North American ALMA Science Center

ALMA Development Project

Pulsars, Magnetars, and Transients with Phased ALMA

Final Report
2017 October

PRINCIPAL INVESTIGATOR: JAMES CORDES

INSTITUTION: CORNELL UNIVERSITY

ADDRESS: 520 SPACE SCIENCES BUILDING, ITHACA, NY 14853

PI CONTACT INFORMATION:

Telephone Number 1 607 255 0608

Email address jmc33@cornell.edu

CO-INVESTIGATORS AND COLLABORATING INSTITUTIONS

Name	Institution	Email
Lindy Blackburn	SAO	lblackburn@cfa.harvard.edu
Shami Chatterjee	Cornell University	shami@astro.cornell.edu
Geoffrey Crew	Haystack Observatory	gcrew@haystack.mit.edu
Gregory Devignes	MPIfR	gdesvignes.astro@gmail.com
Shep Doeleman	SAO	sdoeleman@cfa.harvard.edu
Michael Kramer	MPIfR	mkramer@mpifr-bonn.mpg.de
Joe Lazio	JPL	Joseph.Lazio@jpl.nasa.gov
Kuo Liu	MIPfR	kliu.psr@gmail.com
Scott Ransom	NRAO	sransom@nrao.edu

SUMMARY

The present study developed fast time-domain capability for the ALMA phased-array system that is needed for observations of compact objects in the Galactic center and elsewhere in the Galaxy. ALMA can provide unparalleled sensitivity to a spectral region that has been poorly explored for neutron stars. Observations at mm and sub-mm wavelengths have the potential for providing decisive observational constraints on emission processes from the magnetospheres of neutron stars. ALMA is also important for surveys for pulsars and transients in the Galactic center. An existence proof for mm emission is in the detection of a magnetar (a neutron star with a 10^{14} Gauss field) in the Galactic center at frequencies up to 295 GHz. Pulsars are usually known for their high-energy emission in gamma and X-rays and for their bright, coherent radiation at centimeter wavelengths. Other neutron stars will also be accessible at mm wavelengths if, as implied by the magnetar's detection, there is an upturn in their spectra due to (for example) inverse-Compton scattering of the coherent, cm-wavelength photons. ALMA can play a key role in probing neutron star populations in the Galactic center and using detected objects for the study of spacetime around the central black hole in a panchromatic program.

ALMA observations of pulsars and transient emission from compact objects require high time resolution (sub-millisecond) data from phased ALMA. The development project involved a series of steps needed to acquire such data: definition and implementation of the appropriate phasing mode, providing the signal path from the APP phasing system to Mark 6 recorders, offline resampling and reformatting of data into PSRFITS format, and development of the backend pulsar/transient processing customized to ALMA science contexts. The project leveraged software development for pulsar phased-array modes for the VLA and developments for the Event Horizon Telescope and Black Hole Camera projects.

Feasibility studies conducted under the project used test data obtained in conjunction with ALMA Phasing Project commissioning runs in April 2016 and January 2017. The former provided data used to test the integrity of the Mark 6 to PSRFITS transformation while the latter provided data on the Vela pulsar (J0835-4510), a bright pulsar at centimeter wavelengths. Detection of the pulsar at a significance consistent with the sensitivity (A/T) and bandwidth used is the primary demonstration of the feasibility of pulsar and transient observations with ALMA. It provides the basis for future execution of our primary scientific motivation: searches and follow-up observations of pulsars and transient sources in the Galactic Center. ***We consider the ALMA detection of the Vela pulsar to be an explicit demonstration of the project's success.***

The work was done under a collaboration between Cornell University, the Harvard Smithsonian Astrophysical Observatory, Haystack Observatory, JPL, NRAO, and the Max Planck Institute for Radio Astronomy. The deliverables include a defined APP/Pulsar Operations Flow, a report on the technical development and results of feasibility studies, and a plan for future developments.

Code descriptions and results are available at http://astro.cornell.edu/research/almapsr/

CONTENTS

I. SCIENTIFIC MOTIVATION

ALMA is the premier imaging instrument for high radio frequencies. With a successful completion of the ALMA Phasing Project (APP), the telescope now provides a beamformed output that expands the scientific capabilities of the instrument that enable significant science on compact objects. Apart from enabling mm-VLBI observations, the APP also allows us to use ALMA, independently of other observatories, for time-domain observations of compact objects, in particular pulsars. Despite limited understanding of their emission mechanism, pulsars are proven tools for physical and astrophysical questions, such as studies of the interstellar medium or tests of theories of gravity. Observations with a phased ALMA can play a decisive role in many areas. Firstly, ALMA may allow us to discover and time pulsars in the vicinity of and orbiting about Sgr A*, the super-massive black hole in the Galactic center. Secondly, mm-frequencies may reveal information about the inner workings of the pulsar clock, with potential implications also for the usage of pulsar clocks in other areas. We concentrate here mostly on the use of pulsars as probes for the Galactic Center environment and for the description of black holes in general relativity and other theories of gravity.

Radio and infrared observations have clearly established the presence of a super-massive black hole (BH) in the Galactic center. Following the motion of IR stars orbiting Sgr A* has allowed the determination of the mass of the BH to about 4×10^6 M_{Sun}. Efforts are underway to improve the resolution and astrometric precision of these measurements using optical interferometry, potentially enabling 10 μas astrometry and probing the conditions at a 10^7 km level. In contrast, finding and timing an orbiting pulsar with a modest timing precision of one millisecond will probe physics in the GC at a level of 10^2 km. Some of us have shown (Liu et al. 2012, ApJ, 747,1) that a pulsar in a compact, few-month long orbit will allow measurement of the mass of the BH to a precision of 10^{-6}, the spin of the BH to 10^{-3} or better, and the quadrupole moment to potentially 10^{-2}. This will allow tests of the Cosmic Censorship Conjecture for BHs and tests of the No-Hair theorem. Thereby, these efforts not only share the technological goals of the APP but they are also closely linked scientifically. With the Event Horizon Telescope (EHT) project being the prime science driver for the APP, an APP-enabled pulsar discovery would provide important independent information for determination of the magnitude of the spin and also its 3D direction. This pre-determines the orientation of the accretion disk that needs to be modeled in order to understand the EHT image, promising unprecedented tests of general relativity. In other words, the far-field effects of the BH measured by one or more pulsars will strongly complement the scientific returns from the near-field information provided by EHT imagery.

Extensive efforts have been made by the authors to discover pulsars in the GC using the world's largest telescopes. These include large single dishes (GBT, Effelsberg, Parkes) and also, importantly, arrays, i.e. VLA. Despite the recent success of discovering and monitoring a magnetar (J1745-2900) only about 0.1pc from SgrA*, the large expected population of pulsars has not been found. Nonetheless, based on a large number of independent observational results, some us have concluded that the detectable population of pulsars is large (Wharton et al. 2012, ApJ, 753, 108). Scattering of radio waves in the turbulent GC medium has been suggested as the culprit for the non-detection, but observations of the magnetar reveal a surprisingly small amount of scattering. On-going monitoring of J1745-2900 opens the possibility that the scattering is variable and that a much denser scattering screen may indeed exist closer to the central BH. The best strategy is to combine searches at various frequencies, including very high frequencies, such as Band 3 of ALMA. Only at that frequency may it be possible to detect the short-period pulsars, which can provide timing precision needed for Galactic center studies.

While pulsars have been mostly studied at low frequencies, some of us have observed the magnetar successfully at frequencies as high as 295 GHz and normal pulsars up to 135 GHz. The brightness of the magnetar at ALMA frequencies is unusual for neutron stars. Magnetars are typically transient radio emitters, and it is not clear what triggers the radio outburst or how long radio emission lasts. The example of J1745-2900 therefore demonstrates the possibility that additional transient sources in the GC may be studied with ALMA. Investigation of those additional sources along with the magnetar will provide insight into neutron star emission in general, and the GC environment in particular, by a direct comparison with lower frequency observations. In addition to phasing up ALMA on Sgr A* itself for pulsar searches also required an observing mode that phase-referenced on Sgr A* and phase transfer to the pulsar/magnetar position. Our ALMA Development Project provides the basis for making phased ALMA a general purpose pulsar instrument, allowing the study of a variety of individual pulsars, for instance also for astrometric purposes, with an expected impact on the broad astronomy community. The project is consistent with the larger science case for phased ALMA described in the community document (Fish et al. 2013, arXiv:1309.3519).[1]

II. REQUIREMENTS FOR PULSAR AND FAST-TRANSIENT OBSERVATIONS WITH ALMA

Pulsars and fast transients involve time scales ranging from sub-millisecond to tens of seconds and longer, which reflect their temporal widths as well as their periods (in the case of pulsars). Long data sets of hours are likely needed at ALMA wavelengths because most pulsars are very weak. A notable difference between ALMA wavelengths and centimeter and longer wavelengths is that dispersive propagation is negligible for most cases and the same is true for multipath propagation. Consequently only coarse channelization is needed for removal of any radio frequency interference and to provide redundant detection capabilities.

These basic considerations translate into the following requirements: being able to phase ALMA for sustained periods of time, and acquiring baseband data that can be used to form the necessary data products with coarse channelization and of order millisecond time resolution for longer-period pulsars and sub-millisecond resolution for shorter-period objects. All ALMA bands are currently of interest given that the Galactic center magnetar appears to have a very shallow spectrum and it is possible that ordinary pulsars have spectra that turn up at short wavelengths.

For facility usage, a grid of suitable phase calibrators is needed along with appropriate scripts for acquiring phased array data on VLBI recorders. Software now exists for producing pulsar/transient data products in PSRFITS format and several software packages exist for searching such data for periodic signals and for individual bursts.

III. END-TO-END DEVELOPMENT PLAN

The ADP aimed to acquire test data, demonstrate data integrity of PSRFITS files, analyze PSRFITS files to characterize noise statistics, and to detect a known pulsar using standard methods. The project has produced software modules that start with VLBI data taken in phased-array mode, form PSRFITS files with averaging times and channelization appropriate for pulsar and transient studies, and analyzed the data. This report gives analysis results that demonstrate the achievement of our initial goals.

[1] URL: http://adsabs.harvard.edu/abs/2013arXiv1309.3519F

The Phased ALMA Pulsar Modes study defines pulsar phasing and data modes using the APP and provides software for resampling and reformatting into PSRFITS format for use in pulsar and transient search and timing pipelines. It builds upon the APP system as well as technical developments that provide pulsar data with the phased VLA. The ADP used data from the APP recorded in VLBI Interchange Format (VDIF) and produced PSRFITS files that were processed with pulsar and transient pipelines developed by some of the proposers.
The signal flow is shown in Figure 1. The pipelines require modest but necessary accommodation of the nearly negligible interstellar propagation effects at mm wavelengths in comparison with substantial distortions at cm wavelengths. The goal was to acquire and characterize the noise budget of reformatted APP output data and investigate detection of astronomical signals.

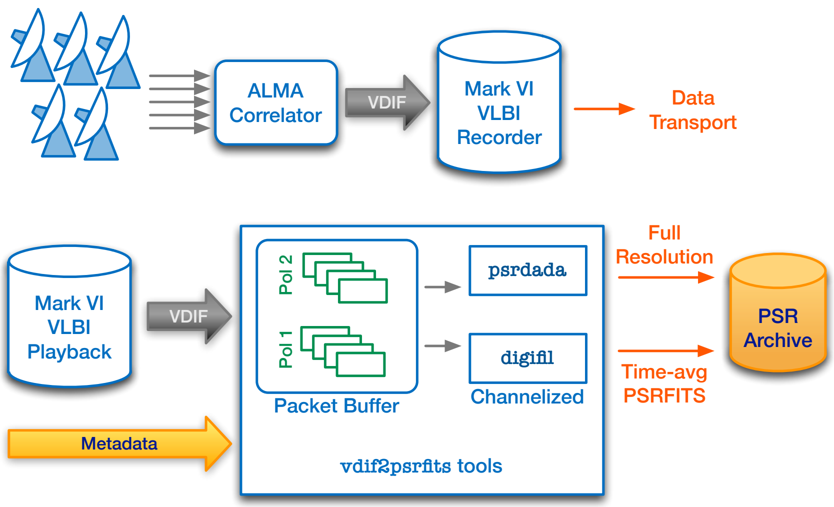

Figure 1. The developed ALMA system for offline processing to produce pulsar/transients data. The phased ALMA voltage data is recorded on VLBI Mark VI disks. In our Study the data were played back and processed at Cornell, Haystack/MIT, and MPIfR/Bonn to produce PSRFITS format pulsar data. The box represents the software suite that converts VDIF format data to PSRFITS data. PSRDADA is a software module that produces full time resolution detected time series. DIGIFIL is a module that produces time-averaged dynamic spectra usable as input to standard pulsar search pipelines.

Concerning Phasing Modes

Three general phasing modes have been considered for operating the ALMA Phasing System (APS) in this study: "Phasing-up-on-in-beam-target", "Phasing-up-on-in-beam-calibrator" and "Phasing-up-on-out-of-beam-calibrator". The first would use the APS on a normal continuum science target (such as Sgr A*) that is bright enough to phase-up the array and allow for deep searches in that vicinity. We were not able to identify a suitable source for testing during the period of this study, but there is no reason that the existing APS could not adequately phase-up and execute such an observation. The third requires modifications to the phasing engine of the APS (which resides in the TelCal software component) so that the phase adjustments steer the beam away from the calibrator and onto the weaker pulsar target to provide constructive phasing of the pulsar signal. This capability was envisioned in APP, but was not implemented for programmatic reasons. The APP phase two project (APP-2 for short, and where distinction is necessary, the original APP becomes APP-1), which has been approved by ALMA NA and which awaits formal approval by the ALMA Board, plans to implement a "source model" capability.

This "source model" mode capability should then allow this sort of pulsar observation to be commissioned and offered in a future observing Cycle. An obvious application would be to use Sgr A* to phase-up a coherent signal steered to the galactic center magnetar. The envisioned implementation would allow specification of the "source structure" via a CLEAN component list. For use in a pulsar observation, one would assign a dominant component corresponding to the calibrator and a weak secondary component corresponding to the pulsar. The plan also includes a more general "source model" (e.g. specified as an image) which would work for more complicated source fields.

The remaining mode (using an out-of-beam calibrator) is effectively implemented in the current observatory code; but it is not currently available for general use as it requires some special setup—this is the so-called "passive" phasing mode—that is not generally offered. The intent of this mode is to phase-up on a bright calibrator prior to the VLBI scan, turn off the active phasing calculations, and then slew to the weaker source and record it while retaining the phase solutions from the calibrator. Provided the weather is very good, the phase solutions should be adequate for several minutes. This mode may be offered in Cycle 6 to lower the threshold for VLBI observations below the current 500-mJy limit required by the phasing system. This mode was in fact used for the test observations which provided the data for this study (see below) and it could be offered to general observers with minor modifications to the observing script and programmatic changes (i.e. it would need to be approved by the ObsMode process, Accepted, and added to the Call for Proposals in some future Observing Cycle with criteria for proposal evaluation sorted out for the Call). The code modifications to the observing script are likely to be worked out in APP-2 (for weaker sources) so this would be the mode that has the greatest chance of early offering once the programmatic issues are addressed.

IV. Software

The project team developed two independent software packages for processing test data. Each built upon software available to team members at Haystack/MIT and MPIfR.

The standard VDIF data output from the Mark 6 recorder during APP observations are 2-bit sampling, single polarization per file and multiple channels (32 in total). Data from two polarisations are stored in individual files. The output format is the same for all four 2-GHz sub-bands. Accordingly, two software data processing pipelines have been developed which may fit different purposes for the usage of the data at a later stage:

Software Pipeline 1: vdif2psrfits (MPIfR):

1). Vdif – psrdada – filterbank / archive
This pipeline first converts the vdif output into PSRDADA format, one of the standard formats used in the pulsar community to store baseband data. The PSRDADA output consecutively attaches two polarizations of the same time stamp, but stores different frequency channels in individual files. It also expands the dynamic range of the original data from 2 to 8 bit. The conversion is done with the ***vdif2dada*** routine developed by the team, and the data converted as such can be readily used for beam-forming (if there are any other telescope datasets available) with the software correlator developed for the Large European Array for Pulsars.

Next, the pipeline uses the standard software package ***dspsr*** to either reduce the data into filterbank timeseries which are readily usable by the pulsar searching software packages ***presto*** and ***sigproc***, or into sub-integration archive data suitable for pulsar timing software ***PSRChive***. The former is done with the routine ***digifil*** and the later done with the routine ***dspsr***. The configuration of the output (time resolution, number of channels, length of sub-integrations, etc.) can be specified as options to these two routines.

2). Vdif – psrfits (search) – archive
This pipeline first converts the VDIF output into search-mode PSRFITS format, one of the standard formats used in the pulsar community. The conversion is done with the ***vdif2psrfitsALMA*** routine, developed specifically for this project by the team. For each frequency channel, this routine takes the time-series within one storage unit in the VDIF file (commonly referred to as a frame), performs an FFT to go to the frequency-

domain, makes detection from voltage to intensity, and collapses the power over the whole FFT spectrum. The detection can be in the form of total intensity, power of a single polarization, or coherence product (easily converted into Stokes parameters) as required.

The PSRFITS output has a basic time resolution of 8 microsecond, and one can give the factor for re-binning. The 32-channel frequency resolution is retained. The output can be directly taken over by ***presto*** to perform pulsar searching, or by ***psrfits_util*** software package to fold into sub-integrations.

Software Pipeline 2: VDIF to PLSR to PSRFITS (SAO)

VDIF to PLSR (vdif2plsr): Each VDIF stream from ALMA contains 2-bit samples for all 32 channels (62.5MHz wide) within one subband and for one polarization. The datastream is divided into frames which each contain 1000 samples for each of the 32 channels (spans 8 microseconds). The first processing step counts the number of samples that occupy each of the four 2-bit states per channel per VDIF frame, and writes (three of the four) counts per packet as well as the VDIF timestamp for each packet to a PLSR file (custom binary format). No output is produced for VDIF frames marked as "invalid" at this stage of the processing.

PLSR to PSRFITS (plsr2fits): This step reads a PLSR file and converts it to a collection of PSRFITS files. The sub-integration data written in PSRFITS format contains average power estimates per channel over each 8 microsecond period spanning a single VDIF frame. The average power estimate is due to a note by Lindy Blackburn, *Calculating Total Power from VLBI Data*, June 1, 2015; see equation (4) therein. This code is derived from lmt2fits by Michael Johnson et al. There is a limitation on the size of each PSRFITS file so that data for each scan may be spread across multiple files.

Step-by-step example: Assuming all dependencies are met and the attached tarball extracted to $SOME_PATH, the following shows roughly the steps involved to go from VDIF to the PRESTO output shown below:

```
# VDIF to PLSR
cd $SOME_PATH/vdif2plsr
make pulsar
./pulsar -1 some_vdif.vdif some_plsr.plsr
# add filler data if VDIF contained invalid frames
cd $SOME_PATH/filldata
python filldata.py -o some_filled_plsr.plsr some_plsr.plsr

# PLSR to FITS
cd $SOME_PATH/plsr2fits
make
export YUPPI_DIR=$SOME_PATH/plsr2fits
./plsr2fits -o some_fits -g 1000 -n 32 -w 62.5 -dt 8 -src J0835-4510
-raj 083520.61 -dej -451034.88  some_filled_plsr.plsr

# Basic PRESTO commands, change as needed
prepdata -o some_dat `ls some_fits*.fits`
realfft some_dat.dat
accelsearch some_dat.fft
prepfold   -o   some_dat   -nsub   32   -accelcand   X   -accelfile
some_dat_accelfile some_dat.dat
```

V. Test Data Acquisition

April 2016

During the April 2016 APP commissioning campaign a few minutes were granted for a test observation on a pulsar. The local time testing slot was arranged for the VLBI observations planned for the VLBA and the EHT, so it was necessary to locate a suitable pulsar target that would be visible shortly before or shortly after the planned observations. We identified PSR J1550-5418 (AXP_1E1547.0-5408) which has a period ~2.07s and proceeded to plan and execute a Band 3 observation using a nearby bright quasar J1617-5848 as the phasing target. The observing plan called for several scans in normal passive mode as well the "fast mode" (see Appendix, *APS Considerations for Pulsar Observations*).

There were a variety of operational issues that yielded baseband data (approximately 10 TB) suitable for searching for spurious signals, but unfortunately it was not suited for a search for the real test pulsar due to incorrect execution of the "passive" phasing mode. In the APP mode, there are two types of scans: VLBI scans (where phasing is possible) and ALMA calibration scans (which must be conducted as normal to support observatory calibration tasks). The ALMA calibrations are conducted in TDM (time division) mode rather than FDM (frequency division) mode. (These terms refer to the programming of the TFB digital filters within the ALMA baseline correlator.) The VLBI scans are specified by a schedule, and gaps of sufficient duration are available for the ALMA calibrations. In this case, the gap between the calibrator and each pulsar observation was large enough that the observing script attempted a Tsys calibration. We learned that mode switches from FDM to TDM and back to FDM actually clear the phasing registers. Thus the phase calibration obtained during active phasing was no longer available for the passively phased scans.

Although disappointing, the session was useful to obtain test data for development of the software processing pipelines and ensured success on the next try.

January 2017

A short test (< 30min) was granted during the "dress rehearsal" for the upcoming VLBI campaign to be held in April 2017. The time slot available allowed a plan to observe a known pulsar, Vela, near zenith. For the phasing, a suitable calibrator (J0828-3731) with a flux ~0.9 Jy was found within 10 degrees of Vela (see Figure 2) using the ALMA Calibrator Catalog (https://almascience.eso.org/sc). A weak phasing experiment in the intervening 9 months established that a 90s "active" scan was adequate for good phasing, and the "passive" scan could be attempted as soon as 30s afterwards with acceptable results. However to be conservative, we used 120s "active" scans with gaps of 60s. Thus scans No0001, No0003, No0005 and No0007 were "active" scans and No0002 (10 min), No0004, No0006 and No0008 were "passive". In all cases, the determined phases were retained within the TFBs. The schedule block was prepared with schedules to allow observation on any night of Jan 29 through Feb 1.

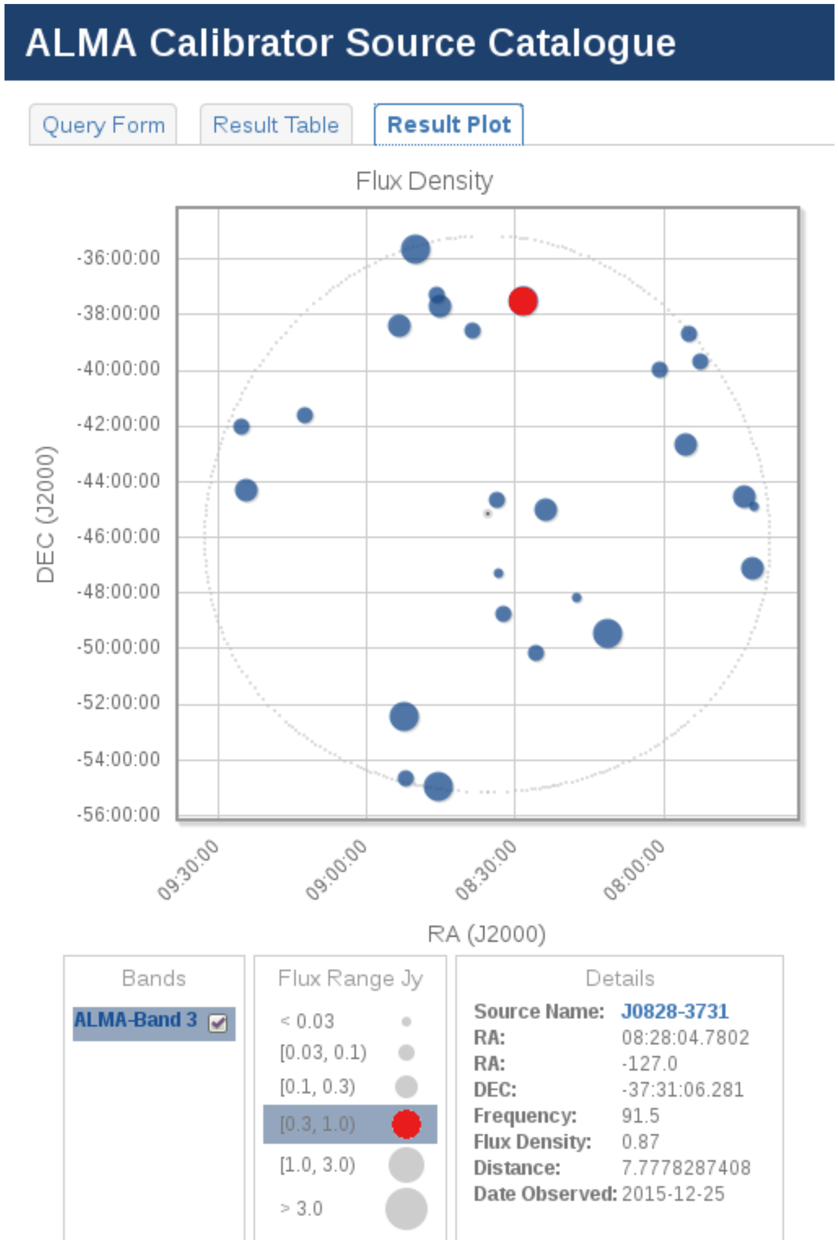

Figure 2: An illustration from the ALMA Source Catalog showing the result of a search for bright quasars near Vela. The brightest source within 10 degrees was chosen, J0828-3731.

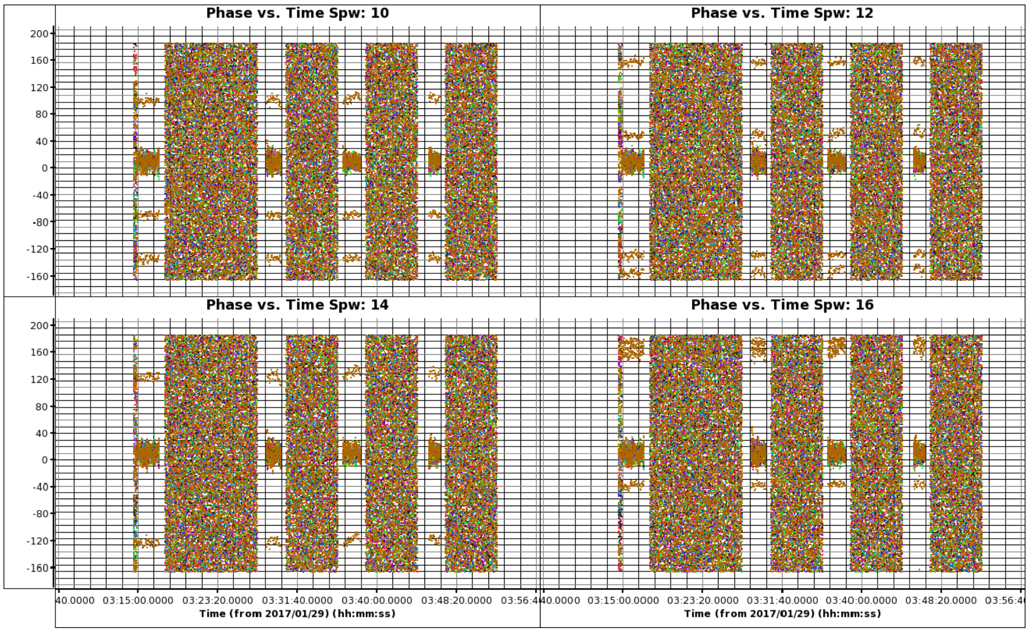

Figure 3: An overview of the Vela observation on Jan 29, 2017. The four panels show the visibility phase with time of baselines to the reference antenna for the XX correlator product for the four spectral windows at 86.268, 88.268, 98.268 and 100.268 GHz. The "active" phasing scans all last for 2 min and are characterized by a tight grouping of phases near 0 phase. The "passive" phasing scans are 10 and 5 min in duration and appear unphased. The three comparison antennas are not phase adjusted and may be seen to carry through the actively phased portions (see text), suggesting that the atmosphere was relatively stable during the passive (pulsar) scan intervals.

In the event, the weather on the first night was suitable, so the observations were made on Jan 29, 2017, with the first pulsar scan beginning at precisely 03h17m20s. The recordings usually start ~40s before the scheduled time (and run correspondingly longer) to allow adequate time for phase-up on active VLBI scans. The weather was excellent (much better than Band 3 requires).

A summary view of the observation is provided in Figure 3. This plot shows visibility phase with time for baselines from the "reference antenna" to all other antennas (of which all were phased except for two "comparison antennas" whose phase may be seen threaded through the plots. Only one correlator product (XX) is shown; the other (YY) is similar. The four subplots (labeled by CASA as "Spw: 10", "Spw: 12", "Spw: 14", "Spw: 16") are for the four 2-GHz spectral windows centered at 86.268, 88.268, 98.268 and 100.268 GHz. The 8 scans are clearly visible, with the 2-min actively phased scans appearing to have a tight grouping of +/- 20 deg around zero (corresponding to the phased antennas) and the longer 5 or 10 min passively phased scans appearing as a block of phase noise, since the pulsar target is not sufficiently bright to make the phases collapse to a band about zero. Note that the array was initially un-phased (a vertical band of arbitrary phase at the left-hand side of all plots) but that phases were retained from one scan to the next. As one can see, the spread of visibility phases at the start of each scan on the phasing target were still tightly grouped, implying that the antennas remained in phase throughout the Vela scans (which are dominated by the receiver noise and thus show no apparent coherence, as expected).

VI. Test Data Analysis and Results

We received 30 minutes of Test Observations with ALMA on 2017 January 29. The Vela pulsar B0833-45 was chosen as the test target source. The array was phased up on a calibrator and then the phases were held fixed while pointing at the target pulsar, as is typical for VLBI phase-referenced observations. Calibrator scans of 2 minutes each were repeated after 10 minutes on the pulsar, and then again after each 5-minute scan on the pulsar. After processing the data with the software pipelines as described in Section IV, we achieved independent detections of Vela with *both* the pipelines. A sample detection is shown in Figures 4 – 6 using both analysis pipelines.

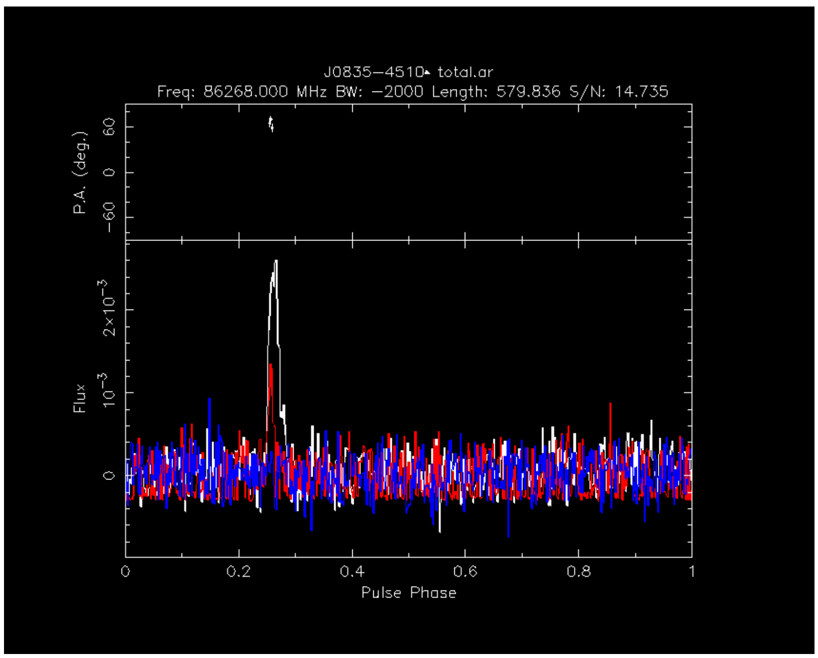

Figure 4: Pulse profile of the Vela pulsar at 86 GHz, using Pipeline 1 to analyze 10 min of data. One cycle of phase is shown after folding the data at the predicted pulsar period.

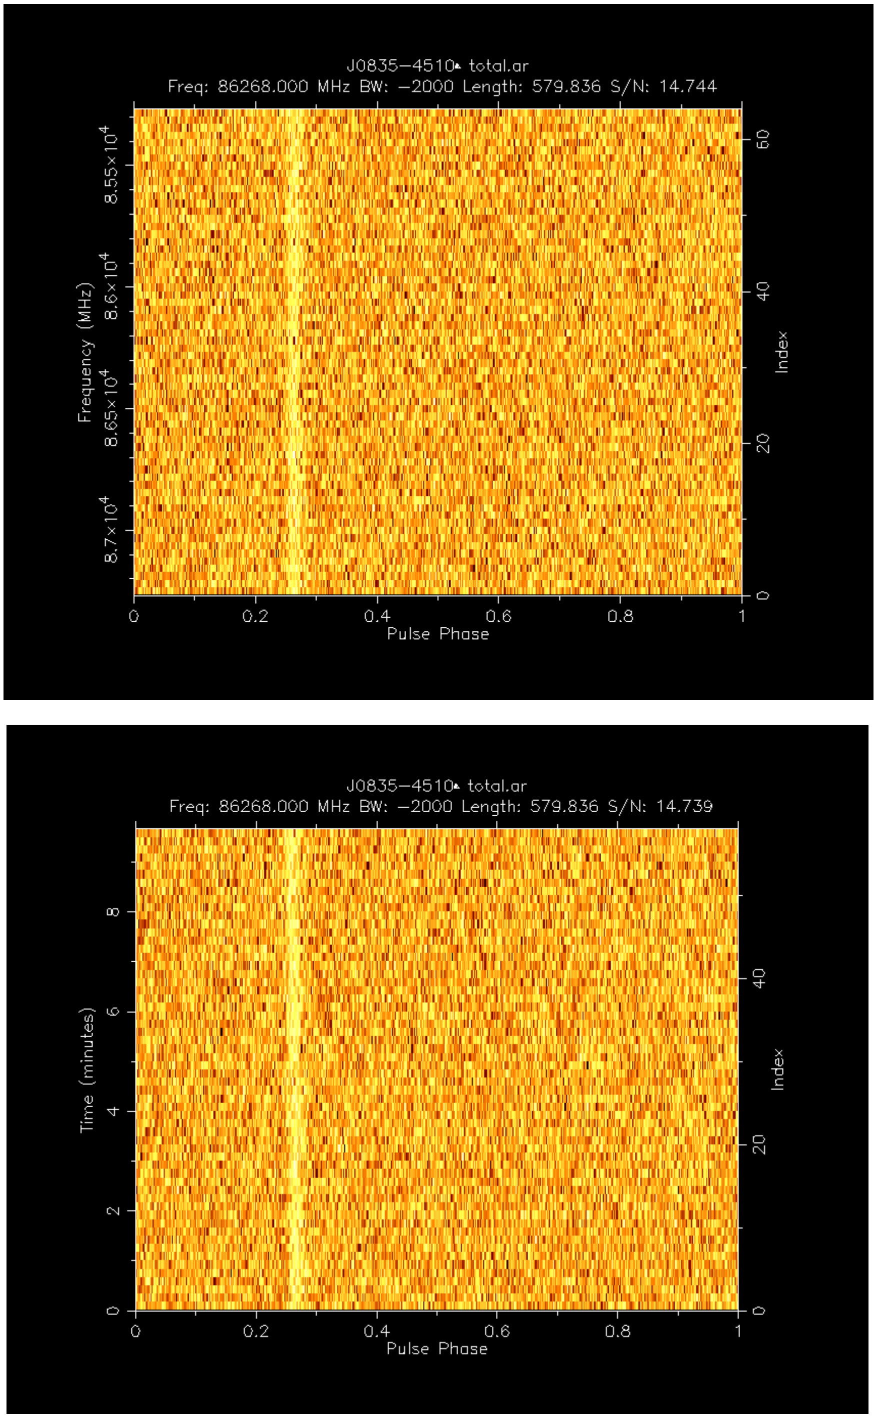

Figure 5: Vela pulsar detection using Pipeline 1, showing the pulse profile as a function of frequency (top panel) and as a function of observing time (bottom panel).

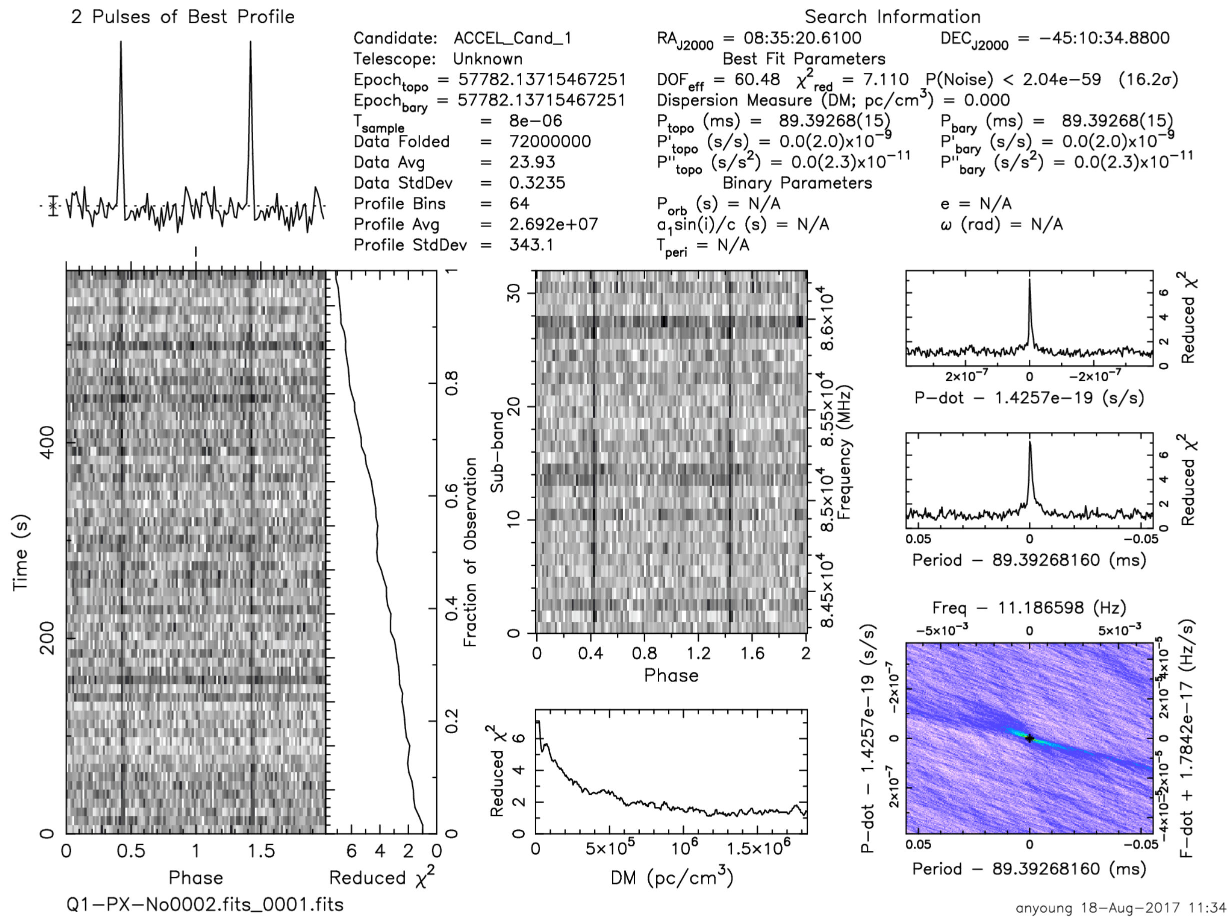

Figure 6: Detection of the Vela pulsar using Pipeline 2, with post-processing of search-mode data using PRESTO. Two cycles of the pulse profile are shown at the upper left, and the pulse profile is shown as a function of observing time (left grayscale panel) and observing frequency (middle grayscale panel).

VII. Next Steps

The goal of this study is the provision of additional ALMA observing modes that enables time-domain studies of phase-connected elements. In order to complete the task, in addition to the technical developments above, we plan to publish all details of this study, including an observing guide (see below), and results that demonstrate the necessary procedures. Simultaneously, all obtained data will also be made available at one (or more) of the institutions participating in this study.

In further developments, we suggest (and offer to participate) in the development of observing modes that offer different phase-up scenarios than that explored in the described work. The successful demonstration of pulsar observations indicates the potential of those modes for pulsar and other time-domain studies with ALMA. We suggest making the described modes available to all users following an appropriate request.

Accordingly, we will set up a pulsar working group within the Event Horizon Telescope (EHT) project. This working group will coordinate the EHT's request to use ALMA data obtained for VLBI observations also for time-domain (e.g. pulsar) studies, once the time-domain mode has been approved as an official observation mode that is available for proposals. These observations will include studies of the Galactic center magnetar and search observations to find pulsars at extremely high frequencies.

VIII. Draft Observation Guide for Pulsar/Transient Facility Observations (when available)

Guide for ALMA Pulsar Data Acquisition and Processing

Pulsar observations and processing are as follows.

Observing:
The first step for getting ALMA pulsar data is to successfully propose for and conduct your phased-array observations. See the ALMA observing guide for details on the observation process. Once the observations have been conducted, the phased-array data will be written to disk. The exact parameters of the data will depend on your particular observing set-up, but in general it will be a few (~GHz wide) channels of raw voltage data written in VLBI Data Interchange Format (VDIF). In order to process the data with standard pulsar software, we need to convert from VDIF to PSRFITS (or SIGPROC) format.

VDIF to PSRFITS:
The next step is to take the raw voltage data in VDIF and convert it to channelized and time-averaged data in PSRFITS format using vdif2psrfits (K. Liu, github link). This step is important for two reasons. First, PSRFITS is a standard data format in pulsar astronomy and just about any pulsar processing software you may want to use will be able to read it. Second, the small number of large (~GHz wide) frequency channels sampled at very high time resolution (~ns) are not ideal for pulsar observations because they are particularly susceptible to radio frequency interference (RFI). Splitting the data into smaller frequency channels means that less data will be lost due to contamination with narrow-band RFI.

The vdif2psrfits code operates as follows. First, raw voltage data are read from the ALMA data files. In VDIF, data are stored in "frames" containing short segments (e.g., 8 microseconds) of data in a single polarization. Frames from the same time are collected for both polarizations and combined to produce the desired coherence product. Next, the combined data are channelized using an FFT filter bank and time-averaged. The number of frequency channels and the output time resolution are free parameters that can be set by the user. For our test observations, we split each 2 GHz sub-band into 32 64 MHz channels with 8 microsecond samples. Finally, the channelized and time-averaged data are written to file in PSRFITS format.
At this point, the data are stored in PSRFITS as channelized time series of Stokes IQUV (or single pol, if desired) and are ready for analysis with your favorite pulsar processing software. For those new to pulsar observing, we will briefly describe the general steps below.

RFI Flagging:
Since pulsar signals can be very faint, we want to remove any RFI signal that could contaminate our results. RFI signals can be narrow-band or broad-band, transient or persistent. It is often too difficult to excise only the RFI signal from the data, so we mostly just mask the segments of the data containing RFI. For example, if we have a narrow-band RFI signal (bandwidth < MHz), we can just mask out the frequency channel containing it. If we have a transient broad-band RFI signal, we can mask out the time samples containing it. Some pulsar processing software packages (e.g., PRESTO) have routines for automatically classifying and masking RFI in a data set.

De-Dispersion:
The propagation of radio waves through the cold electron plasma of the interstellar medium produces a frequency dependent time delay called the dispersion delay. The dispersion delay is

$$t_{DM}(\nu) = 4.15 \text{ ms DM } \nu^{-2}$$

where the dispersion measure (DM) gives the integrated electron density along the line of sight to the pulsar in units of pc cm^{-3} and the frequency ν is in GHz.

The delays across the ALMA band can be quite small due to the high observing frequencies. For example, the delay between 86 GHz and 100 GHz for a pulsar with DM = 1800 pc/cc is only 0.26 ms. Thus the dispersion delay can be ignored in most cases except when very fine time resolution is needed for large DM objects (e.g., in the Galactic center).

After the frequency channels have been shifted according to the appropriate dispersion delay, the data in all the frequency channels are summed together to produce a 1-dimensional time series. This time series can then be used to search for new periodic signals (or transients) or to fold the data on a known pulsar ephemeris.

Folding / Searching:
Given a known pulsar ephemeris, you can fold the time series data (de-dispersed at the known dispersion measure) to produce an average pulse profile. If you don't know the ephemeris of the pulsar, you can blindly search for it. De-disperse over a range of trial dispersion values, generate power spectra from each time series, and search for periodic signals.

IX. Resources and Products

Phased ALMA Pulsar and Transients Project – Documents, Software, Products:
http://astro.cornell.edu/research/almapsr/

Links to Software:
vdif2psrfits on GitHub.
vdif2psrfits (frozen version hosted at Cornell, August 2017).
vdif-plsr-psrfits (frozen version hosted at Cornell, August 2017).

Appendix

APS Considerations for Pulsar Observations

A full description of the APS is to be found in the forthcoming paper, "The ALMA Phasing System: A Beamforming Capability for Ultra-High-Resolution Science at (Sub)Millimeter Wavelengths" (to be submitted to PASP, presently). Likewise a complete description of the ALMA observatory (regarding technical details) is to be found in the proposer's technical handbook which is updated for every observing Cycle. (The current planning cycle is Cycle5, so see https://almascience.nrao.edu/documents-and-tools/cycle5/alma-technical-handbook/view).

Here we merely review and point out some of the salient features of operation. The data from the receivers is digitized at 4000 Msamples/second and truncated to 3 bits (which preserves 96% of the signal) for transmission to the ALMA correlator(s). The APS uses the "baseline" correlator (the ALMA Compact Array [ACA] is used with the 7-m dishes); these signals are received by "station cards" at the front end of the baseline correlator where they are processed according to the correlator mode of operation. In the APS system, the station cards program their 32 tunable filter banks (TFBs) to FDM mode, assigning each TFB to one 62.5 MHz frequency sub-band with the plan that all 32 may be stitched together to form a 2-GHz continuum observing band. Since the TFB response is not flat at the edges of their sub-bands, the baseline correlator discards the outer 1/16 of the each sub-band (1/32 on each side) and stitches the remaining 15/16 of the frequency spectrum into a 1.875 GHz usable spectral window. When viewed as ALMA data, it appears to have 1920 spectral points—although to keep the archived data to a manageable size it is averaged by a factor of 8 for 240 spectral points across the continuum band. Note also that the TFBs resample the incoming 3-bit data into a 2-bit data stream. The ALMA software that picks up the correlator data "dumps" also takes care of appropriately scaling the data. Finally the TFB frequency channel centers are exactly 15/16 x 62.5 = 58.59375 MHz apart which is not a particularly convenient number for VLBI (and probably not for pulsar work either).

In normal operation of the APS, the 2-bit data from a set of *N* 12-m antennas (the "phased array") is digitally added and scaled again to 2 bits to form a "sum" antenna signal which may be fed back into the ALMA baseline correlator (with a delay of 192 ns which requires compensation) for on-line diagnostics of the signal. These diagnostics requires that one or more antennas of the operational array be left outside of the "phased array" and are termed "comparison antennas". (as an aside, we comment that the APS can also make use of the 7-m antennas, but there is marginal utility in so doing.) However the main purpose of the "sum" signal is to be recorded for VLBI, or in our interest, high time resolution pulsar data. The significance of the sum signal depends greatly on the phasing of the "phased array". When all the of the phased antennas are brought into phase with one of their number, the so-called "reference antenna", the signals interfere constructively and the VLBI recording is sqrt(*N*) times more sensitive, and performs much as single antenna of size 12 x sqrt(*N*) m. It was a special consideration in the design of the TFBs that they contain a "phase register" which allows an external agent (i.e. the APS) to insert phase shifts dynamically during the course of correlator integrations.

The normal or "slow" phasing loop is governed by the time required to feed the correlator data to the TelCal phasing engine where the phasing solutions are calculated. The protocol and bandwidth of the underlying hardware require ~8 s to transfer the data; having sufficient data the to phase-up on all targets of interest (i.e. brighter than 500 mJy) requires an additional 8s of on-source integration. Adding in the dump time (a few seconds) and the time to calculate and apply the phasing adjustments (less than one second) requires a phasing "loop" of the order ~20 s before phasing adjustments may be inserted into the TFB phasing registers. The timeline is as shown in Figure 4. In that figure, the top bar represents the target VLBI scan (or pulsar recording). On each cycle of the "slow" loop, a correlator sub-scan (usually containing 4 4-s integrations) is passed to the TelCal phasing engine, which reports back corrections. Note that these corrections are made during the start of the next sub-scan, so typically the first integration needs to be discarded. Also shown is a "fast" loop that uses the water vapor radiometer (WVR) data to make adjustments for wet-atmospheric delays. While it is possible to run the two types of loops independently or together this has rarely been done. However, there is some sense to run the "fast" loop alone during "passive" scans to allow operation in wetter conditions than otherwise.

The VLBI recordings are organized by VLBI scan which is scheduled to start at a particular time and run for a set duration. Each of these recordings consist of a time-stamped sequence of VDIF (VLBI Data Interchange Format) packets, which each contain a small header and a larger batch of 2-bit data. The packets in use at ALMA use the "extended data header version 2" which is fully documented at:
http://www.vlbi.org/vdif/docs/alma-vdif-edv.pdf.
The Mark6 recorders in use at ALMA distribute these packets according to a so-called "scatter-gather" plan to guard against data loss due to disk failures. Decoding the recorded data requires some understanding of the "scatter-gather" format (which gathers ~10 MB of packets into blocks and then puts time-adjacent blocks on to different disks within a recorder module) to reconstruct the packet stream. However at that point it is sufficient to read the sequence of 8032B packet headers and sort the data bits into the 32 frequency sub-bands. The packet header represents time as seconds from some epoch (start of semester) and a packet counter. Each packet also has a serial number to be sure of consistency. The 32 2-bit TFB sub-bands are time ordered, so each 8 ns sample occupies one 8B octet with 1000 such per packets at a rate of 8us/packet or 125000 packets/second.

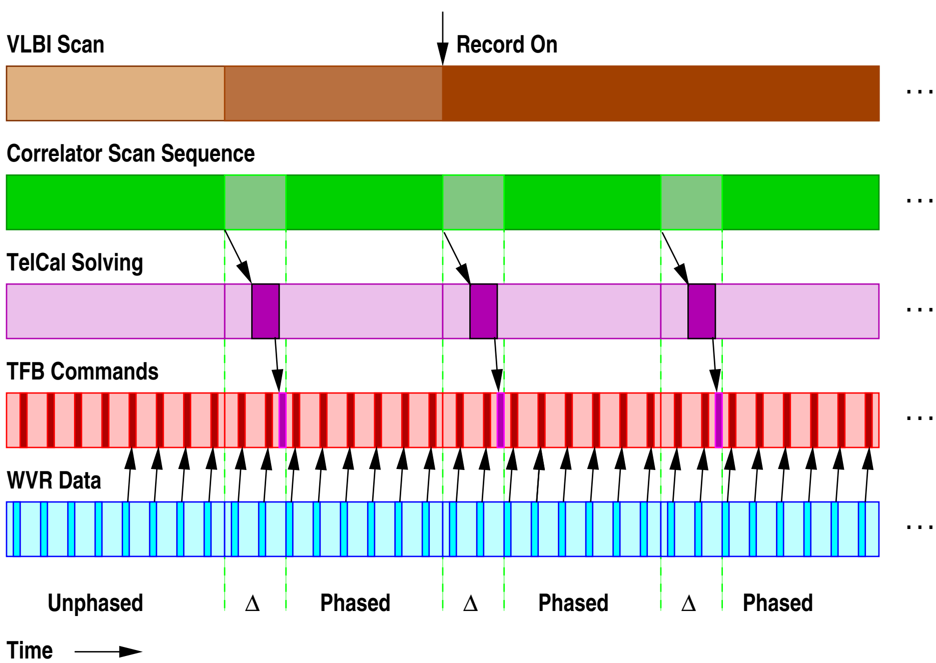

Figure 7: This figure schematically illustrates the phasing loops used by the APS. Time flows to the right in a correlator sub-scan sequence of data acqusition (darker green) and dump intervals (light green). The target VLBI scan is at the top; a two correlator sub-scans are usually adquate to achieve a phased-up signal for recording. The phase solutions are calculated in the TelCal software component (purple) and returned shortly before or after the start of the next correlator sub-scan at which time phasing adjustments may be made (red), closing the "slow" loop. Asynchronously, water vapor radiometer data (WVR) is available approximately every second at which point water vapor delay corrections may also be applied for an open "fast" loop.