VLBA Sensitivity Upgrade Memo 50

Using the new fast VLBA fiber networks

Walter Brisken, Jay Blanchard & Barry Clark

04 Mar. 2020

Abstract All VLBA sites will soon have fiber network connectivity to the DSOC in Socorro, NM, at speeds of 200 Mbps or higher. This document is meant to explore the various use cases for this new infrastructure and to serve as a near-term roadmap for development.

1 Background

In its early days, data connectivity to VLBA sites was limited to modem speed, approximately 65 kbps. This was sufficient for monitor and control. A special mode was implemented to allow fringe checks to be performed using very small amounts of data. Between 2005 and 2010 these links were updated to a capacity of 1.4 Mbps, allowing for the increased monitor data rates from the RDBE system to be accomodated and for most software updates to be sent over the network rather than on physical media. One notable use of these links was the demonstration of near-real-time correlation of two spacecraft orbiting Mars. Between 2010 and 2018 some further improvements were made. USNO funded deployment of 1 Gbps service to Pie Town and Mauna Kea to support daily low latency determination of UT1-UTC. The Owens Valley antenna was connected to 10 Gbps capable network, deployed as part of the "Digital 395" project; initial service commenced in 2015 at 20 Mbps with later upgrade to 200 Mbps.

As of this writing, a project to raise network speeds to 200 Mbps at all sites is nearly complete. New infrastructure being deployed should be usable at up to 10 Gbps, however service costs are too high at this point in time to justify such speeds.

1.1 UDP and TCP

When using the Internet to transmit data, usually one of two low level protocols is used. User Datagram Protocol (UDP) is the simplest. Packets formed at the source are simply transmitted with a destination. Packets received at the destination are guaranteed to be complete, but they are not guaranteed to arrive in the order sent and there is no guarantee they are delivered at all. On a well constructed, closed network operating within bandwidth and packet rate limitations, the packets should not be lost and should arrive in order. Transmission Control Protocol (TCP) ensures packet delivery in the order that packets are sent through use of a return channel that acknowledges receipt of packets. TCP makes it simple to create reliable network streams but due to the round-trip communication, networks with high latency will result in reduced throughput. TCP has the advantage of adapting its throughput in response to other traffic by reducing packet rate when packet loss is detected. Generally speaking, UDP-based data transmission mechanisms will be faster than TCP-based ones, but such systems either need to have additional packet loss detection mechanisms or be tolerant of packet loss.

Sent packets that exceed the Maximum Transmission Unit (MTU) in size are split into multiple packets and reassembled at the other side. This fragmentation process increases the likelihood of lost packets as every one of the split packets must arrive at the destination for the packet to be complete. Typical MTU values on long-haul networks are about 1500 bytes. The long-haul links on the VLBA have MTUs that range from 1372 to 1472 bytes.

2 Current status

2.1 Fiber status

The NSF provided funds (here called "CSA-F"; see section A) for the VLBA to deploy high-speed networks to each of the VLBA antennas and to operate these links for a limited period. At the end of CY2019, the operational bandwidth at the VLBA stations was as follows:

Site	Location	Bandwidth (Mbps)
BR	Brewster, WA	200
FD	Fort Davis, TX	200
HN	Hancock, NH	1.4
KP	Kitt Peak, AZ	1000
LA	Los Alamos, NM	1.4
MK	Mauna Kea, HI	1000
NL	North Liberty, IA	200
OV	Owens Valley, CA	200
PT	Pie Town, NM	1000
SC	St. Croix, USVI	200

Connectivity at LA and HN is expected to be brought up to the 200 Mbps level by end of June and August, 2020, respectively. CSA-F funds are expected to allow continued operations at the above bandwidths through the end of August, 2023.

2.2 Recent and established uses of the networks

2.2.1 Monitor and control

The primary use of the data links between Socorro and the VLBA sites to date has been exchange of monitor and control files to support operation.

2.2.2 USNO daily observations

Mauna Kea and Pie Town have had 1 Gbps connectity for several years now in support of the USNO's need for low latency UT1-UTC observations. These observations continue to run daily and typically transfer 300 GB of data per station directly to the USNO correlator in Washington D.C. This data transfer typically makes use of the `tsunami` program, which is a UDP-based transmission protocol.

2.2.3 Near real-time fringe checks

In late 2017 a suite of tools called `nrtfringediagnostic` was developed to allow for Near Real-Time (NRT) evaluation of coherence and timing. These python tools are used in sequence to generate a low bandwidth observing file, transfer the data to Socorro, correlate it, and clean up the bulky data. This has been routinely used after maintenance periods as an additional check that the VLBA antennas are operating correctly. Its operation is limited to the narrowest bandwidths (typically dual-polarization at 1 MHz bandwidth) due to data transfer limitations; most VLBA antennas were using 1.4 Mbps connections when this was implemented. The data transfer makes use of secure copy (`scp`), a TCP-based encrypted data transmission protocol which is not appropriate for high-speed, long-haul data transport due to the impact of latency on TCP throughput.

2.2.4 Network monitoring

A widely-used network monitoring tool, iperf[1], has been deployed in server mode on the x-cube units at each VLBA site. Every 15 minutes, when enabled, an iperf client program is started from a machine in Socorro

[1] See `https://iperf.fr/` .

and is used to measure the UDP link capacity, in both the sending and receiving directions, to each VLBA site. This tool collects the results in a database and has a web tool that can be used to show performance history. This tool has been valuable in testing and accepting the upgraded networks and in monitoring them and detecting performance regressions. Because this monitoring process makes full use of the network bandwidth it can interfere with actual use of the network links. Some coordination between use of links and monitoring will need to be developed.

2.2.5 Data transmission from Mauna Kea

In July, 2019, the VLBA site technicians at the Mauna Kea site had their access to the site blocked by protestor activity. During this month-long period during which no data modules could be shipped to Socorro, the 1 Gbps data link was used to bring back as much data as possible to minimize correlator backlog and extend recording capacity. At the time the link was operating at only 650 Mbps due to hardware limitiations. While productive use of the data link occurred, it was even more valuable as a learning experience:

- At 2 Gbps and lower record speeds, data transfer can read from the Mark6 module at the same time data is being written.
- At 4 Gbps, a module cannot be read while recording, but different module in the same Mark6 unit can be read for data transmission. Careful coordination between operations and data copy management was required.
- `jive5ab`[2] was used for data transmission. This program makes use of UDP-based Data Transfer Protocol (UDT[3]). This program worked very well for copy from a Mark6 module to a filesystem in Socorro. Interrupted data transfers were able to be resumed with minimal data transfer redundancy.
- `jive5ab` was also used for copy from remote module at Mauna Kea to local Mark6 module in Socorro. This worked well, but with the exception that data transfers resuming is not possible with the currently deployed version of `jive5ab`.
- Frequent monitoring of the network using iperf was an enormously beneficial tool. At the times when this tool runs the data rate drops significantly due to fierce competition for bandwidth.
- With UDP data transfer, choosing a MTU that is appropriate is critical. A slightly over-sized MTU will drastically reduce performance. A slightly under-sized MTU will slightly reduce performance.

Since mid August, when the access to Mauna Kea was restored, a new router and switch were installed that resulted in reliable data transmission speeds at 950 Mbps, effectively the advertised link speed. The reliability of the link was improved as well. Also since this time a Mark6 expansion chassis was installed at Mauna Kea allowing for up to four modules to be mounted at one time. This allows for significantly improved logistics and simpler coordination between observing and data transfer.

2.2.6 Diagnostics

Finally, on occassion, some recorded baseband data is hand-transferred to Socorro for diagnostic purposes. In these cases `scp` or `m5copy` (controlling `jive5ab`) is typically used.

2.2.7 Software updates

Links from the DSOC in Socorro to VLBA sites are used to update software. The previos standard 1.4 Mbps throughput allow up to about 10 GB per day, which is sufficient to transmit even large software updates, but over a time period which could be hours to days long. Higher speed networks won't change the approach to software updates but will make them significantly quicker.

[2] See https://www.jive.nl/~verkout/evlbi/jive5ab.html .
[3] See http://udt.sourceforge.net/ .

3 Use cases at 200 Mbps

200 Mbps is the target data rate for the developments and services funded by CSA-F. In practice, due to various overheads, a throughput of 195 Mbps is realized and should be considered the actual ceiling data rate. This data rate is about 22 times lower than the current peak data acquisition rate of 4096 Mbps when VDIF framing overheads are considered. The average VLBA data rate is currently a factor of about 3 lower than this, but as 4 Gbps observing becomes more routine, the average data rate could nearly double.

The following use cases are well suited for data transmission at the 200 Mbps rate.

3.1 Interferometric pointing

A very compelling use case for moderate data rates is interferometric determination of the VLBA pointing equation. As of now, the VLBA exclusively uses total power for determination of pointing offsets. In this mode a 10 point pattern is executed; 5 of the pointing positions are required for the fit, and the other 5 are used to determine gradients in elevation and time of the background level. In interferometric pointing (which is exclusively used at the VLA), only the 5 fit pointing positions are required as the interferometer is insensitive to the background. Interferometric pointing would greatly decrease the brightness threshold for the pointing sources and stands to yield much more robust solutions with increased diagnostics. In the case of spectral line pointing, no special line-free channels need to be recorded.

Interferometric pointing comes with the following challenges and required developments:

1. A tool to generate interferometric pointing observing files will be needed.
2. A source list appropriate for interferometric pointing will need to be developed (maybe start with the EVN fringe finder catalog?)
3. An operational real-time correlation setup will need to be developed.
4. A pipeline that processes the correlator output and generates a `.fit` file will be needed.
5. At least 3 antennas must be in a subarray for unique gain determination with a fourth antenna desired for robustness and diagnostics. This will complicate capture of pointing information over some parts of the sky for the antennas most distant from the array center.
6. Perhaps two pointing modes will be required: one where all antennas execute the same pattern at the same time and another where one antenna stays fixed on target with the other(s) executing the pattern. Experimentation may be required to determine which model works best.

An extension of interferometric pointing that will be of great value is interferometric determination of the FRM parameters: focus and rotation.

It should be noted that determination of antenna gains will still require single dish pointing.

3.2 Inteferometric referenced pointing

Referenced pointing operates in largely the same manner as off-line pointing, but with the requirement that pointing offsets are solved and applied within a short (minute timescale) period of the measurement. Such pointing measurements must be robust and reliable as the sensitivity of subsequent observations depends on accurate results. In addition to the development considerations for interferometric off-line pointing, one would need to consider the following:

- A mechanism to turn data transmission on and off as pointing sources are observed will be needed.
- A mechanism to prepare the correlator for correlation of pointing scans will be needed.
- Software to determine pointing corrections from correlated data without human input will be needed.
- A mechanism to convey the pointing updates to the executors running at each site will be needed.

3.3 Real-time fringe check

As noted in section 2.2.3, some software infrastructure already exists to perform some low latency fringe diagnostics on the VLBA. This `nrt` toolset is used on a weekly basis after site maintenance to ensure that the full analog path is complete, that oscillators are stable, that timing is within operational limits and that receiver polarizations are correct. This code was developed when typical network speeds were two orders of magnitude smaller than they are now. A new real-time replacement for this would allow for considerably higher bandwidths, longer tests, and less data management hassle.

3.4 Low latency correlation of data quality observations

Each week, usually following maintenance periods, a data quality observation (dq series) is performed. This observation observes with each receiver band and exercises the S/X dual-band capability. These observations exercise a wide range of bandwidths, currently up to 2 Gbps, exceeding the capacity of the fiber networks for each station. However, a post-observation data transfer at 200 Mbps could be performed in approximately 5 hours, allowing next-day confirmation of performance at all observing bands. A lower bandwidth or shorter duration version of these data quality observations could allow for much faster data transfer and lower latency.

An initial demonstration of such a capability has already been developed for exploratory purposes. This was performed with a coordinating script which takes an experiment code and scan number and produces correlator output as quickly as it can given data transfer speeds.

3.5 Low latency correlation of all fringe finder sources

Almost all VLBA schedules contain a scan on a very bright source every few hours which are often used as bandpass calibrators, but which also can serve as diagnostic "fringe finder" sources, which can be critical at correlation time if there are large unexpected clock offsets. A system could be developed to transmit part or all of each fringe finder scan back to the VLBA correlator for automatic correlation and reporting to operations.

4 Use cases at 1 Gbps

As of writing, three VLBA antennas have access at a 1 Gbps. Pie Town and Mauna Kea are connected at this rate in support of the USNO's need for low latency daily observations on one baseline. The Kitt Peak antenna was connected at 1 Gbps as part of CSA-F due to marginal or no cost increase compared to 200 Mbps access.

Because of the different data products and timescales for data transmission, not all of the following may be simultaneously operational at a 1 Gbps data rate.

4.1 Real-time science observing

Fast Radio Burst detection experiments could be run on live streaming data. The underlying mechanism for searches of this kind (which were first implemented by the VFASTR project) is to look for coincedent dispersed pulses. VFASTR makes use of correlator intermediate products for this search. A dedicated pipeline would have more control over search parameters and could stand to be more effective (albeit at lower bandwidth).

A more advanced version of this capability could have spectra formed at each antenna within a local computer. One would want millisecond time resolution and 100 kHz spectral resolution, demanding about 80 Mbps data transfer rates for full bandwidth observing, which would be conceivable to implement using the 200 Mbps connections.

4.2 Reduced reliance on disk transport

This use case is a generalization of that described in section 2.2.5. Stations with network connectivity at the 1 Gbps level could make use of the network to reduce by about half the amount of data that needs to be shipped to Socorro. The benefits would include less wear and tear on disk modules, reduced servicing of Mark6 units at the VLBA sites, and potential for lower average latency of correlation. This could be implemented in various ways. One model would be to have two modules loaded in the Mark6 expansion unit at the VLBA site. Operations would always be downloading data from one of the modules in this expansion unit and would alternate recording on a project-by-project basis between the other module in the expansion unit and the modules in the primary Mark6 unit based on module usage. Data would be transferred either to a disk array or a Mark6 module at the correlator and would be deleted at the station end after the data transmission has been verified.

4.3 Continuous transmission of at least one VDIF thread

The VLBA can currently generate data in one of two data formats: Mark5B when using the PFB personality (at 2 Gbps only) and VDIF[4] when using the DDC personality. VDIF is the presumed only format to be used in next generation backends. With the VLBA's use of the VDIF format, data is organized so that each packet of data contains a time series of data from one channel which is identified through its "thread id". This organization of data makes it very easy to select one or more baseband channels and even to send data from different channels to different destinations. Currently the VLBA is limited to 2 bits per sample and up to 128 MHz of Nyquist sampled data per baseband channel, which amounts to 512 Mbps of data (excluding framing overhead). With a 1 Gbps link, at least one full baseband channel could be sent in real-time to Socorro allowing for real-time diagnostics. Observations employing smaller channel bandwidths could stream more channels back.

Various diagnostics could be performed with such data streams. The most obvious is to cross-correlate in real-time and perform a fringe fit when a calibrator is being observed as a continuous diagnostic of array performance. Alternatively, real-time spectra could be formed as a mechanism to monitor interference or switched power could be demodulated as a means of detecting timing errors.

If the `vlitebuf` utility is used as the receiver at the correlator, many different processes could happen simultaneously without the need for redundant data transmission.

5 Use cases at 5 Gbps

With a network bandwidth of 5 Gbps, all currently implemented VLBA observing modes could be transferred in real-time to the correlator. There are ambitions to raise the VLBA data rate so a 10 Gbps or faster connection would be required in the future to achieve full-time real-time operation. At the time of writing, no plans exist to develop connectivity at this capacity, however it remains desired. As such, only a few high level concepts are described. Should full bandwidth connectivity become realistic, these use cases and others will be reconsidered.

5.1 Single antenna at 5 Gbps

With one VLBA antenna connected to the correlator (or more, but not quite all), the data recorder(s) could be moved from the antenna(s) to the correlator. This could allow the current operations model to continue but without module shipment from the antenna site(s).

5.2 Full VLBA at 5 Gbps

Once all of the VLBA antennas are connected at their peak data rate, real-time correlation can become reality. A likely issue related to this is that aggregating the full bandwidth from the array, which would be

[4] See https://vlbi.org/vlbi-standards/vdif/ .

in excess of 40 Gbps, into Socorro may be a challenge. However, it could well be that the VLBA correlator could be moved to an equipment rack near a networking hub, which could be anywhere in the US.

5.3 VLBA antennas to join VLA antennas

High speed data links could allow data from a subset of VLBA antennas to be fed into the VLA's WIDAR correlator. This would amount to a "super Pie Town Link". Having four or five of the inner VLBA antennas join the VLA would allow some ngVLA scientific pathfinder experiments to be executed. Considerable engineering effort on the WIDAR correlator would be required, something that is outside the scope of this prioritization process. The concept is noted here for completeness. It should be noted that considerable investment in new electronics and fiber links capable of 64 Gbps would be required for the VLBA to match the bandwidth of the VLA when when observing at the VLA's widest supported bandwidth.

6 Technical approaches and limitations

This section describes the promising technical approaches that could play a role in enabling the use cases described above. Given that the upgraded links have not been exercised much to date and the software tools are new to us and evolving, it could be that final implementation deviates substantially in some ways.

6.1 Greater throughput to VLBA correlator

At the time of initial writing there was a network bottleneck between the 10 Gbps DSOC external network interface and the VLBA correlator that in real-world sitations leads to a maximum aggregate incoming data rate of 1 Gbps. It seems this situation can be remedied by "channel bonding" multiple 1 Gbps interfaces to increase the throughput. This is a rather inexpensive and simple fix that should increase the connectivity to the 3 Gbps level.

Since the first drafts of this memo were circulated, the channel bonding has been implemented, but as of this Jan 21, 2020 version has not yet been tested.

6.2 VDIF data over UDP

The VLBI Data Interchange Format (VDIF) is well suited for transmission over Internet. In the case of UDP transmission some care must be taken to avoid unnecessary packet loss due to fragmentation. The RDBE's that generate VDIF data at the VLBA and at GBT always generate data frames with 32 bytes of header information and 5000 bytes of data. These 5032 byte frames are much larger than the MTUs. Experimentation over the VLBA links has demonstrated that pre-emptively reducing the data frame sizes to fit within an MTU greatly increases transmission reliability. Because of the requirement that VDIF frames be a multiple of 8 bytes in length, the only option is to divide each 5032 byte frame into five 1032 byte frames (32 bytes header and 1000 bytes of data).

6.3 Data transfer using jive5ab

As noted above, `jive5ab` uses a UDP-based data transmission called UDT which is both fast and robust against dropped packets due to additional logic that detects dropped packets and requests re-transmission. `jive5ab` is a very general program that purports to "take data from some source and put it somewhere else". It is the program that the VLBA runs on its Mark6 recorders to perform the data recording. A running instance of `jive5ab` runs as a service and takes commands over a TCP interface exposed at a user-specified port number. Discussions to date have centered on using `jive5ab` for non-real-time[5] data transfer from VLBA sites to the correlator in Socorro. In order to isolate the data copy process from the recording process, a second instance of `jive5ab` will be started on each of the Mark6 recorders and an instance will

[5] `jive5ab` may be used for real-time data transfer as well, but this has not been explored in any depth at this time.

be started on each of the correlator servers and Mark6 units, many of which have internal storage that can be used to store temporarily the transmitted data. Once these instances of `jive5ab` are running, a separate control program (such as the `m5copy` program that comes with `jive5ab`) can be used to start, monitor, and stop these data transfers.

6.4 Real-time data transfer

While `jive5ab` may ultimately provide real-time data transfer service, another path using NRAO-developed tools may serve as a more convenient and versatile starting point. A demonstration of real-time VLBA correlation was made using correlator infrastructure developed for the VLITE process. At the station end, a small utility `raw2udp` runs on the x-cube unit, receiving VDIF frames produced by the RDBEs. This program then retransmits the data via UDP packets to an instance of `vlitebuf` running on one of the software correlator servers. `vlitebuf` receives the data and exposes it as a file through the Linux FUSE system. Because the data looks like a file, the correlator can proceed to correlate as if the data were a real file on the filesystem.

Demonstration of this approach indicates that a UDP-based data transport system at rates up to (and maybe exceeding) 128 Mbps is practical. Some hurdles were presented that any production system will need to address:

- The 5032 byte data frames were fragmented upon send, resulting in a low rate (1% or so) of dropped frames at 64 Mbps, but a rather high rate (20%) when operating at 128 Mbps. The solution was to split the 5032 byte frames into 1032 byte frames as noted above.
- Many processes need to be running in concert. Failure of any one could cause correlation to drop an antenna or fail completely. Some sort of process management needs to be introduced. In a fully deployed system, it would be ideal if all software could be started, in a dormant state, when each computer boots. A managing process would then configure and activate the various code when appropriate.
- The data transfer process should not be running when data rates exceed link capacity. Otherwise monitor and control of the antenna will be compromised.

The `raw2udp` program has the option to select a subset of VDIF threads (corresponding to baseband channels) to be transmitted, and each thread can be sent to a different process at the correlator. This flexibility may allow for some useful capabilities and real-time diagnostics.

6.5 Downsides of real-time correlation

Real-time correlation will come with some limitations:

- Predicted earth orientation parameters (EOPs) will need to be used in lieu of the final values which are typically available within about 5 days of observing. There is a task in AIPS that can correct for this.
- Extrapolated, rather than measured, clock offsets will be used. The remedy and consequences are similar to the EOP case.
- Multiple-pass correlation is not possible. The advent of the multiple phase center mode within the DiFX correlator obviates much of this need, however, there are cases where correlator passes with different spectral resolution are needed or where iteration is needed. It should be noted that it could be possible to both record data (at the sites or at the correlator) and process in real-time, which could provide a means for multiple-pass correlation.
- Some correlator modes that require heavy processing (e.g., massive multiple phase center) or high output data rates (sometimes seen with very high spectral and frequency resolution) will not be able to keep up with the incoming data rate.

- Pulsar gating requires up to date pulsar timing parameters that are often available only after an observation has concluded.
- Participation in global, GMVA, and IVS experiments where non-VLBA antennas are included could be problematic for multiple reasons including: 1. need to aggregate in real time data from all the antennas into the same correlator facility, 2. need to have accurate clock offsets for antennas that may not frequently participate in VLBI, and 3. manual editing of correlator configuration files is often needed for non-VLBA antennas to accomodate setups that are not reflected in the `.vex` file.

These limitations will not affect the utility of real-time diagnostic correlation but will need to be considered before the VLBA moves to an entirely real-time operating mode.

6.6 Integration of data transfer into VLBA datacopy software

The VLBA datacopy utilities are a collection of Python scripts that orchestrate and track movement of VLBI baseband data. They were initially developed for reading data for a specified VLBA project from a Mark5 or Mark6 module to a collection of filesystems attached to the correlator and then possibly from those locations to external disks. The advent of high speed data transfer directly from VLBA sites to these filesystems leads to a natural extension of the VLBA data copy system. The database that tracks the data is useful in creating filelists for correlation and for finding and removing data after it is no longer needed. A new utility within this collection (to be called `dctransfer`) would likely make use of `jive5ab` running on the Mark6 recorders and on the correlator computers and Mark6 units. An update to the `dcupdate` utility (which updates the database with knowledge of data locations after an observation completes) would be needed to indicate the location and properties of the files to be transferred.

6.7 Network parameters in vlbaparm database

In order to allow for improved interoperability of various bits of software and to maintain in a central location many of the relevant network parameters, it would be useful to add to the vlbaparm database the following parameters for each antenna's data link:

- The nominal link speed
- The MTU for the link
- The default destination server running `vlitebuf`

7 Priorities going forward

Based on the VLBA technical meeting discussion held on Jan 21, 2020, the highest priority developments were deemed to be (starting with the most urgent):

- Networking parameters relevant for real-time and near-real-time operation into `vlbaparm` database
- Operational real-time correlation
- Operational near-real-time data transfer and correlation
- Real-time interferometric pointing

Top level requirements for each of these developments are described below. Also listed are some thoughts on how to proceed based on limited experience to date.

Formal requirements gathering and status updating of these four capabilities will be performed within JIRA at https://bugs.nrao.edu/browse/VLBASYS-756.

7.1 `vlbaparm` database parameters

The `vlbaparm` database stores information relevant for each VLBA antenna, such as parameters for the pointing model. It was decided to add parameters related to use of the high-speed fiber networks to this database as well, and to encourange all software making use of the networks to use parameters from the database.

Param. name	Units	Description
ETRANSFERMTU	bytes	Maximum UDP packet size to send
UPLOADRATE	Mbps	Maximum transfer speed (correlator to station)
DOWNLOADRATE	Mbps	Maximum transfer speed (station to correlator)
ETRANSFERRATE	Mbps	Maximum download speed to use
REALTIMEMAXRATE	Mbps	Maximum data rate to allow in real-time
REALTIMERATE	Mbps	Current portion of total bandwidth allocated to real-time
REALTIMEHOST	(string)	Hostname of computer at correlator to receive from this station
REALTIMEPORT	(integer)	UDP port number on receiving computer
REALTIMEPATH	(string)	Default storage location for data to be e-transfered to correlator
JIVE5ABPORT	(integer)	Port number for control of data transfer jive5ab instance

The `entity` associated with each parameter will be the two-letter antenna code. Some of the parameters deserve a bit more explanation:

ETRANSFERRATE This parameter indicates the nominal bandwidth assignment for non-real-time transfers. It should generally be less than the `DOWNLOADRATE` to allow for other communication processes, such as monitor and control, to occur.

REALTIMERATE This parameter should be set when any real-time transfer is to begin, and when it ends, providing information to any other process that may start up about the bandwidth reserved for real-time operation. This is seen as a future-looking concept which may or may not be implemented in the database.

REALTIMEPATH This represents the filesystem path containing data downloaded from the station in question. To retain convention used by the VLBA Datacopy software, a subdirectory called *stn_proj/* should be created to contain the data, where `stn` is the name of the station, and *proj* is the project code, including any epoch codes. For example, `br_bb388o`. Management of data within this filesystem is not within the purview of this exercise; methods for handling the data will need to be developed and evolve through use.

JIVE5ABPORT `jive5ab` is the software used to record data on the Mark6 units at the site. The program can also be used for a host of other data transfer purposes. It is believed at time of writing that running an independent copy of `jive5ab` at each site with a matching copy on the respective `REALTIMEHOST` computer at the correlator could form an effective means for non-real-time data transfer. Ultimately this could also form the basis of the real-time data transfer.

A command-line utility should be developed to query these parameters in a structured manner and allow updated to them. Two concepts exist for this: either modify `parminator`, an EVLA parameter database tool, to work with the `vlbaparm` database, or to write a special-purpose tool for this (which would be similar to the `vlitedb` program affiliated with the VLITE project.

7.2 Real-time correlation

This item refers to operation of a DiFX correlator in Socorro (or other termination point for VLBA networks), possibly (and even likely) on the same cluster that runs the off-line correlation. An initial use of this would be execution of a reduced-bandwidth data quality test, to be run after maintenance each week. This would ensure proper timing, receiver polarization, and general usability of the array over a range of observing configurations.

Initial requirements:

1. Real-time correlation should be triggered by the VLBA operator, either using a GUI or documented command line procedure
2. Real-time correlation should support data rates up to 128 Mbps from each station
3. Default time and spectral resolution appropriate for continuum science
4. Only VDIF-based data (from DDC personality) needs to be supported
5. It should be possible to archive real-time correlated data; such data should be identifiable as such (e.g., through use of job names that include "rt" or similar)

Reference design components include:

- Software running on the Mark6 or X-cube at the site that transmits VDIF packets over UDP to a host machine at the correlator; a minimal program called `raw2udp` has been demonstrated as cable of doing this, but this functionality could be moved into `xcuberec` or `soft_switch`
- `vlitebuf` software running on a machine at the correlator, exposing the received data as files and allowing for significant buffering
- The DiFX correlator running in the same way it does for the VLITE system at the VLA: file-based datastreams are used, and the environment variable `DIFX_FILE_CHECK_LEVEL` should be set to `NONE` to prevent seeking within the `vlitebuf`-exposed "files".
- Management software that starts and stops the processes, creates DiFX file sets, and creates `.FITS` file(s) at the end

Additional potential features of lower priority:

- Overrides for default correlation parameters
- Derived parameters (e.g., delays determined via fringe fitting) can be put into the VLBA monitor database
- Automatic execution of real-time correlation based on a scan intent in the `.vex` file
- Channel selection: select a subset of recording channels for real-time correlation
- Automatic, near-real-time analysis of correlated data (e.g., pointing solutions, fringe detection, RFI plots, ...)
- Operation on Mark5B formatted data (from PFB personality)

7.3 Near-real-time data transfer and correlation

This item refers to fetching recording data from a Mark6 module that is currently mounted in a Mark6 unit at a VLBA site and storing it in a convenient location for subsequent processing (either inspection of data packets on a per-antenna basis, or for near-real-time correlation).

Initial requirements:

1. This capability should be available for the VLBA operators to execute

2. Data to download should be specified by the following user-supplied inputs:
   - Project code (including epoch), or a `.vex` file.
   - Antenna list (default: all that are in the project)
   - Scan name(s)
   - Maxumum length (in GB) to download per scan (default: unlimited)
3. The program should warn if existing downloaded data already exists

Reference design components include:

- `fringeCheck.py` – a reference implementation of a system that automates more or less the entire down-load and correlate process.
- `jive5ab` runs on both the station Mark6 unit and a corresponding machine at the correlator.
- After conclusion of correlation, `difxPlot.py` is run to plot some output diagnostics.

Additional potential features of lower priority:

- There should be an option to resume an interrupted download
- The program should provide an estimated download time, based on file size and `DOWNLOADRATE` parameter

7.4 Real-time interferometric pointing

Initially this would provide the capability to execute off-line pointing in a more accurate manner. Compared to the total-power based pointing that is now used exclusively, interferometric pointing offers greater immunity to marginal weather conditions and ability to use weaker sources. Later this may evolve to support referenced-pointing.

Initial requirements:

1. Perform correlation of pointing observations in real-time
2. Analysis of the pointing data would happen as a separate process
3. A list of suitable sources covering the sky $\delta > -40°$
4. New software to convert the correlator output to a form that can be processed by `fit`, or similar

This functionality will leverage the real-time correlation capability described above. Software will need to be developed to convert the correlator output into a format that can be used for pointing offset determination. Additional potential features of lower priority:

- Determine the pointing offsets immediately after correlation completes, or even as correlation proceeds
- Implement referenced pointing using interefrometric pointing

8 Acknowledgements

Thanks to Claire Chandler, Alan Erickson, Jeff Long, Matt Luce, Craig Walker, Vivek Dhawan, Ken Kellermann, and James Robnett for useful feedback and the gathering of folks at the VLBA Technical meeting on Jan 21 for their contribution to the discussion.

A CSA-F Statement of Work

The NSF proposal that resulted in the CSA-F funds contained within its Statement of Work a description of three project phases leading up to the goal of eventual combined VLA and VLBA operations. The CSA-F funds allow for the completion and initial operation of the first of these phases. The three phases were described in the proposal as follows:

A.1 Phase 1

In the first phase of deployment, all VLBA antennas will be fibered. The capacity of new fiber should be no less than 10 gigabits per second (Gbps). In this initial phase the network will be operated at approximately 200 megabits per second (Mbps), a capacity much greater than currently possible (typically 1.4 Mbps over copper cables to the sites), but much less than the fiber capacity. This limited use situation will still require hard drive recording for the majority of observations. However, the much improved connectivity will immediately allow three new capabilities:

1. Vastly improved diagnostics (through an existing nearly automatic "fringe check" process)
2. Support of some time-critical observations
3. Support for over-the-wire software and operating system upgrades.

All of these capabilities are important in support of both scientific and partner users of the VLBA.

A.2 Phase 2

In Phase 2, regular VLBA operations moves to network data transport rather than shipping hard drives. The data transmission rate on the fiber is increased to match or exceed the maximum operating bandwidth of the VLBA. Depending on the particular portion of network, this upgrade could be as simple as requesting higher access rates on the existing infrastructure, or it could require deployment of new VLBA networking equipment on the data path. Additionally, the inbound bandwidth of the processing center must be increased to the aggregate data rate of all participating antennas. Currently this is 20 Gbps with an increase to 40 Gbps anticipated in the near future. In addition to completing the network infrastructure upgrades, a moderate amount of software will need to be developed which will allow for buffering data transport and storing raw data at the processing center in cases where real-time processing is not possible. Phase 2 will enable the following capabilities:

1. Real-time or near-real-time data processing; scientific results in hours, not weeks
2. Elimination of hard drive shipping from routine VLBA operations
3. Improved integration of the VLBA into international VLBI networks

These capabilities are largely improvements to the VLBA operational model which will lead to improved flexibility and sustainability of operation through reduced staffing requirement (by approximately 1 FTE).

A.3 Phase 3

Phase 3 is the full integration of the VLBA with NRAO's JVLA and/or anticipated Next Generation Very Large Array (ngVLA). Full integration would result in the following:

1. Single telescope array, operated with flexibly configured sub-arrays
2. Increased commonality of NRAO/LBO radio astronomy instrumentation
3. Reduction of operations footprint

4. Integration of JVLA or ngVLA into international VLBI networks

These benefits will further improve sustainability of VLBA operations through reduced staffing (by approximately 4 FTEs) and offer additional capabilities to scientific users. Full integration with the ngVLA is seen as critical for maximal exploitation of the range of interferometer baselines that would be offered by the combined instrument.