antenna would minimize the right minus left phase difference, `PRTSY` to print statistics of the values in SysPower (`SY`) tables from the EVLA, `FIXAN` to correct errors or change the array center in antenna tables, `SPCOR` to correct an image cube for spectral index and/or primary beam, `FIXRL` to fix data for mislabeled polarizations in some antennas, and `TARS` to check the Faraday Rotation Synthesis algorithm (task `FARS`) with simple user input data. New verbs include `NAMEGET` to fill in the file naming adverbs completely to aid in writing procedures and `IM2HEAD`, `IM3HEAD`, `IM4HEAD`, `IMOHEAD`, `Q2HEADER`, `Q3HEADER`, `Q4HEADER`, and `QOHEADER` to display headers selected by the second, third, and fourth input and the output name adverbs, respectively. New `RUN` files include `OOCAL` to enable full self-calibration with the spectral-index and other options of the `OOSUB` task and `LINIMAGE` to build a `FLATN`ed spectral cube while separating spectral windows to improve performance in the `IMAGR` portion.

`31DEC09` contains a significant change in the format of the antenna files, which will cause older releases to do wrong things to data touched by `31DEC09` and later releases. `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`. Magnetic tape logical unit numbers changed with `31DEC04`. You are encouraged to use a relatively recent version of *AIPS*, whilst those with EVLA data to reduce should get the latest release.

Mountain lion

Apple has again announced a new revision of OS X, this time called Mountain Lion (10.8). The VAO in Socorro has made a Mac laptop available to us for testing. It was found to install and run the existing `MACINT` load modules. Several `$SYSUNIX` procedures required minor modification since the new version of `ps` issues a warning message when it is used by a process with a library path set. On lion last year we tried using the latest Intel compiler (12.0.4 dated 20110503) with the x86_64 architecture. That version ran a little bit faster than the 32-bit version, but not enough faster to justify making a new `MAC64` "architecture" with a new NRAO computer to support it.

UV-data calibration and handling

Editing

`RFLAG` has become the standard editing task to remove RFI from most data sets. It has received significant attention in the last six months. The spectral flagging mode was changed to use a sliding median window filter of user-set width. Previously it used, effectively, a full-width window applying to all channels. This caused the disappearance of options to flag full spectra because the rms was too high or too low. An option to expand any flags by ± `FPARM(6)` channels replaced them. In the spectral mode, `RFLAG` automatically deleted all channels not included in ICHANSEL. Now such channels are retained, but the default ICHANSEL is now all channels. A number of errors were corrected which caused the task to write many more flags than it should, effectively flagging already flagged data. Other errors had to be corrected for the mode in which RFLAG is run to write a new flag table and then immediately plot the consequences of that table on the data set. The problems involved pre-clip flux levels as well as file status.

Following detailed self-calibration, users often delete those data associated with "bad" amplitude and phase solutions. An option A&P has been added to SNFLG to perform this operation, removing the need for a very repetitive use of task EDITA. TVFLG and SPFLG were corrected to retain and use the source number correctly when flagging a single source extracted from a multi-source data set.

Calibration with models

The task OOSUB will subtract a model from, or divide a model into, a *uv*-data set. It has options to apply the model in a frequency-dependent fashion using a model of the primary beam and images of spectral index and curvature. It also has options to omit some of the model Clean components from the computation depending on their position within the primary beam. A new procedure has been written to use OOSUB to do self-calibration. The procedure divides the model into the data with OOSUB and then invokes CALIB to determine the gains. The resulting SN table is copied back to the input data set and the temporary, divided data set is deleted (if requested).

A number of detailed aspects of applying models in tasks like UVSUB and CALIB were also addressed in the last six months. A user may wish to apply some algorithm to the CC tables of some of the facets of an image, creating CC tables of version numbers greater than one. The software will now allow the use of the highest CC table version in each facet if INVERS = 0 or to use a specific version number, ignoring those facets lacking that version number, when INVERS = n. Previously, the software could get confused and either die or change the modeling to use the Cleaned images rather than Clean components. Neither was desirable. When using edited CC tables, gridded model would get stuck in the table and only use those components that preceded the first flagged component. Models can now come from small images or ones which are not a power of two on a side. The former require DFT rather than gridded modeling, but a bug kept the code from knowing about image size when it had to make that decision. The latter often perform well with gridded modeling, but the sizes of the grids had to be increased to the next power of two for correct function.

Miscellaneous *uv*-data matters

BPASS no longer supports (*i.e.*, messes up) the BIF and EIF options.

SETJY now can determine the source velocity at the first data sample and write that into the source table with OPTYPE='VCAL'.

VLANT now has a DOINVERS option to allow corrections to the CL and AN tables to be undone. While a new CL table produced by VLANT can simply be discarded, the changes to the antenna coordinates are made in place so this option is required to undo them. Errors for the EVLA in the sign of the Y correction and the zero point of hour angle were found and fixed.

FRING and RLDLY can now solve for delays dividing the IFs into quarters to support the new 3-bit EVLA system.

CLCOR neglected the fact that the GMRT uses the opposite phase sign convention from other telescopes. This was corrected for a number of OPCODEs, most importantly the one to apply corrections for changes in antenna and source positions.

FITLD had the order of the adverbs revised to clarify which of the many adverbs apply only to VLBI IDI-format data sets and which are more general.

FIXRL is a new task to rearrange data on the polarization axis when some antennas have been connected incorrectly (swapping R and L usually).

INDXR now has a frequency-dependent VLA gains file which it applies to the CL file in a frequency-dependent manner.

ELINT now displays the rms of the fit on plots. Adverbs QUAL and CALCODE were added to help separate out the desired calibration sources. Problems reading over-the-top tables and large source tables were corrected and the task was forced to ignore flagged solutions.

SNREF is a new task to examine the statistics of the right minus left phases in an SN table as a function of the choice of reference antenna. It should help to select the antenna having the most stable R-L and hence providing the best polarization calibration.

MORIF is new task which can divide a data set into a larger number of spectral windows (IFs) which may help deal with spectral index issues at very low frequency and may help performance in data sets with a large number of spectral channels (see LINIMAGE below).

PRTSY is a new task to print the extrema and statistics of the EVLA SysPower (SY) table as a function of antenna and IF. It should help provide parameters for use in TYSMO and other tasks.

FIXAN is a new and dangerous task that re-writes the coordinates of antenna tables. It has OPTYPEs to switch between Earth-centric and site-centric systems correctly and to correct an apparent site-centric system computed by simply differencing Earth-centric coordinates. This last is wrong since site-centric systems must have the antenna (X, Y) based at the local longitude.

Imaging

AIPS imaging task IMAGR received only a little attention in the last six months. The biggest change was to a double-precision pseudo array processor; see below. This will result in more accurate images, especially near the corners where round-off error was magnified by the large correction for the Fourier transform of the gridding function. The uniform weighting routine was also worked on to insure that all summing variables are double precision floats or integers. Large EVLA bandwidth synthesis data sets were overflowing the counters in uniform weighting. An option was added to average spectral channels as the data are read in the on-the-fly, baseline-length dependent time averaging done in writing the work file. This will reduce the size of the work file and, in many cases, will not compromise the accuracy of the images.

When IMAGR is used to make a spectral-line cube, the calibration routines are started up several times per channel both for the input data which uses a flag table and for the work file which does not. These calibration start-up routines used to comment on the use of the flag table on every twentieth call. This led to confusion since it would say sometimes "using flag version n" and other times "using no flag table." The low level routines were changed so that flag table use is reported once after each change of the file being read unless that file is a scratch or work file for which FLAGVER was never set. For cubes, this means that the flag file usage is mentioned at the start of each channel. If the flag table version changes, as in RFLAG, then that change is also reported.

Tasks SCMAP and SCIMG loop over self-calibration and imaging (and even data editing) within a single task. The were improved by adding BOXFILE to define Clean boxes, IM2PARM to allow auto-boxing, and a parameter to allow negative components to be used in the model and restart of the Clean. Flag tables were ignored for input data which were not compressed, while they were used for compressed data. This oversight was corrected.

A new RUN file and procedure called LINIMAGE was written to help speed the creation of spectral image cubes. It loops over spectral window (IF) copying the IF to a scratch file and then invokes IMAGR on that rather smaller file. This is much faster than extracting each channel one at a time from the full file. After IMAGR, the facets are FLATNed. When all IFs have been imaged, the headers are examined and either MCUBE or FQUBE is run to put all the IFs back together. A data set with a very large number of spectral channels may be run through MORIF to prepare it for LINIMAGE.

Display

LISTR was changed to do system temperature and gain as determined from a SysPower table in the 'GAIN' listing. Adverb XINC was added to the 'LIST' and 'GAIN' listings as well as an adverb to change all times to local sidereal time rather than IAT.

POSSM was changed to read the BP table to find a list of times present in the file to allow it to use SOLINT properly when plotting BP tables.

IM2HEAD is a new verb to display the header like IMHEAD, but for the file pointed to by the adverbs IN2NAME, IN2CLASS, IN2SEQ, and IN2DISK. IM3HEAD, IM4HEAD, and IMOHEAD are the same but for the third and fourth input name and the output name adverbs.

Q2HEADER is a new verb to display the header like QHEADER, but for the file pointed to by the adverbs IN2NAME, IN2CLASS, IN2SEQ, and IN2DISK. Q3HEADER, Q4HEADER, and QOHEADER are the same but for the third and fourth input name and the output name adverbs.

Slice plots along a FQID axis can usually be plotted on a linear frequency axis by placing the plot points appropriately. ISPEC and RSPEC were changed to do this. Alternatively, the slice can be interpolated to regular frequencies as it is computed. SLICE, ISPEC, RSPEC, and BLSUM were changed to do this. The slice file format had to be changed slightly to provide the necessary information for subsequent plotting.

IMEAN was given a DOPRINT option to control printing and particularly to write a compact format listing the window, minimum, maximum, mean, rms, and flux. Using this option, looping over a plane at a time with BLC and TRC, would give a useful text file for a spectral cube.

CCNTR was changed to allow limits on the flux of plotted components, to allow plot symbols to scale with component flux, and to ignore flagged rows in the CC table rather than dying.

TV code was changed to remember if the TV has been opened or not. Previously, failing plot tasks could hang waiting for plot finishing on a TV display that was never going to respond.

Image and *uv* analysis

Faraday rotation synthesis

AIPS' Faraday rotation synthesis package has continued to be tested by committed users and corrected by the *AIPS* group. The main FARS task was changed to convert from real/imaginary to amplitude/phase correctly, which was a particular problem when residual images were output. The Clean code was changed to do no iterations at some pixels due to the value of FLUX and to output zero in those cases. An option to shift the phases to the original λ^2 set was added. The handling of APARM(4) was clarified and its display in the history file corrected. The DOFARS procedure was changed to use more appropriate default output classes.

A new task, named TARS, was created to test the algorithms used by FARS. It is nearly identical to FARS but takes its input data from a text file which may be prepared more easily by people wishing to test this new algorithm.

The output of FARS may be compromised if the input Q and U polarization cubes are affected by frequency-dependent amplitude changes due to the source spectral index and to the primary beam of the interferometer. The new task SPCOR was written to correct the Q and U cubes for these effects. One may use SPIXR on a I cube to determine images of spectral index that may be appropriate to the Q and U images. The I spectral index image output of MSMFS in CASA may also be used.

AFARS is the task which reads the FARS output images to find the position on the rotation measure axis of the maximum amplitude and then write images of the rotation measure and of the amplitude or phase at that position. It was corrected to account properly for BLC, to avoid the parabolic fit estimation when it did not have all required values due to blanking and edge effects, to use the header coordinates properly rather than guesses which could be wrong, and to output the blanking value when the input rotation measure row was all blank or zero. The help file was improved to clarify which files are required as inputs.

UV-plane model fitting

The monster task `OMFIT`, written by Ketan Desai many years ago, received some much overdue attention. The task contains some very high quality mathematical routines which implement a wide variety of source model types and a joint model-fit and self-calibration. It is a bit difficult to use, but once convergent models are found, it can be very useful. A few rather simple corrections were required to get it to work at all. The input model is read by a `KEYIN`-based routine which had to be corrected to handle parameters longer than 8 characters. The task was corrected so that it would work on single-source input files. Michael Bietenholz provided suggested code changes to add a `WEIGHTIT` parameter adjusting the data weights, to normalize phase-only self-cals correctly, to display "real" chi-squared values, and to include a spherical shell model type. Commentaries on how to use `OMFIT` were also provided by Michael for inclusion in the explain file.

The other *uv*-plane model fitting task `UVFIT` also received some attention. It was changed to allow up to 60 components in the model. All of the code in the task now supports this parametric upper limit after the last of the hard-coded 4's were rooted out. The computation of phase now uses the proper geometric routines which are slower but very much more accurate for components well separated from the phase center. The input text file describing the model may now include a comment string which is included in the output summary text file.

Other analysis changes

COMB was corrected to adjust the units of each of the input images so that units of Jy/beam become units of Jy per the same beam. Previously two images in Jy per different beam were added, subtracted, divided, etc. producing outputs which may have had more to do with the differences in beam than in source.

DTSIM is a *uv*-data simulation task which was corrected to use right-handed antenna geometries, to use correct and current-best antenna locations for known antennas (VLBA and EVLA), to use sensible and many more user-controlled defaults, to create an NX table, and to write the data in the standard axis order. The help file was overhauled to tell the truth about things including whole areas previously omitted. A useful example replaced an overly simple one. This task may actually be useful.

SAD now uses a robust determination of the rms to inform the histogram fit of the rms, producing a much more reliable estimate of the overall uncertainty.

General

Double-precision pseudo array processor

When a four-byte integer is used to count, it will overflow after $2^{31} - 1$ or 2147483647. When a four-byte, IEEE-format real variable is used to sum 1.0's, it will stop increasing after 16777216. With the EVLA, bandwidth synthesis easily exceeds the latter limit and can exceed the former. AIPS' pseudo-array processor ("AP") code is a software emulation of the old hardware array processors we used to own. It isolates computationally intensive operations in a modest collection of "Q" routines based on a shared data space. This structure allowed us to change the entire pseudo-AP from single-precision to double-precision real and integer for the shared data and, with some editing, to complete the change so that all variables inside the AP are extra precision. Variables received from and sent to the non-Q calling routines are normally still single-precision, although new routines `QDGET` and `QDPUT` were written to allow full-precision data transfers.

This has a number of consequences. The most important is that the number of words available in the pseudo-AP memory is halved, or the memory in bytes required is doubled if the number of words is kept constant. This required some changes in the verb `SETMAXAP`. Tests a couple of years ago suggested that performance of a double-precision AP was a about 10% worse than a single-precision version, presumably due to memory cache issues as much as anything. The computed results, however, are more accurate even when double-precision was not needed to avoid overflow. Convolving functions and their Fourier transforms are computed in the AP code and so will be much more accurate in gridding *uv* data and then correcting in the image plane for that gridding.

Other general-interest changes

Lustre is a high-performance file system used at NRAO and elsewhere. Such systems do not keep information about file sizes with the files themselves and so can mis-inform *AIPS* about the file size. Low-level routines were changed to pause and try again twice before dying when it appears that a task wants to read or write beyond the end of file.

Source table format was updated to include columns for the observed (or pointing) right ascension and declination. The current RA and Dec columns are used for the phase stopping position, which may be quite different.

Frequency table format was changed to include a column for `BANDCODE`, an 8-character string meant to identify which receiver was used for each IF for each frequency identifier.

ZEXIT.C is a new routine to allow a non-zero exit code to be returned to the calling procedure. It is useful in stand-alone programs, especially the *AIPS* pre-processor program `PP`. It replaces the non-standard Fortran system routine `EXIT`.

NAMEGET is a new verb which may be useful in writing procedures. It uses the first image name parameters (`INNAME`, *et al.*) to find the image or data set that an *AIPS* task would find with those name parameters and to fill in the remaining ones so that there are no defaulted or wild-carded ones.

Patch Distribution for `31DEC11`

Important bug fixes and selected improvements in `31DEC11` 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 sub-directories below this. As bugs in `31DEC12` are found, they are simply corrected since `31DEC12` 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 `31DEC11`, which means that a MNJ run on `31DEC11` after the patch will fetch the corrected code and/or binaries rather than failing. Also, installations of `31DEC11` after the patch date will contain the corrected code.

The `31DEC11` release has had a number of important patches:

1. Imaging code had problems with large gridded modeling. *2012-01-16*
2. `UVFIT` had many areas dimensioned for the old limit of 4 components, rather than the new limit of 20. *2012-01-16*
3. `FARS` did not shift residuals back to the original λ^2 correctly and did not add residuals into amplitude/phase outputs correctly. *2012-01-16*
4. `AFARS` did not handle `BLC` correctly, producing wrong rotation measure images. It also stated incorrect information in the history. *2012-01-16*
5. `SQASH` had an error in history writing that led to aborts. *2012-01-16*
6. `FITAB` omitted an essential keyword for compressed UV data. *2012-01-16*
7. `CLCOR` ignored the fact that the GMRT uses an opposite phase convention. *2012-01-16*
8. `VLANT` did not correct EVLA data properly. *2012-03-01*
9. `SCMAP` and `SCIMG` did not apply the flag table to the data entering the self-cal step *2012-03-07*

Recent *AIPS* and related Memoranda

All *AIPS* Memoranda are available from the *AIPS* home page. There is one new memoranda in the last six months.

117 *AIPS* **FITS File Format**
Eric W. Greisen, NRAO
June 28 2012 (latest revision)
AIPS has been writing images and *uv* data in FITS-format files for a very long time. While these files have been used widely in the community, there is a perception that a detailed document in still required. This memo is an attempt to meet that perception. *AIPS* FITS files for *uv* are conventions layered upon the standard FITS format to assist in the interchange of data recorded by interferometric telescopes, particularly by radio telescopes such as the EVLA and VLBA. An Appendix providing a serious introduction to the subject of FITS format is also provided.

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 *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 2012, there have been a total of 837 IP addresses so far that have accessed the NRAO cvs master. Each of these has at least installed *AIPS* and 281 appear to have run the MNJ on `31DEC12` at least occasionally. During 2012 more than 194 IP addresses have downloaded the frozen form of `31DEC11`, while more than 596 IP addresses have downloaded `31DEC12`. The binary version was accessed for installation or MNJs by 303 sites in `31DEC11` and 600 sites in `31DEC12`. All of these numbers are significantly lower than those of 2010 at a comparable date. The attached figure shows the a comparison of the numbers from all of 2011 and from 2012 for the first half of the year. From bottom to top the numbers plotted are the cumulative unique IP address to download the `NEW` version, to download the `TST` version, to access the `cvs` web site, and the total unique addresses. It is strange that the downloads in 2012 are if anything slightly ahead of 2011, while the cvs access and hence the total numbers in 2012 are way behind those in 2011. It makes one suspect an as yet unidentified change in the way the `cvs` numbers are counted.

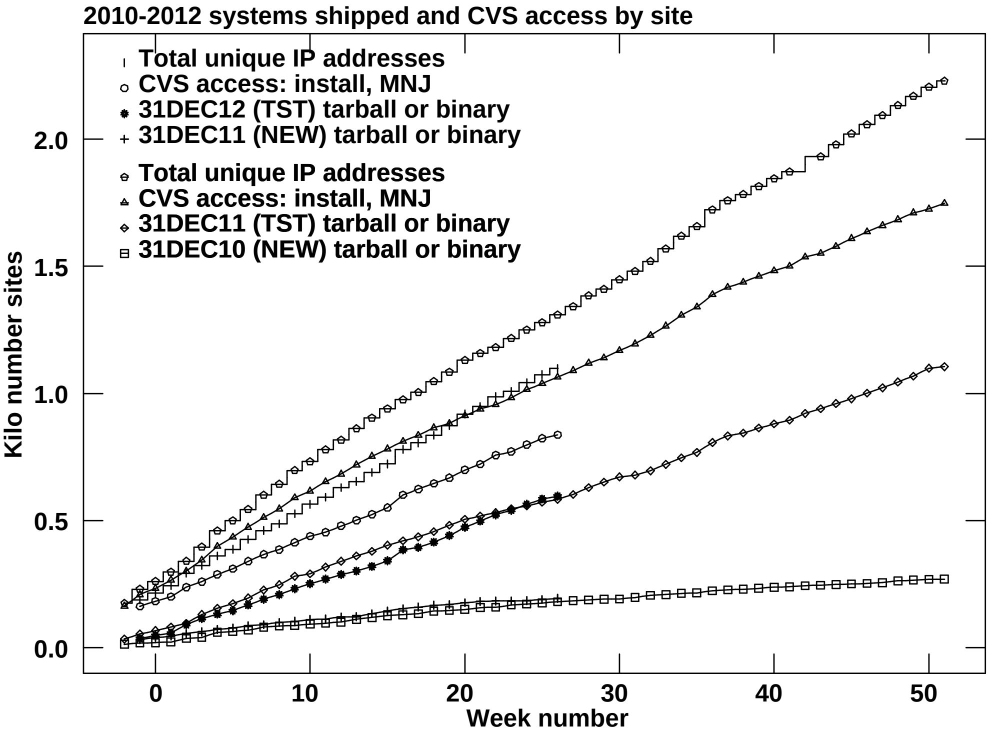

A personal note

The rumor that your editor is about to retire is not true. While I am certainly old enough, I am having too much fun and believe that I am still too useful to retire any time soon. *AIPS* will continue to be supported for the forseeable future.

Warning

Recent *uv* FITS files written by `CASA` and containing ALMA data have had significant errors in the contents of the antenna and frequency tables and the frequency information in the main header. The antenna error will primarily affect the use of `UVFIX` to recompute projected baselines. The frequency error will certainly affect plotted spectra (reversed axis labeling) and will have a small effect on computed images. These errors are known to be in release 3.4.0; FITS files written by earlier releases should be checked closely. `CASA`'s `exportuvfits` program is now being worked upon and these issues are expected to be missing from the next release. These errors do not appear to arise with EVLA data as written, via `CASA`, from the archive.

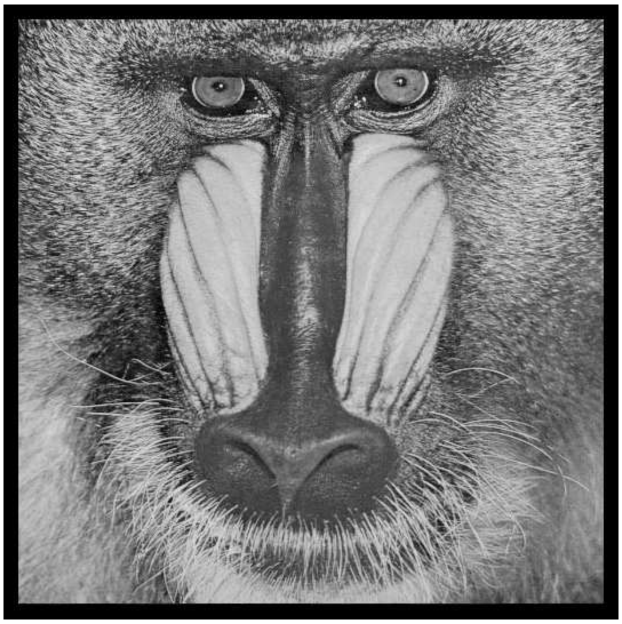

June 30, 2012

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