EVLA Memo XXX
Phased array observations at the EVLA; Updated

Adam Deller & others

February, 2011

Abstract

Phasing the EVLA is a high priority for a large number of potential users, particularly VLBI and pulsar astronomers. This memo describes the modes of phased array operation which will be made available at the EVLA, the input necessary to "drive" the EVLA in phased array mode, and the outputs which will be produced. Finally, the subsystems that will contribute to producing the phased array outputs, with a particular focus on the "Y27" single subarray mode, which will be the first available, are described, and the necessary developments in each area are summarized.

Contents

Definition of acronyms used in this memo

ALC	Automatic Level Control
CASA	Common Astronomy Software Applications – software package for reducing EVLA data
CBE	Correlator back end – the cluster of nodes which captures data from the baseline boards, performs further accumulation and formatting, and then sends visibility data and metadata to the archive, MCAF & TelCal
DiFX	Distributed FX software correlator used for the VLBA
MCCC	Master Correlator Control Computer
MCAF	Metadata Capture and Format, the software component responsible for gathering visibilities and metadata and storing the formatted data in the archive
NGAS	Next Generation Archive Server
OPT	Observation Preparation Tool
OSRO	Open Shared Risk Observing
PST	Proposal Submission Tool
RSRO	Resident Shared Risk Observing
RXP	Retiming, Crossbar & Phasing. The FPGA on the baseline board responsible for forming the phased array sum.
SEFD	System Equivalent Flux Density
SDM	Science Data Model, all the non–visibility metadata in the archive
TelCal	Telescope Calibration, the software component of the EVLA monitor & control system that derives calibration quantities from the correlator visibilities
VCI	Virtual Correlator Interface, a document exchanged between the executor and the Configuration Mapper
VDIF	VLBI Data Interchange Format
VEX	VLBI EXchange (an observation description file format)
VLBI	Very Long Baseline Interferometry
XML	Extensible Markup Language
Y1	A mode of operation where a single EVLA antenna used for VLBI whilst the remainder of the array is used for normal observations.
Y26+Y1	A mode of operation where up to 26 EVLA antennas are phased on a weak target and a single EVLA antenna continuously tracks a calibrator source (reducing the need for slewing)
Y27	A mode of operation where up to 27 EVLA antennas are phased (no subarrays)

1 Introduction

Forming a phased array requires the individual antenna voltages to be delayed to a reference position and coherently summed. In the case of ideal case of perfect phasing, identical antennas and independent noise (the weak–source limit), the sensitivity gain produced is linear with the number of antennas summed. In comparison, incoherent summation of the individual antenna signals (by detection before summation) leads to a sensitivity gain proportional to the square root of the number of antennas. Incoherent summation, which is not useful for VLBI, is not the subject of this development effort and will not be supported at the EVLA in the near term. See Moran & Dhawan (1995) for more details on phased arrays.

The field of view of the phased array is considerably smaller than the primary beam of the individual antenna elements, in the ratio D/B, where B is the maximum baseline between elements and D is the diameter of the antenna elements. For this reason, it may be desirable to form more than one phased array beam (e.g., for astrometric observations with an in–beam calibrator).

The EVLA will be used in phased array modes for two main types of science: VLBI, and time domain studies of pulsars and transients. In both cases, the phased array time series of voltages will be sent to one or more disk recorders (the Mark5C recorder – see the latest memo by the Mark5C collaboration: NRAO/Haystack 2008). The phased array output stream could in principle be accepted and processed by some other device, such as a pulsar backend. However, such applications will not be considered in this memo. It is noted here that the EVLA phased array data will be produced in a standard format, as discussed in Section 4.1, and thus any device which can read such a format and connect via 1 or 10 Gb ethernet could be easily integrated into the system.

2 Modes of operation

Many of the most common modes of operation in which phased array data will be used from the EVLA require a subarraying capability. Subarraying capabilities, however, are not envisaged to be available in the near (< 1 year) future. Thus, the first mode of phased array operation to be implemented will be "Y27", which is a single subarray, with all (available) antennas phased. The remainder of this memo will focus on the developments which are necessary to make this mode of operation possible.

However, for completeness, the cases where phased array observing will require subarrays are briefly considered here. Whilst no effort will be made initially to support these modes, the design of the phased array system will attempt to make their later integration as simple as possible. Ultimately most phased array observations with the EVLA are likely to come from one of the following categories:

- a "Y1" mode, where a single EVLA antenna is used for VLBI, while the remaining 26 antennas are used for a normal EVLA observation (two subarrays).
- a "Y26 + Y1" mode, where 26 antennas are phased on a VLBI target source and a single

antenna tracks a calibrator, reducing the need for slewing during phase referencing (two subarrays).

- a hybrid of the above, where the majority of the EVLA is used for normal observing, but M antennas are phased on a VLBI target and one antenna tracks a phase reference source (three subarrays).
- An "A×Y1" mode, where A antennas are used in "Y1" mode, and the remaining 27−A are used for normal EVLA observing (A+1 subarrays).

All of these cases can be considered to be simplifications of the generic many subarray case, and so the general case is enumerated here:

- The array is divided into N subarrays, each of 1–27 antennas.
- Each subarray which is being phased forms M phased array beams, where M may vary between subarrays.

The total number of phased array outputs from the correlator in this case will be P, where $P = \sum_{n=1}^{N} M(n)$. In practice, the need to record or analyze each of these phased array beams simultaneously will limit P to a small number. The hard upper limit to P is 256, since there are 128 baseline boards, and each baseline board can phase two "subband polarizations". Typically a single baseline board would produce dual polarization data for one subband frequency. As noted in Section 5.6, the initial upper limit for P will be 20 dual polarization beams, since there will be 20 baseline boards with the requisite network connections. In the "Y27" mode which will be the first made available (N=1), we will consider the case of making either one or two phased array beams from a single subarray available.

3 Configuring phased array operations

Ultimately, the EVLA is controlled through the "executor" described in Chapter 9 of the EVLA project book. The executor takes observing scripts (which for normal EVLA observations are generated by the dynamic scheduler) and translates these into commands for the various EVLA subsystems (antenna, correlator, etc.) Some of the details of the configuration are delegated to other programs – e.g., the correlator configuration is handled by the Configuration Mapper program.

For pulsar and transient observations, where the EVLA is the only array that must be controlled, the normal path of configuring the instrument using the OPT[1] should suffice. However, this has required some additional functionality for the OPT. New features will include:

- The ability to specify that phasing solutions should be generated at a given scan (an "intent"). The intent "CALIBRATE AUTOPHASE" was used.

---

[1] http://science.nrao.edu/evla/observing/opt.shtml

- The ability to specify that phasing solutions should be applied (both both the correlated WIDAR output and the phased array output). This was made available as a "modifier" called "APPLY LAST PHASE/DELAY CALIBRATION"

- The ability to configure the output data format (number of subbands, number of bits, destination IP/MAC address etc.; configured via the resource configuration tool in the OPT, so associated with the correlator configuration used)

- The ability to turn phased array data generation on/off – this has been tied to the resource configuration (so if a resource which configures the output data format is selected for a scan, phased array data will be generated for this scan, otherwise not).

- The ability to specify the number of phase centers, and the offsets from the pointing centre for a given scan (initially defaulted to one phase centre at the pointing centre, general case to be determined)

For VLBI observations, however, the schedule for all participating telescopes is generally produced using NRAO SCHED[2]. Duplication of a schedule for the EVLA via the OPT would be a time–consuming, error–prone and unpopular option. Thus, a means to control the EVLA using information produced from SCHED is necessary.

The VEX format file produced by SCHED contains most of the necessary information (source, scan, LO/IF, number of bits etc) to configure the telescopes included in the VLBI observation. Indeed, as part of the VLBA upgrade program, VLBA antenna control will be gradually converted to "EVLA–style" scripts, which will be produced from the VEX file using a program called vex2script. This approach could be extended to include the necessary extra information for the EVLA, such as which sources to phase the array on.

In both cases, the commands to start and stop the recorder(s) should be automatically embedded in the generated script. As discussed in Section 4.1 below, the recorder start and stop commands should eventually be possible at very fine time granularity, although initially a granularity of one second will suffice. In general, the decision of which bitstreams to send to which recorders (once more than one is available) will be made on the basis of simple heuristics. The Mark5C units themselves can automatically decide which bank to record to, and specific configuration at this level will not be required, although should be made possible.

The hardware implementation of the phasing logic does not allow for antennas to be weighted in the phasing process[3]. Thus, stations with poor performance (pointing, system temperature, rapidly varying phase etc) must be completely removed from the phased array output once their unweighted addition is no longer beneficial. This and some other aspects of configuration of the EVLA for phased array observations cannot be handled automatically by the executor. These aspects which will need to be under operator control will need a separate interface with which the operators can interact. This is discussed further in section 5.2.

---

[2] http://www.aoc.nrao.edu/software/sched/index.html

[3] stations could effectively be weighted by changing the gain in one of the upstream filter stages, but since the station board outputs are requantized to 4 or 7 bits, the available headroom is limited, and in the interest of simplicity this option is not considered at present

Since the station signals cannot be weighted in the phased array summation, it is important that there are no large deviations in the signal levels at each station from the ideal values. An rms deviation of less than 1 dB would be considered acceptable, leading to an SNR loss of <5%. Currently, the gains for each stations are set during the first scan at a new frequency band, which corrects the signal levels at that time. However, it is possible that the signal levels at different antennas will drift over time, which would lead to an unsatisfactory, arbitrary weighting of antennas in the phased array sum. A capability for auto–leveling of antenna signals exists but is not yet production ready. The stability of the relative gains between antennas should be investigated and if the current "set and forget" performance is unsatisfactory, the auto–leveling tool should be developed to usable status.

The phased array sum signal can (and indeed must) be scaled before requantization to ensure optimal state counts. This is controlled by a low–level process that is discussed in Section 5.4.

4 Phased array outputs

In addition to the phased array baseband data, additional calibration and metadata must be produced and recorded. This section describes the format of each output product of the phased array system.

4.1 Baseband data

The baseband data will be formatted as VDIF data packets. VDIF is described in detail in Whitney et al. (2009). VDIF is a flexible format in which the quantized baseband data samples from different subbands can be interleaved in a single "thread" (where a thread is a series of ethernet packets, each with a data header), as was generally performed in older VLBI formats, or can remain in separate, parallel "threads". The EVLA will operate in the latter mode, where the phased array datastream for a single subband (as produced by a single baseline board) will be placed in a single VDIF thread. Each of the RXP chips on a given baseline board will phase one polarization of one subband, and these will be sent as separate VDIF "threads". Thus, each polarization of each subband will be separately packetized and sent when dual polarization data are being produced by a single baseline board.

The VDIF datastream is completely self–contained, and requires no further meta–data. The destination for the packets will be set at configuration time by the Configuration Mapper. By default, the destination will be made on the basis of simple heuristics, such as sending all the data to one recorder if the data rate is sufficiently low, or otherwise subdividing by beam[4]. Initially, of course, only one recorder will be available and the destination values could be hard coded in the Configuration Mapper. Ultimately, it would be desirable for human intervention at the vex2script stage (before the commencement of observing) to allow the

---

[4] It is important that these heuristics are coordinated with those applied to other telescopes in the array, since it will make correlation considerably easier if the division is identical at all antennas

destinations to be configured manually. In this instance, the executor would be required to pass additional information on to the Configuration Mapper.

A desirable feature for the baseband data output would be the ability to start and stop the production of VDIF data packets on a short timescale. A mundane application of this ability would be stopping recording during slew times, allowing more economical usage of disk space. Future applications are foreseen such as so called "burst mode" recording for VLBI, where the data source exceeds the available record speed for short periods, followed by a time where no data is produced. Potential applications include fast phase referencing at high frequency and potentially some pulsar studies. However, this is not an immediate driver.

In order for the VLBI position and *uvw* values to agree, it is desirable for the fringe–stopping position to be settable (ideally, to the point at or at least near the geometric center of the collecting area of the antennas). This is particularly relevant to the future usage of "Y1" and "Y26 + Y1" – if the Y1 antenna is far from the array centre, the VLBI *uvw* values assigned would be inaccurate. Initially, all phased array beams will be delay– and fringe–stopped to the array phase centre (meaning only a single VLBI position will be needed for the EVLA). Ultimately, however, it would be desirable for this to be a configurable value via the observing script, which would impact the delays being generated and disseminated by the executor. The delay stopping position would also have to be logged and made available to the VLBA monitor system.

An alternative approach (which would also be complementary if both were implemented) is to generate corrections to the VLBI *uvw* values based on the difference between the geometric array centre and the fringe stopping point. This could be generated by the evla2vlba script discussed in Section 4.2.2, but software would be required to integrate these corrections with the VLBI dataset, either at correlation time or in post–processing. Such corrections to the *uvw* values are not an immediate requirement.

4.2 Monitor and alert data

VLBI observations require both alert data (for use during observing to allow operator intervention) and monitor data (for use after observing to improve the quality of the correlated VLBI dataset through flagging and calibration). Alert data is transmitted in real time to VLBA operations in Socorro, while monitor data is stored on–site for a period of time and later transferred to the VLBA archive.

4.2.1 Alert data

For EVLA phased array operations the requirements for alert data are very limited, since most "normal" alert conditions will be handled by the the existing EVLA infrastructure and operators. The only new piece of hardware which must be monitored is the Mark5C recorder. Alert data will be generated and monitored using identical software to the Mark5C recorders at other VLBA sites, and no additional effort will be required for adaptation to the EVLA.

4.2.2 Monitor data

The required monitor data from the EVLA for VLBI phased array operations can be divided into three categories: flags, weather, and system temperature. The science archive from the standard EVLA processing during phased array operations will contain all of the necessary information to produce these monitor data. Accordingly, a CASA script will be developed to extract flagging, weather and system temperature information for the phased EVLA from the EVLA archive in an automated fashion immediately following an experiment, and store it in XML format files which can be compressed and sent to the VLBA monitor database system. This script will be referred to hereafter as evla2vlba.

With the introduction of new devices such as the Mark5C, the VLBA plans to begin running a second, parallel database for these new sources of monitor data. This database will be fed by a "loader" process which operates on XML format files which are very similar to those produced by the EVLA. Thus, no further effort will be required from EVLA software once the XML monitor files have been written by evla2vlba and sent to the monitor loader.

The flagging data will be relatively simple to generate. Flags for the phased array sum can be generated based on antenna–based flags for antennas which were contributing to the phased array sum. This will require the logging of commands which alter the composition of the phased array, as discussed in Section 5.2. These logs should be treated as normal array metadata, and be available from the archive. Heuristics will need to be developed to generate a single phased array flag from the individual antenna flags. A simple default scheme such as flagging if more than one contributing antenna is flagged for any reason should suffice in most instances, but evla2vlba should allow the algorithm to be configurable. Correlator chip flags with information specific to the cross–correlation results should be ignored.

EVLA weather data will be logged in the EVLA archive in a format similar to that used for VLBA antennas. It is probable that no format changes will be required for weather monitor data, and the EVLA XML files can forwarded directly to the VLBA monitor loader.

System temperature monitor data for a single antenna (although this mode is unlikely to see widespread use until the introduction of subarrays) will already be available in XML form in the EVLA archive, and thus as with weather, a simple forward to the VLBA monitor should suffice for "Y1" observations. System temperature is one of the VLBA monitor data products which will be migrated to the new VLBA database in 2010, with the introduction of new digital backends at the VLBA. The format will be identical in both cases; a time series of the "on" power, the "off" power, the estimated gain and the estimated noise calibration temperature. It should be noted that the system temperature extraction and logging for the EVLA is not yet mature and stable, although this will surely be the case by the time "Y1" observations become available.

System temperature monitor data for the phased array, however, is considerably more difficult to obtain than for single antennas. The SEFD of the individual antenna elements (itself dependent on both the system temperature and the telescope gain) and the accuracy of the phasing combine to determine the SEFD of the phased array output. A number of schemes exist for estimating the phased array SEFD given the available EVLA dataset, including the correlation data and switched power measurements. The scheme implemented should make use of the existing data products of the EVLA and should provide a useful and robust estimate of the SEFD at all times, both on strong and weak sources, and should require no further post–processing.

The system temperature calibration provided for the phased VLA was the ratio of antenna temperature (T_A) to system temperature (T_{sys}). This was then converted to an SEFD in post–processing after loading and flux–calibrating the VLA dataset and copying a known flux value to the VLBA dataset. This was both laborious and produced incorrect gain calibration for weak sources, which had to be manually excised.

Accordingly, the system temperature calibration to be generated from the EVLA will solve for the T_A to T_{sys} ratio of the phased array sum *only* for sources on which the array is being phased. The evla2vlba script will make use of the most reliable flux density calibrator in the dataset to scale the T_A to T_{sys} ratio (as calculated by summing the correlation coefficients for baselines to the reference antenna and correcting for the source contribution) immediately to an SEFD. No further weighting is necessary in the summation of the correlation coefficients because no weighting is applied to the individual antenna bitstreams during the phased array sum.

This initial calibration will then be linearly interpolated across scans where the array is not being actively phased and further corrected using three observables:

1. When the array returns to a calibrator source after a scan or scans on a target, the magnitude of the phase corrections will be used to estimate the decorrelation over time due to dephasing of the array during the non–calibrator scans. Initially a simple assumption such as a linear drift of phase for each antenna with time will be used.

2. The reference gain curve for each antenna in the phased array will be used to modify the interpolated SEFD based on the change in gain for each individual antenna with elevation.

3. The switched power measurements for each antenna in the phased array will be used to further correct the derived SEFD based on the measured variation in the individual antenna system temperatures during the target scan.

As noted above, it is not clear how mature the switched power measurement and logging will be for the EVLA during initial phased array operations, and hence the final correction based on switched power may not be applied immediately. The first two corrections should be simple to calculate and apply, however. The final phased array SEFD will be stored in an XML file as a function of subband, polarization and time, and sent to the VLBA monitor loader.

There are two desirable extensions to the system temperature monitoring for the EVLA. The first is the ability to specify the storage of monitor data (for both the single antenna and interferometric pathways) from subbands other than those recorded, for the case where the recorded bands are affected by narrow–band RFI. The second is the ability to produce the summed switched power for contributing antennas in the "Y27" case, as this could be used as a check against the interferometrically derived SEFD.

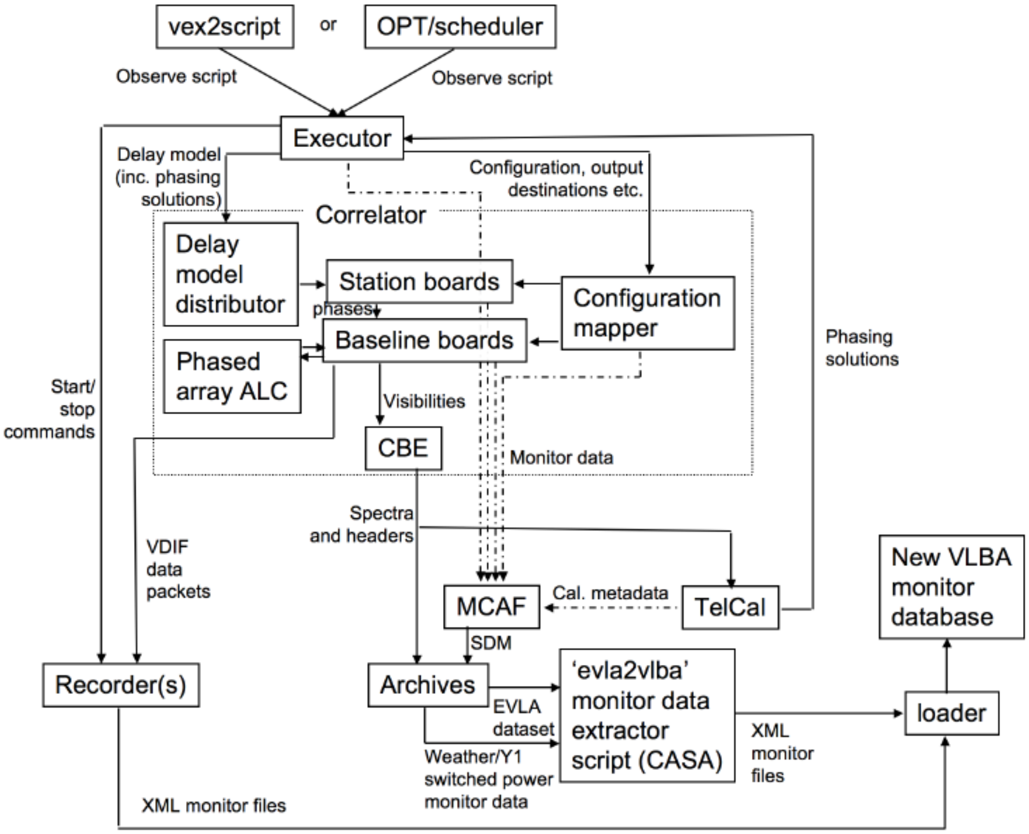

Figure 1 A functional overview of the subsystems involved in phased array observing, where non–standard EVLA functionality must be used.

5 Detailed description of subsystems

Figure 1 shows the functional layout of the subsystems involved in phased EVLA observations. Figure 2 shows the physical location and connection of the hardware components which are used in phased EVLA observations, and indicates where possible the location of the functional subsystems. These subsystems are now described in detail.

5.1 Configuration via the executor

The executor will be responsible for distributing configuration commands to the various components involved in the phased array generation and recording. These include the correlator station boards and baseline boards, the CBE, MCAF, TelCal and the Mark5C recorders. The only new element to be controlled, the Mark5C, will have a "virtual MIB" interface and hence may be controlled in the same manner as other hardware components of the EVLA.

Additional functionality has been added to the executor to understand the new configurations that will be needed for phased array observing (i.e., setup of the baseline boards to phasing mode, configuring the IP address for the VDIF packets, starting and stopping

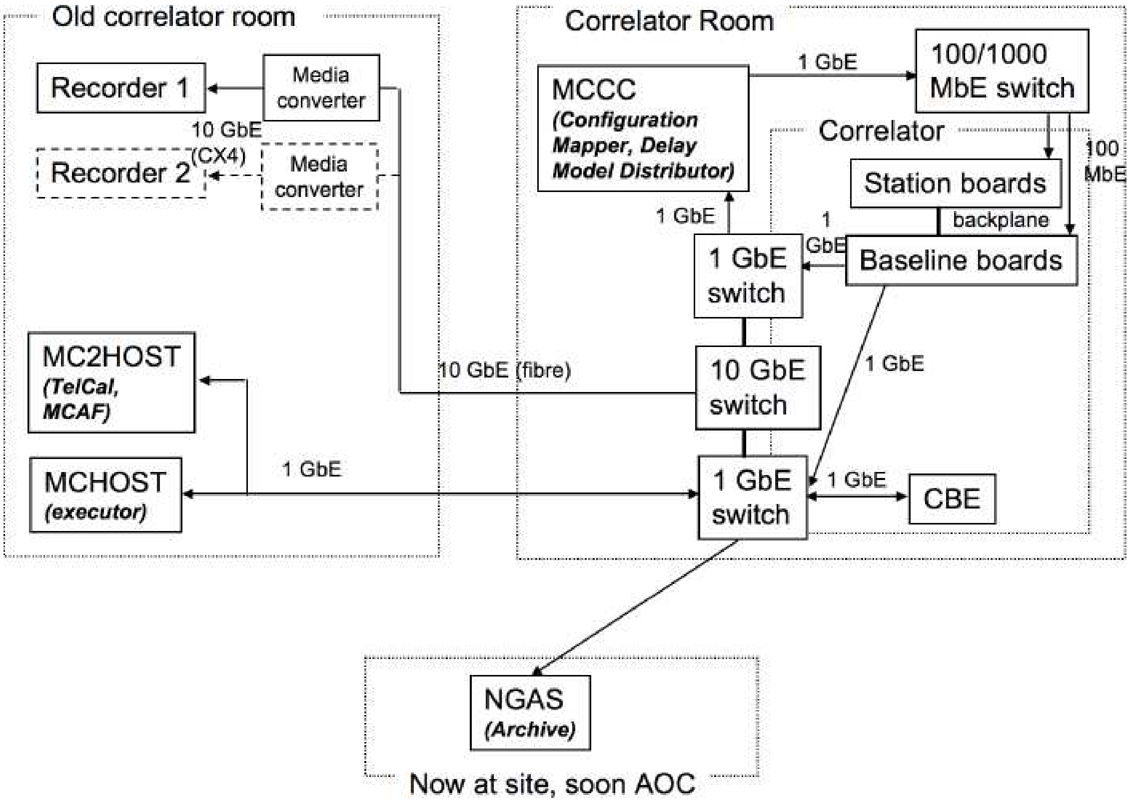

Figure 2 A physical overview of the hardware involved in phased array observing.

the recorder, etc.) Presently, since only one recorder is available, many of these configuration values have been set to defaults. Commands which are only required at the start of an observation (such as setting to phasing mode) do not require an extremely fast reaction. Reconfiguring the destination address for VDIF packets and recorder start/stop should be possible on short timescales, however. Initially, a reconfiguration on the next second tick would be acceptable – ideally, sub–second response times (particularly for recorder start/stop) would be desirable for the reasons discussed in Section 4.1.

The phasing solutions are recalculated and applied at the end of every subscan[5]. This is a natural interval as currently TelCal inspects the data at this level of granularity. A single phasing solution is applied until the next update – no attempt is made at predictive interpolation. The subscan interval can be specified by the observer in the OPT, and a similar customization for the SCHED–based configuration path should be made available, although initially a sensible default will suffice. Since the CBE may take as long as several correlator integrations to close a subscan and notify TelCal, there may be a small latency between the end of a subscan and the calculation of phasing solutions. The calculation of phasing coefficients will also take a short amount of time, dependent mostly on the amount of data (hence number of integrations and number of channels) that must be acquired by

---

[5] For a detailed description of the breakdown of an EVLA observation into scan, subscan, integration etc, see Morgan et al. (2004)

TelCal. Steps should be taken to ensure this net latency does not exceed 5 seconds, or half the subscan length, whichever is shorter, and the actual time of the updating of phasing coefficients by the Executor should be on an integration boundary and should be logged.

If very short phasing intervals are desired, the reduction in subscan duration to $\ll$10 seconds would lead to an unnecessarily large number of data in the EVLA SDM. Accordingly, a longer term goal is to generalize the calculation of phasing coefficients by TelCal to occur on timescales shorter than the subscan duration. This would also require changes to the Executor to check for new TelCal solutions on a similarly short timescale. This extended functionality is not a high priority.

The pathway for setting the reference antenna for phasing calculation is not yet determined, but the reference antenna is likely to be embedded in the script read by the Executor and propagated to the SDM and TelCal. Thus, vex2script should allow the reference antenna for the EVLA to be set. It may also be desirable for vex2script to annotate the original VEX file to note the reference antenna position (in the case of “Y1” observing, to make this information easily available to the VLBA correlator).

Currently, the executor is limited to sending one delay model per sampler. Currently, for multiple phased array beams to be implemented, the subbands must be taken from different samplers. Ultimately, it would be desirable for this restriction to be lifted, which would require changes to the executor. In any event, the executor should eventually be able to phase the array on a different position to the current pointing position, which will be required for multiple phased array beams when one beam is placed on the inbeam calibrator.

Presently, the delay model table which will be stored in the SDM can only store one delay model per antenna. It would be desirable to generalize this to allow for different delay models at each subband, since this is what can be actually implemented. However, for geodetic applications (the main point of interest where accountability of the EVLA delay model is crucial) multiple phased array beams are unlikely to be used, and hence the current architecture is probably sufficient.

It may be desirable for the executor to be able to automatically reconfigure the composition of the phased array sum (via a message to the Configuration Mapper) based on real–time metadata, such as flagging information. This would necessitate the passing of flag information from MCAF to the executor. Such a self–selecting phased array is not an immediate priority.

Finally, it should be possible to turn off the noise calibration (switched power) injection. This will be necessary for time domain (i.e., non–VLBI) studies. This should ultimately be selectable by the the executor script (as written by the scheduler or SCHED) but may initially be controlled via the non–executor configuration discussed below.

5.2 Configuration not handled by the executor

Some aspects of configuration for phased array observations cannot be handled by the executor. An example of this is the removal of poorly performing antennas from the phased array sum. This requires a command to be sent from the Configuration Mapper to the RXP chip with the new antenna mask for the phasing sum. This and other observation–

dependent changes to the configuration will be applied by the operator. This will require the development of some form of stand–alone program which can interact with, at a minimum, the Configuration Mapper and MCAF. It is envisaged that this would take the form of a phased array page for a GUI. The ability to deselect and reselect antennas for the phased array sum is the primary, and initially only, functionality required. The necessary VCI schema changes have already been made, so the low–level infrastructure is in place, but the interactive component remains to be developed.

Initially, it will be acceptable for an added antenna to be immediately (i.e., on the next one second tick) included in the phased array sum. A desirable default feature would be for the added antenna to be held back from addition to the phased array sum until after the next phasing solution calculation, allowing the operators to schedule the addition of an antenna at any time without degrading the phased array output before the antenna phase is corrected.

Further in the future, the ability to change the destination of the output VDIF data packets would be desirable, as would the ability to manually stop and start the Mark5C recorders. The ability to turn the noise calibration (switched power) signal off or on would also be useful, if this functionality is not already available to the operators elsewhere.

Any commands issued via this interface should be logged (in a standard location and standard XML format, and included in the metadata saved in the archive) so that it is possible in post–processing to know the composition of the phased array at all times. This may be most transparently handled by sending messages to MCAF, and allowing MCAF to handle the actual logging. This is a design consideration for the implementors of this GUI.

5.3 Phasing solution calculation and application

The CBE provides the correlator output (spectra, headers etc.) which are intercepted by the TelCal system. TelCal will derive, amongst other quantities, the antenna–based phase corrections necessary for correct antenna phasing. These solutions are passed back to the executor, which distributes them to the correlator by way of adding them to the delay model being propagated.

The calculation of phasing solutions should be performed during scans which are marked with the intent (an EVLA concept) of “`CALIBRATE_AUTOPHASE`”. The solution interval will be one subscan, as discussed above. The solutions are propagated to the executor and applied on all subsequent scans with an “apply last phase/delay calibration” modifier (modifiers are also an EVLA concept) until new solutions are derived. As discussed above, the maximum latency from the end of a subscan with intent “`CALIBRATE_AUTOPHASE`” before the application of new phasing solutions should be 5 seconds or half the subscan duration, whichever is lesser. The application of the phasing calibration is simultaneous in the “normal” correlator pathway (leading to visibilities) and the phased array output, and the value of the corrections shall be noted and stored in the SDM, along with the reference antenna number.

Initially, calibration will be performed on a “per–subband” basis, assuming a point source model at the pointing centre. Since the total available bandwidth is generally much larger than an individual subband, it would be desirable to perform calibration using the entire

EVLA bandwidth, allowing fainter (potentially in–beam) calibrators to be used. Since these calibrators would not necessarily be at the pointing centre, TelCal would be required to understand the use of models for calibration scans, as well as stacking subbands for phase solutions. The observation script generator (SCHED+vex2opt or the manual entries into the OPT) would need to provide this additional information (source model and the intent to solve using multiple subbands for given scans) which must then be passed along to TelCal by the executor.

5.4 Phasing at the baseline board

The antenna calibration phases are applied for both the normal correlation pathway and the phased array output at the baseline board. For the phased array output, this occurs in the RXP chip, while for the normal correlation pathway the phase corrections are performed in the correlator chip. All unmasked antenna bitstreams are phased using the provided delay and phase correction solutions, and then summed. The summed signal is requantized (to 1,2,4 or 8 bits) and sent from the RXP chip to the GigE chip on the baseline board to be sent to the Mark5C recorder(s).

Both the input levels to the phased array sum and the number of antennas being summed will affect the signal level at the output of the requantizer, and hence alter the requantizer gain required to obtain optimal state counts for the quantized phased array sum baseband data. Thus, it is necessary to continuously monitor the power in the quantized phased array sum (which can be measured using a register on the baseline board) and adjust the requantizer gains as necessary to maintain optimal state counts. This is effectively a software automatic level control (ALC) loop. The total requantizer gain is determined by two stages in the RXP chip – a first stage which selects a window from the internal adder tree, which can effectively be considered coarse gain, and a second stage which is the settable requantizer gain.

This power monitoring and requantizer gain control has been implemented in a low level piece of software similar to those used to capture and store the switched power measurements. These values are configured in the VCI file used to set up the correlator. The desired output RMS and time constant can be changed, but the current default values seem to perform acceptably.

The baseline board will produce VDIF formatted data, so no further modification is required and the data can be switched directly to the recorders. The configuration of the baseline boards to send the packets to the correct IP address is discussed above in Section 5.1.

One baseline board will be used for each sub–band that will be phased. Each baseline board can output up to ∼950 Mbps over the 1 Gb ethernet link to the main switch. This can comfortably handle the widest bandwidth (128 MHz) for a single sub–band polarization at the default VLBI bit resolution (2 bits).

As discussed below, each Mark5C will be capable of recording up to 2 Gbps, equal to 256 MHz of dual polarization data at 2 bit sampling. This implies that no more than 4 baseline boards will be required per Mark5C, assuming wide subbands (and hence 512 Mbps per baseline board). Since initially 1 (and subsequently 2) Mark5C units will be available at the EVLA, a total of 8 baseline boards will be sufficient to completely supply the available recording capacity in the most common observing mode. Network connectivity for the baseline boards is discussed in Section 5.6.

In order to provide multiple "Y27" phased array beams, independent baseline boards will be connected to identical inputs but given different phasing solutions. Due to the limited recording bandwidth, this is unlikely to place a strain on the available (ethernet–connected) baseline boards. It is possible to trade off bandwidth against number of phase centers, but the initial implementation (which considers a maximum of two phase centers) will not enter this domain. Similarly, multiple narrower sub–bands could be used, requiring more baseline boards – this will explicitly not be supported where more than 20 baseline boards are required, at least until additional network connectivity is obtained.

5.5 Pulse calibration tone extraction

The station boards of WIDAR are capable of extracting and measuring pulse calibration tones. The most useful proxy for the pulse calibration of the phased signal would be the pulse cal tone measurements from the reference antenna for the observation, or from any one single antenna along with the measurements of the differential delay applied between this antenna and the reference antenna. This could be obtained from the standard EVLA data products produced by MCAF. However, the injection and measurement of pulse calibration signals are not seen as a high priority for phased array observations.

5.6 Networking

The phased array data streams will utilize the same network switch which connects the CBE nodes. This network switch is located in the main correlator room. Data is sent from the baseline boards used for phasing to the recorders using standard ethernet protocols. *the following is out of date – I believe all of the baseline boards have been connected??*. 20 baseline boards will initially be connected to the main switch using copper 1 Gb ethernet connections. Provisions have been made for up to 32 baseline boards to be connected in this fashion before further major networking purchases must be made. The initial 20 boards will allow up to 10 Gbps of data to be sent to recorders – well in excess of the current maximum record rate. The VLBI recorders will be connected to the switch using 10 Gb ethernet, and the switch will aggregate the packets from multiple baseline boards (sub–bands) to a single recorder as necessary.

As discussed in the following section, the 10 Gb connections to the main switch are fiber, but the connection on the Mark5C recorder is CX4 (copper). This has been overcome using a fiber–to–CX4 media converter at the Mark5C recorder.

As discussed in Section 5.1, the networking configuration will be performed by the executor, and in the initial implementation with one Mark5C may be set to default values. The baseline boards will receive the IP address of the recorder to which their VDIF packets shall be sent. The Mark5C machines will simply record every packet they receive.

5.7 Recording

Initially, a single Mark5C recorder will be available at the EVLA. Provisions have been made for a second recorder, but it will not extend the available functionality until either a) two simultaneous beams can be phased from "Y27", or b) subarrays are made available on the EVLA.

The Mark5C recorders will be located in the old VLA correlator room. This decision was made on the basis of minimizing foot traffic into the WIDAR correlator room, as recording at high data rates will require frequent disk module changes. As noted above, this imposes an additional cost due to the media converter required to enable the CX4 connector on the Mark5C to accept 10 Gb ethernet data from the 10 Gb fiber coming from the switch. The cost of this media converter is around $750/unit, and is thus small compared to the cost of upgrading the existing Mark5A units to Mark5C compatibility.

The recorders will be started, stopped and configured using commands from the executor. The Mark5C units require minimal configuration; the selection of a bank to record to (of the two available) and the setting of start and stop times. The Mark5C units will run a "MIB–emulation" program, allowing them to be controlled, and to produce status updates, like any other element of the EVLA system. The Mark5C units will do no filtering/inspection of the packets they receive, and thus the baseline boards will assume the responsibility of sending the VDIF packets only to the correct Mark5C to be recorded.

6 Development plan: requirements, timeframe and test plan

The near, mid and long term requirements for phased array operations are discussed below. At the conclusion of this section, Table 4 summarizes the timeframe over which these requirements will be implemented. This section is intentionally brief and is intended for reference purposes – for the details associated with each development item, see the preceding sections.

6.1 Near term items

The highest priority is the development of a system which can produce a single "Y27" phased array beam. The most straightforward path towards controlling the EVLA for this case is to develop a pathway from the VEX output produced by SCHED. This is absolutely essential for routine VLBI operations, and minimizes the requirements for EVLA resources which are currently under great demand supporting OSRO operations. It also makes time domain studies possible (in a RSRO–like sense) although the configuration path is obviously far from ideal, and not dynamic. Table 1 lists the minimum development which must be completed to enable this mode of phased array operations. This development should be completed on a timeline commensurate with the availability of the Mark5C at the EVLA.

In order to expedite the availability of this mode, it will be acceptable to set some parameters (such as the phasing solution interval, the baseline boards used for phasing, the destination IP address for VDIF packets etc) to default values.

Table 1. Highest priority items to enable "Y27", scheduled by SCHED+vex2opt

Requirement	Subsystem(s) affected
Produce source, scan and config info in a form that can be ingested into the OPT	vex2opt(new)
Default WIDAR setup determined for low latency in phased array mode	vex2opt
New control information in script to be parsed and disseminated correctly	executor **(done)**
Ability to control time and solution interval (scans via intents) for phasing calculation	TelCal
Phasing corrections stored and available in SDM	TelCal, MCAF
Software ALC for phased array sum requantizer gain	low–level software **(done)**
Phased array sum composition can be altered by the operators	New GUI page?
Flags, weather monitor data copied/aggregated from EVLA archive to VLBA loader	evla2vlba (new script)
System temperature monitor data calculated from archive and sent to VLBA loader	evla2vlba (CASA component)
10 Gb networking available (new switch in correlator room, ≥4 baseline boards connected)	networking **(done)**
At least one Mark5C installed in old correlator room, connected to 10 Gb ethernet	recording **(tested, final Mark5C not yet avail**
DiFX software correlator must handle VDIF data in multiple threads	DiFX (correlator) **(done for file-based)**

Table 2. Intermediate–term requirements for "Y27"

Requirement	Subsystem(s) affected
OPT support for single beam phased array observations (phasing scans etc)	OPT, scheduler **(done)**
Tcal based system temperature extraction for a single antenna	evla2vlba
Tcal based system temperature extraction for "Y27"	evla2vlba
VLBI configuration support for multiple beams	SCHED, VEX, vex2script
System support for 2 simultaneous phasing positions	executor, TelCal
Installation of second Mark5C recorder	recording
Control over destination of VDIF packets	OPT, SCHED, vex2script
System support for inbeam phase calibration (source models)	TelCal
OPT support for multiple beams (sources with additional positions)	OPT, scheduler
Preparation support for EVLA–only phased array observing	PST

6.2 Mid term items

Once an initial phased array capability for the EVLA has been obtained, the focus will turn to expanding the capabilities of the "Y27" system by enabling configuration of the array via the normal OPT channel and adding the possibility of recording two beams at different phase centers simultaneously. At this time, it would also be desirable to generalize "hardcoded" values from the initial implementation, and to add the ability to phase on a model other than a point source at the phase center. Both additional functionalities (OPT configuration and two simultaneous beams) should be made available by the end of 2010. Table 2 lists the developments that will be required for this additional functionality, in rough order of priority.

6.3 Long term items

The main long–term goal of for phased array observations at the EVLA is the support of phased arrays on subarrays, and in particular the "Y26 + Y1" option to minimize slewing during fast phase referencing at higher frequencies. However, the specifics of the implementa-

Table 3. Long–term requirements for “Y27”, “Y1”, and other modes

Requirement	Subsystem(s) affected
Support only one subarray being phased (e.g., “Y1”)	executor
Control of delay position of subarray	executor
Configuration support of multiple simultaneous phased subarrays	SCHED, VEX, vex2script
Execution support of multiple simultaneous phased subarrays	executor
Post–processing support of multiple simultaneous phased subarrays	evla2vlba

Table 4. Milestones for phased array development

Milestone	Completion date
All hardware & basic software in place for phased array observations	??
EVLA–only test observations in phased array mode	done
VLBI test observations including phased EVLA	done
Completion of monitor data pipeline for phased EVLA	??
Release of phased EVLA for RSRO–style observing	??
Completion of OPT support for phased EVLA	done
OSRO availability	??
“External” support for two simultaneous phased EVLA beams (SCHED, VEX...)	Q4 2011
“EVLA” support for two simultaneous phased beams (OPT+correlator)	Q4 2011
Support for multiple phased arrays in EVLA subarrays	Q1 2012

tion of any subarrayed phased array observations will be heavily dependent on the measures in place to support subarraying generally at the EVLA.

The majority of the work to support split phased array operations will be common to the work of supporting any subarrays at the EVLA and as such will not be considered here – working subarrays is considered a prerequisite to beginning work on these long term goals. If only one subarray is being phased (e.g., the “Y1” case) no additional work is required beyond the correct propagation of antenna masks to the baseline boards. At this time, control of the location of the delay centre of the (sub)array is desirable.

Phasing multiple subarrays simultaneously is more challenging, since resources such as Mark5C recorders must be shared. More information must also be provided via the configuration path (SCHED, VEX), and the logging of monitor and calibration data via evla2vlba is also complicated considerably. Table 3 summarizes the long term requirements for phased EVLA observations.

6.4 Timeframe summary

Table 4 summarizes the milestones which have been set for phased array operations with the EVLA and their intended completion dates.

Table 5. Test plan summary

Target date	Description	Subsystems with new functionality tested	
1 April 2010	On–chip test of RXP phase rotation	RXP chip	done
15 April 2010	EVLA–only test on bright pulsar	Executor, correlator, TelCal, recorder	done
1 May 2010	VLBI test on bright calibrator pair	SCHED/vex2script, evla2vlba	outstanding
1 July 2010	VLBI test using a variety of bandwidths	SCHED/vex2script, correlator, TelCal	outstanding
1 December 2010	OPT test on bright pulsar	OPT	done
1 March 2011	VLBI test with in–beam calibrator	SCHED/vex2script, executor, correlator, recorder	outstanding
1 May 2011	EVLA–only test of pulsar pair	OPT	outstanding
1 November 2011	Phase calibrator model test	Executor, TelCal	outstanding
1 February 2012	“Y26 + Y1” test	SCHED/vex2script, executor, correlator	outstanding

6.5 Test plan

A sequence of test observations are proposed below to verify the correct operation of the EVLA in phased array mode. They are summarized in Table 5 and described in detail below.

6.5.1 “On–chip” tests of the RXP phase rotation

These initial tests will be hand configured and are intended to ensure the correct implementation of the phase rotation in the RXP chip. No VDIF packets will be produced, and so it is not necessary for the Mark5C to be available.

The first of these tests will “phase” a single antenna whilst observing a strong, narrow spectral line source. The antenna should be chosen to have the maximum possible “fshift”, which is the net frequency offset imposed by the delay model, LO offset etc. The fshift should be several times larger than the width of the spectral feature being observed. The output of the RXP chip will be manually routed to the correlator chip array, and the resultant autocorrelations will be inspected for correct frequency placement and absence of any spectral ringing.

The second of these tests will use both RXP chips on a single baseline board to “phase” a single antenna. The outputs of the two RXP chips will be routed to the correlator chip and cross–correlated, but the phase rotation in the correlator chip will be disabled. This output will be compared to the visibilities for this baseline produced by a second baseline board which correlates all antennas in the usual way. If the phase rotation applied in the RXP chips is correct, the two sets of visibilities should be identical.

6.5.2 Initial EVLA–only bright pulsar/maser test

The first systems–level test will utilize a strong pulsar or maser to ensure that configuration of the system in phased array mode can be performed correctly, that the phasing works as expected and that the recording system is functioning properly. No more than 30 minutes will be required, and 256 MHz of dual polarization data will be recorded (in four 64 MHz bands, requiring 4 baseline boards).

This test can be scheduled using SCHED or the OPT, even though it will be an EVLA–only experiment. It is acceptable if some hand–editing of the observe script is required. The array shall be phased on the calibrator B0355+508 and then pointed at the bright pulsar PSR J0332+5434, or the maser W3OH. PSR J0332+5434 has a peak flux density of several Jy at 1400 MHz and a pulse width of 10s of milliseconds, and thus single pulses from this system will be easily visible to even a single EVLA antenna. Likewise, W3OH has a peak flux density of $\sim$60 Jy and will be easily visible on a single antenna. In either case, the accuracy of the phasing on the target can be estimated by imaging the regular EVLA dataset.

Since the pulsar emission will be subject to scintillation (the scintillation time and bandwidth at 1400 MHz are 10 minutes and 13 MHz respectively) it is not feasible to simply compare the strength of the on and off pulse signal in the accumulated phased array output as a measure of whether the sensitivity of the phased array signal is as expected. Accordingly, antennas will be removed from the phased array sum over a timescale much shorter than the pulsar scintillation time. The resultant decrease in S/N of the pulsar signal will be compared to the theoretical expectation. A reasonable approach would be to reduce the phased array membership from its initial value (expected to be of order 24 antennas) to 12, 6, 3 and finally 1 antenna at 10 second intervals. In the event that the operator GUI for altering the phased array sum is not yet operational, the RXP chip GUI can be used (requiring knowledge of which baseline board is being used, and the mapping of antennas to baseline board inputs).

This approach can also be taken with the maser scans, although in this case the timing of the addition/subtraction of antennas to the phased array sum is not critical.

All available antennas would then be restored to the phased array sum and the procedure repeated. Several repeats would give additional robustness against scintillation effects. The array should then be rephased on the calibrator, and the pulsar/maser sequence repeated.

An important attribute which the VDIF output data should possess is reasonable state counts for the 2 bit data. It is possible that this test may be run before the software ALC for the phased array sum is complete – in this case, the state counts may not be optimal. If the software ALC is not complete before this test is run, the state counts should be verified after the ALC completion on one of the later planned tests, or else this test should be repeated.

6.5.3 VLBI test of single beam

Building on the initial test, a short VLBI observation is required to verify the correct operation of the VLBI observation configuration pathway (SCHED + vex2script) and the post–processing of monitor data (evla2vlba). Naturally, this provides a means to obtain a more accurate estimate of the correctness of the phasing. As with the previous experiment, a short amount of data (no more than 30 minutes) should suffice. At least 3 other VLBA antennas should be scheduled in order to obtain closure quantities – it would be preferable to schedule the entire VLBA (at least 8 antennas, allowing for maintenance or other downtime) to increase redundancy.

A suitable VLBI calibrator which is also an acceptable EVLA calibrator in the given array will be used as the phase calibrator, and the schedule will include several scans on this source interspersed with scans on other nearby VLBI calibrators. One scan will be made on a primary EVLA flux calibrator to provide evla2vlba with all the information it requires to produce system temperature monitor data. The frequency of the test is unimportant – L, S, C or X band would be acceptable.

6.5.4 Test different bandwidths

Extending the previous test, another VLBI observation will be used to ensure that all supported bandwidths (1, 2, 4, 8, 16, 32, 64 and 128 MHz) can be phased and recorded correctly. At this stage, it is not expected that the EVLA can transfer calibration from a previous wider bandwidth subband to a narrower subband – the phasing solution will be repeated for each bandwidth recorded. Due to the number of different bandwidths to be tested, this may take somewhat longer than the preceding test – up to 90 minutes may be required.

6.5.5 EVLA–only bright pulsar/maser test scheduled with the OPT

This test will be identical to the first bright pulsar/maser test, but should be entirely scheduled using the OPT. No hand–editing of any files should be required. The reduction and analysis will be identical to the previous pulsar observation.

6.5.6 Two–beam VLBI test

This test will be used to verify that two phased array beams can be generated and recorded simultaneously. If no additional Mark5C recorder has been installed at the EVLA at the time of this test, it should be repeated after the installation of the second Mark5C recorder. The test will be short (no more than 30 minutes required) and will be conducted at 1400 MHz. As with the other VLBI tests, a minimum of 3 VLBA antennas are required, but all available VLBA antennas would be preferred.

The target for this test will be a pair of VLBI calibrators which lie within half a primary beam at 1400 MHz. Suitable calibrator pairs include:

- J0449+1121 and J0448+1127
- J1153+4931 and J1152+4939
- J1310+3220 and J1310+3233
- J1344+6606 and J1343+6602
- J2254+0054 and J2253+0105

The EVLA will be phased on a suitable nearby calibrator and then observe the target for 5 minutes at each of three different pointing centers (calibrator 1, calibrator 2, and their midpoint) before rephasing. Each phased array beam will be correlated with the VLBA antennas in a separate correlator pass and analyzed individually in an identical manner to the earlier VLBI tests.

6.5.7 Two–beam EVLA–only test

A similar test to the two–beam VLBI test will be required to verify the configuration of the EVLA for dual–beam operation using the OPT. A similar methodology to that described above will be employed, but using pulsars in place of VLBI calibrators. Suitable pairs of bright pulsars are rare at northern declinations – the most appropriate target pair is PSR J1954+2923 (8 mJy) and PSR J1955+2908 (1.1 mJy).

6.5.8 Inbeam phasing test

This test is intended to verify that phasing on an in–beam calibrator is possible. There are two main components to this mode of operations – the use of a model allowing calibrators not at the phase center to be used for phasing, and the transfer of calculation of phase solutions using stacked subbands. The former requirement is obviously required for maximum pointing flexibility when using an inbeam calibrator; the latter is needed to allow the use of fainter calibrators.

An identical observational setup to that used for the two–beam VLBI test may be used for this test. A field should be chosen with one bright VLBI/EVLA calibrator and one fainter VLBI calibrator (to simplify the model that will be used for the EVLA phasing). As with the two–beam VLBI test, no more than 30 minutes will be required, and the EVLA should be phased on a known good external calibrator before beginning the observations of the target field containing the "inbeam" calibrator.

6.5.9 Initial phased subarray testing ("Y26 + Y1")

The first test of the phased subarray capability of the EVLA will be to ensure that the "Y26 + Y1" mode of operations works correctly. This test would replicate the earlier single beam VLBI test, but with the single Y1 antenna remaining on the primary EVLA phasing calibrator for the duration of the experiment. In addition to the correct functionality of the online EVLA systems, this test will ensure that the extended functionality in the SCHED/vex2script configuration route and evla2vlba post–processing is working correctly.

References

Moran, J. M., & Dhawan, V. 1995, in Astronomical Society of the Pacific Conference Series, Vol. 82, Very Long Baseline Interferometry and the VLBA, ed. J. A. Zensus, P. J. Diamond, & P. J. Napier, 161–+

Morgan, T., Ryan, K., Sowinski, K., & Waters, B. 2004, EVLA High Level Software Design, Tech. rep., NRAO, www.aoc.nrao.edu/evla/admin/reviews/software/EVLADesign.pdf

the Mark5C collaboration: NRAO/Haystack. 2008, The Mark5C specification, Tech. Rep. http://www.vlba.nrao.edu/memos/sensi/sensi12.pdf, NRAO/Haystack

Whitney, A., Kettenis, M., Phillips, C., & Sekido, M. 2009, in Proceedings of the 8th International e-VLBI Workshop. 22-26 June 2009. Madrid, Spain Published online at http://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=82., p.42, 42–+