

 Open access • Journal Article • DOI:10.1086/658907

DiFX2: A more flexible, efficient, robust and powerful software correlator

— [Source link](#) 

Adam T. Deller, Walter Briskin, Chris Phillips, John Morgan ...+7 more authors

Published on: 05 Jan 2011 - [arXiv: Instrumentation and Methods for Astrophysics](#)

Topics: [Software](#)

Related papers:

- [DiFX-2: A More Flexible, Efficient, Robust, and Powerful Software Correlator](#)
- [AIPS, the VLA, and the VLBA](#)
- [DiFX: A Software Correlator for Very Long Baseline Interferometry Using Multiprocessor Computing Environments](#)
- [DiFX: A software correlator for very long baseline interferometry using multi-processor computing environments](#)
- [The NRAO VLA Sky Survey](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/difx2-a-more-flexible-efficient-robust-and-powerful-software-2ejtpxcoh9>



CHICAGO JOURNALS



DiFX-2: A More Flexible, Efficient, Robust, and Powerful Software Correlator

Author(s): A. T. Deller, W. F. Briske, C. J. Phillips, J. Morgan, W. Alef, R. Cappallo, E. Middelberg, J. Romney, H. Rottmann, S. J. Tingay, R. Wayth

Reviewed work(s):

Source: *Publications of the Astronomical Society of the Pacific*, Vol. 123, No. 901 (March 2011), pp. 275-287

Published by: [The University of Chicago Press](#) on behalf of the [Astronomical Society of the Pacific](#)

Stable URL: <http://www.jstor.org/stable/10.1086/658907>

Accessed: 12/03/2012 04:33

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at

<http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The University of Chicago Press and Astronomical Society of the Pacific are collaborating with JSTOR to digitize, preserve and extend access to Publications of the Astronomical Society of the Pacific.

<http://www.jstor.org>

DiFX-2: A More Flexible, Efficient, Robust, and Powerful Software Correlator

A. T. DELLER,^{1,2} W. F. BRISKEN,¹ C. J. PHILLIPS,³ J. MORGAN,⁴ W. ALEF,⁵ R. CAPPALLO,⁶ E. MIDDELBERG,⁷
J. ROMNEY,¹ H. ROTTMANN,⁵ S. J. TINGAY,⁴ AND R. WAYTH⁴

Received 2010 September 20; accepted 2010 December 27; published 2011 February 23

ABSTRACT. Software correlation, where a correlation algorithm written in a high-level language such as C++ is run on commodity computer hardware, has become increasingly attractive for small- to medium-sized and/or bandwidth-constrained radio interferometers. In particular, many long-baseline arrays (which typically have fewer than 20 elements and are restricted in observing bandwidth by costly recording hardware and media) have utilized software correlators for rapid, cost-effective, correlator upgrades to allow compatibility with new, wider-bandwidth, recording systems and to improve correlator flexibility. The DiFX correlator, made publicly available in 2007, has been a popular choice in such upgrades and is now used for production correlation by a number of observatories and research groups worldwide. Here, we describe the evolution in the capabilities of the DiFX correlator over the past three years, including a number of new capabilities, substantial performance improvements, and a large amount of supporting infrastructure to ease use of the code. New capabilities include the ability to correlate a large number of phase centers in a single correlation pass, the extraction of phase-calibration tones, correlation of disparate but overlapping sub-bands, the production of rapidly sampled filter-bank and kurtosis data at minimal cost, and many more. The latest version of the code is at least 15% faster than the original (and in certain situations, many times this value). Finally, we also present detailed test results validating the correctness of the new code.

Online material: color figures

1. INTRODUCTION

Development of the Distributed FX (DiFX) software correlator began in 2005, primarily for usage with the Australian Long Baseline Array (LBA) as part of a sensitivity upgrade program (Deller et al. 2007), where it entered production usage in 2006. It is an FX-style correlator (see, e.g., Chikada et al. 1987; Thompson et al. 1994; Romney 1999) designed to run on modern CPUs under Linux or Mac OS X. The basic principles of radio astronomy cross-correlator fundamentals will not be rederived in this article, which focuses on the particular implementation of the DiFX software correlator. We direct the reader to the preceding references for a thorough explanation of FX-style correlator functionality and to Deller et al. (2007) for a

comprehensive description of the specific implementation of this functionality in the DiFX code.

The DiFX code is accelerated using vector arithmetic libraries: specifically, the Intel Integrated Performance Primitives (IPP) library,⁸ and the distribution across multiple nodes is enabled using implementations of the Message Passing Interface.⁹ The advantages brought by the adoption of a newer, more flexible, correlator architecture were enumerated by Deller et al. (2007) and included greater flexibility in the setting of correlation parameters, lower cost, rapid development, ease of maintenance, and upgradability (both in hardware and software). In keeping with this final point, development of the DiFX software correlator has continued rapidly since its first public release in 2007. Since that time, many new features and performance improvements have been merged into the DiFX codebase, which cumulatively are sufficient to merit a major version increment for the DiFX package, which we designate DiFX-2. Where necessary, any release of DiFX prior to DiFX-2 will be referred to generically as DiFX-1.x, acknowledging that this time period spanned a series of release numbers. Some of the features presented here were made available early in the DiFX-1.x series and have been available for several years, while others were added recently and are only available in DiFX-2.

¹ National Radio Astronomy Observatory, PO Box O, Socorro, NM 87801.

² UC Berkeley, 601 Campbell Hall, University of California at Berkeley, Berkeley, CA 94720.

³ Australia Telescope National Facility, CSIRO Astronomy & Space Sciences, PO Box 76, Epping, NSW, Australia.

⁴ International Centre for Radio Astronomy Research, Curtin University, GPO Box U1987, Perth, WA, Australia.

⁵ Max-Planck-Institut für Radioastronomie, Postfach 20 24, 53010 Bonn, Germany.

⁶ MIT Haystack Observatory, Westford, MA 01886.

⁷ Ruhr-Universität Bochum, Astronomisches Institut, NA 7/73, Universitätsstra. 150, D-44801 Bochum, Germany.

⁸ See <http://software.intel.com/en-us/intel-ipp/>.

⁹ See <http://www.mcs.anl.gov/research/projects/mpi/>.

DiFX has been adopted by a number of leading VLBI facilities in addition to the LBA. Specifically, the Very Long Baseline Array (VLBA) operated by the National Radio Astronomy Observatory (NRAO¹⁰) in the US has retired its hardware correlator, which was designed in the 1980s, and migrated completely to DiFX. In addition, the Max Planck Institute for Radioastronomy (MPIfR) has begun routine operation of DiFX in parallel with the existing Mark4 hardware correlator operations and plans to phase out its Mark4 hardware correlator by the end of 2010. The specific needs of the LBA, VLBA, and MPIfR have driven the development of many of the new capabilities of DiFX, including the features discussed in this article, which allow entirely new areas of long-baseline science to be undertaken. All of these developments have been made available to all current and potential users.

Almost as important for new users of DiFX, considerable effort has been made to improve the documentation and online resources available for installing and testing DiFX.¹¹ Two mailing lists are available to seek or disseminate information regarding DiFX: one list reaches the entire DiFX community, while the second is focused specifically on code developers. Prospective users of DiFX are directed to the online resources, where further information is available on how to obtain the code from the central repository. Finally, improved version control has been implemented since the early days of DiFX-1.x, and tagged releases of frozen code are made available on a regular basis. The current stable version of DiFX-2 is DiFX-2.0.0, and the current version of the previous series is DiFX-1.5.4.

In this article, we describe the new DiFX-2 functionality in § 2. The performance improvements which have been provided in the accompanying code changes are listed and quantified in § 3. Additional code that provides functionality for DiFX-2 that is not directly part of the correlator itself is described in § 4, and the validation testing undertaken for DiFX-2 is described in § 5. Future work is discussed in § 6, and our conclusions are presented in § 7.

2. NEW FEATURES

2.1. FITS-IDI and Mark4 Format Output

Initially, DiFX-1.x only supported the RPFITS¹² file format, which was the format historically used by the LBA. However, RPFITS is not a standard FITS format and has limited support in most radio astronomy postprocessing packages. Accordingly, in version 1.5.0 the ability to produce FITS-IDI format correlator files (as produced by the VLBA and the Joint Institute for VLBI in Europe [JIVE] hardware correlators) was added to DiFX.

¹⁰ The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc.

¹¹ See <http://cira.ivec.org/dokuwiki/doku.php/difx/start>.

¹² See <http://www.atnf.csiro.au/computing/software/rpfits.html>.

Unlike the RPFITS files written by early versions of DiFX-1.x, the FITS file is not directly written by the correlator, but is translated from a DiFX binary output written by DiFX-2 after correlation has completed. This eliminates the need to link large FITS libraries into DiFX and simplifies and speeds the output writing process. Support for the RPFITS format has been withdrawn in DiFX-2.

Geodetic observers typically use specialized postprocessing software such as HOPS,¹³ which is closely tied to the visibility data format produced by the Mark4 hardware correlator (Whitney et al. 2004). In order to facilitate the use of DiFX-2 for geodetic observations, an additional translation program has been written to produce these Mark4 format visibility data sets from the binary DiFX output. The ability to import DiFX format output data directly into the HOPS geodetic postprocessing package is new in DiFX-2.

2.2. Native Mark5 Interface

The Mark5 recording media series (Whitney 2003) is widely used among VLBI networks, with the VLBA, the European VLBI Network (EVN), the Korean VLBI Network (KVN), and global geodetic arrays all utilizing the system. Initially, however, DiFX-1.x could not read data directly from Mark5 disk modules, as it required the data to be accessible from a standard Linux file. Correlating Mark5 data therefore required a tedious intermediate step of exporting the files to a Linux file system, which imposed additional overhead in time and storage space. Since version 1.5.0, DiFX has had the ability to read Mark5 disk modules natively, using the application interface available from the Mark5 vendor. This eliminates the need to export data from modules to standard files, streamlining the correlation process and reducing the need for large amounts of standard disk storage. The support of VLBA and Mark4 data formats on both standard Linux files and Mark5 modules has been extended to include Mark5B, and limited support is already in place for the next-generation VDIF format described by Whitney et al. (2009).

2.3. Phase-Calibration-Tone Extraction

Phase-calibration tones can be injected at the front end of a radio astronomy antenna in order to provide a convenient means to estimate instrumental delays. For applications such as geodesy, phase-calibration tones are heavily relied upon. While some radio telescope arrays extract, average, and store phase-calibration-tone information at the antenna, others rely on the correlator to perform this important function. Accordingly, a flexible phase-calibration-tone extraction system has been added to DiFX-2. The phase-calibration extraction in DiFX-2 can be configured to extract any number of tones, unlike many existing

¹³ See <http://www.haystack.mit.edu/tech/vlbi/hops.html>.

hardware implementations such as those at the VLBA stations (which provide two tones per sub-band). Preliminary results comparing DiFX corrections with those extracted at the VLBA stations show agreement to $\ll 1^\circ$ (corresponding to \sim femtoseconds at 43 GHz) and also show that the computational overhead of extracting all the phase-calibration tones present¹⁴ is $\sim 5\%$. A detailed analysis of the performance of DiFX on geodetic observations, including the verification of DiFX-2 phase-calibration extraction and the production of Mark4 format visibility data, is deferred to a future publication (Morgan et al. 2011, in preparation).

2.4. Spectral Selection and Averaging

Once the data have been channelized (the “F” portion of the FX algorithm), it is possible to discard segments of the spectrum that hold no interest for the current observation. The main application of such spectral selection is to zoom in on widely separated spectral features such as masers that are contained within a wide bandwidth. Use of this new feature in DiFX-2, which we generically term “zoom mode,” reduces the load on the cross-multiply/accumulate (“X”) portion of the FX algorithm and, more importantly, reduces the amount of data that must be returned to the manager node for long-term accumulation (see Deller et al. 2007). This allows very high spectral resolution to be obtained without overloading the correlator interconnect, without generating unduly large amounts of data to be written to intermediate disk results and later discarded (as was the case with all versions of DiFX-1.x).

An alternate application of spectral selection, which we generically term “band-matching,” facilitates the correlation of heterogeneous recorded bands by subdividing wider bands recorded at some antennas to match narrower bands recorded at other antennas. This is a particularly useful feature for correlating infrequently used antennas with nonstandard VLBI back-end systems, which are unable to produce bands compatible with other VLBI systems. A current example involves geodetic correlations where 32 MHz bands recorded at most stations are correlated against 16 MHz bands recorded at the Plateau de Bure interferometer. The correlation of upper sideband data with lower sideband data (covering the same spectral range) is also supported.

Spectral averaging allows multiple spectral points to be averaged after correlation, but before the visibilities are returned to the manager node. This is useful in two cases. The first is when the desired final spectral resolution is low. In this instance, generating such coarse spectral resolution directly by means of a very short Fourier transform is not efficient, so use of an optimally sized transform (typically with a length of ~ 256 points) is preferred, with spectral averaging before the visibilities are

returned to the manager. As with spectral selection, this saves correlator interconnect and disk-space resources.

The second, and more important, usage of spectral averaging is with the multiple-phase-center correlations described in the following section. In order to avoid bandwidth decorrelation effects as described subsequently, very high spectral resolution is required initially. After the visibilities have been shifted to the desired phase center, however, they can be averaged to a standard VLBI resolution. Spectral averaging is critical for this application, as the return of multiple copies of very high spectral resolution visibilities would completely overwhelm the correlator interconnect.

2.5. Multiple Simultaneous Phase Centers

Because of the very high fringe rates inherent in VLBI observations, the use of standard frequency (hundreds of kilohertz) and time resolution (seconds) leads to an extremely small (several-arcsecond) field of view (see, e.g., Middelberg et al. 2010). We hereafter refer to a narrow field of view resulting from standard VLBI correlation parameters as a “pencil beam.” Thompson et al. (1994) provide a detailed explanation of the challenges inherent in wide-field imaging. While improving the temporal and spectral resolution allows somewhat wider fields of view (and was indeed one of the main drivers of DiFX development), this carries an increasingly problematic cost in expanded data volume. Mapping even one-tenth of the primary beam of the VLBA at 1.4 GHz with no more than 10% decorrelation due to time and bandwidth decorrelation requires 8 kHz frequency channels and an integration time of 0.1 s—which yields visibility data sets > 2 Tbyte for a typical 12 hr VLBA observation at current bandwidths.

An alternative to mapping large swaths of sky (which, in any case, are almost entirely empty at VLBI resolution at centimeter frequencies) is to image small areas around known sources. This can be accomplished by shifting the phase center of the correlation (a uv shift) to the location of known sources and averaging visibilities to obtain manageable-sized data sets, which can be used to produce pencil beams at new locations (see, e.g., Lenc et al. 2008). A uv shift is implemented by calculating the baseline-based differential geometric delay between the desired and applied phase center, converting to a phase rotation for each visibility by multiplying this delay by the associated sky frequency, and rotating the visibility phases by this value. Morgan et al. (2010) examine the problem of uv shifting in more detail, including the detailed calculation of the necessary delay shifts. The drawback of using this approach after correlation is the necessity of generating an initial visibility data set that is as large as that required for a single large image. The intermediate data volume problem is therefore comparable with that experienced with the single-large-image approach, and the I/O cost of writing visibilities to, and reading from, disk is substantial.

If implemented within the correlator, however, the twin problems of I/O and storage volume are solved, because the

¹⁴ Spaced at 1 MHz intervals, so typically 8 or 16 tones per sub-band.

intermediate data products (the high spectral resolution visibilities held at the processing nodes) do not need to be transmitted from where they are calculated and are never written to disk. Obtaining sufficiently high time resolution is trivially implemented—the time division multiplexing within DiFX (see Deller et al. 2007) already provides time resolution better than that required in most cases, but the ability to uv shift and average after a shorter, user-specified, time has also been implemented. For P phase centers, the processing nodes transmit P normal-sized (postaverage) collections of visibility results back to the manager node, and P normal-sized data sets are ultimately written to disk. The impact on performance of this feature is relatively small, due to the fact that the uv shift/average operations need only be carried out relatively infrequently (compared with the multiplications, additions, and Fourier transforms required for the regular correlation process).

The visibility amplitudes and weights are corrected for time and bandwidth decorrelation online, before the visibilities are written to disk. The resultant P pencil beams can be reduced and imaged using standard tools. Presently, these corrections are not tabulated and saved, since no postprocessing software exists that could parse and use this information. The information is readily available, however, and could be formatted and written out when suitable postprocessing becomes available.

Thus, as long as low-resolution finder catalogs are available, VLBI-resolution surveying is possible with DiFX-2 with minimal overhead. Figure 1 illustrates the use of a low-resolution image and the VLBI data sets that would result from a multiple-field correlator pass. Middelberg et al. (2010) have already used this new capability to carry out pilot VLBI survey observations in the Chandra Deep Field South, and these observations were instrumental in the development and refinement of the new correlation mode. Section 3 describes in detail the performance impact of adding multiple phase centers to a correlation. Section 5 shows the verification that uv shifted visibility data sets have no residual phase or amplitude errors. This feature is new in DiFX-2.

2.6. Correct Model Accountability

The RPFITS output format used initially by DiFX-1.x had no means to store an accurate representation of the delay model applied at the correlator. The transition to FITS-IDI in version 1.5.0 has made correct model accountability possible, and an accurate representation of the applied delay model is now stored in two binary tables—the IM table and the MC table. These tables store the same sampled model polynomials used by DiFX-2 and the applied clock model and can be used by postprocessing software such as AIPS¹⁵ to accurately make changes to the phase center of the correlated data set or to correct for antenna position errors, Earth orientation parameter errors, and the like.

¹⁵ See <http://www.aips.nrao.edu/>.

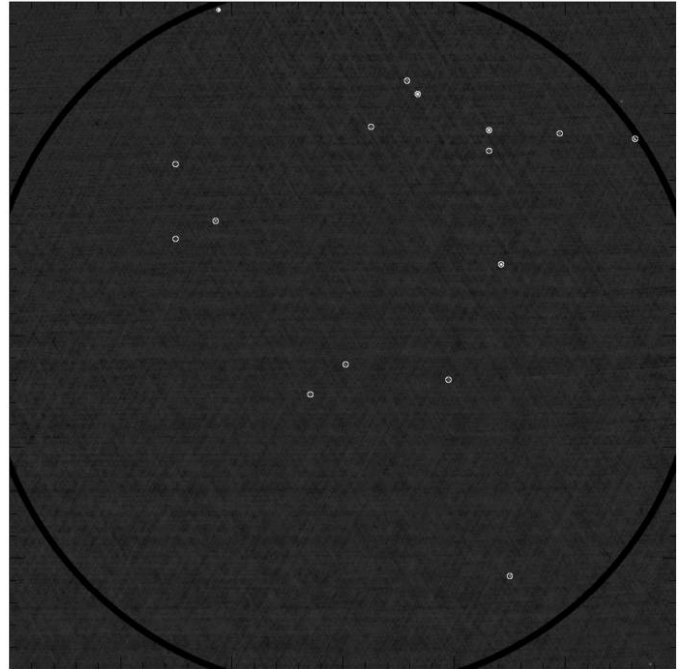


FIG. 1.—Example of a representative finding field centered on $07^{\text{h}}45^{\text{m}}07.270^{\text{s}} + 33^{\circ}40'37.52''$ (from the FIRST [Faint Images of the Radio Sky at Twenty cm] survey—<http://sundog.stsci.edu/>). The bold black ring shows the $31'$ primary beam of a 25 m dish at 1600 MHz, and the small white rings show the individual pencil beams that would be placed on known sources. The pencil-beam diameter is displayed as $12''$, at which point the cumulative time and bandwidth decorrelation from 0.5 MHz channelization and 4 s averaging reaches 10%.

2.7. New Data Monitoring Tools

2.7.1. Autocorrelation Filterbank Spigot

In order to facilitate searches for transient signals, an autocorrelation “spigot” has been added to DiFX-2. This spigot supplies the autocorrelations from each antenna at a user-specified time and frequency resolution, by means of a UDP (User Datagram Protocol) multicast message. The additional computational load is negligible, since the antenna autocorrelations are already calculated as a matter of course, and the additional load of sending the multicast messages is negligible for all but very short integrations. The messages are sent with a simple plain-text header, allowing (one or more) analysis programs to capture, time-order, and inspect what are essentially N independent but time-aligned filter-bank data streams, where N is the number of antennas. This feature has been available since version 1.5.1, but is considerably improved in DiFX-2. An example of the two-dimensional dynamic spectrum that is obtained (for each antenna) is shown in Figure 2.

At the VLBA, additional functionality has been added to allow feedback from the analysis programs, allowing them to request that small time ranges of baseband data be extracted after the correlation has finished and written to disk elsewhere.

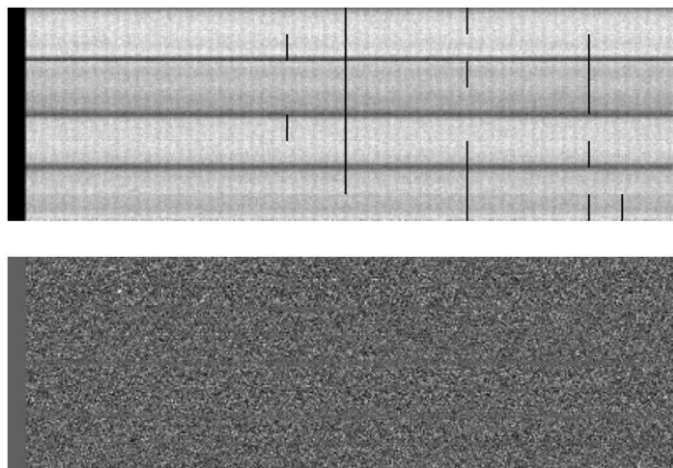


FIG. 2.—Example of the autocorrelation dynamic spectrum produced from the Brewster VLBA antenna, with grayscale intensity representing autocorrelation signal strength. These data were captured commensally during an observation in 2010 June. Time runs horizontally covering a period of 1 s, and the 64 MHz of bandwidth (consisting of eight concatenated 8 MHz sub-bands spanning 1350.49–1414.49 MHz in right circular polarization) runs vertically. A single pixel is 2 ms and 500 kHz. *Top*: The raw filter-bank output. The imprint of the 80 Hz noise calibration signal present in VLBA data is clearly visible, as is that of the sub-band bandpasses. Data lost during times of network congestion appear as zero amplitude (*vertical black lines*). *Bottom*: The processed filter-bank data presented to transient detection code, which has been filtered to interpolate missing data, remove bandpass shapes, and remove the noise calibration signal.

Thus, the detection pipeline can trigger baseband data “grabs” based on the autocorrelation filter-bank data, permitting a detailed analysis at full time resolution after the correlation has completed. For the first time, this offers the possibility of full-time commensal observations on VLBI arrays. This facility is currently being used for a commensal transient search of VLBA data in support of an upcoming large transient survey on the Australian Square Kilometre Array Pathfinder (ASKAP). The Commensal Real-time ASKAP Fast Transient (CRAFT) survey is described by Macquart et al. (2010), while the VLBA fast transients pipeline is described by Wayth et al. (2011).

2.7.2. Station-based Kurtosis Estimation

Spectral kurtosis (Nita et al. 2007) can be calculated for radio filter-bank data to estimate the form of the probability density distribution function for each filter-bank channel. Since radio-frequency interference (RFI) generally corrupts the form of the probability density distribution of the filter-bank data, spectral kurtosis can be a powerful and inexpensive method of identifying spectrally confined RFI. Recently, Deller (2010) tested a simplistic implementation of spectral kurtosis calculation within DiFX and used it to identify previously unnoticed rapidly-time-varying RFI at one VLBA station. The test implementation described in Deller (2010) has been updated to the fully correct spectral kurtosis calculation in DiFX-2, and the results are

now made available using the same “spigot” architecture used for the autocorrelation filter bank described in § 2.7.1. The spectral kurtosis values are used to identify RFI before the filter-bank data are passed through the transient detection pipeline.

2.7.3. Real-Time Visibility Monitoring

The final new data monitoring tool included in DiFX-2 is a TCP-based visibility monitor server, which was first deployed in version 1.5.2. Enabling this feature causes DiFX to send copies of the visibility data through a TCP network connection to a monitor server, which sorts and sends selected visibilities to connected clients for real-time processing and/or display. This feature enables data quality assessment during correlation, which is particularly useful for the verification of correct array and correlator setup during non-disk-based observations (eVLBI).

2.8. Other New Functionality

A number of other minor yet useful features have been added to DiFX-2. The first is the ability to compensate for local oscillator (LO) offsets on a frequency-by-frequency basis. An LO offset at a station leads to continuously wrapping phase with time (constant across all spectral channels) on all baselines to the affected station. In DiFX-2 an appropriate phase compensation is made after channelization in the same operation as fractional-sample correction (see Deller et al. 2007). This allows correction of LO offsets up to a small fraction of the spectral channel bandwidth (so typically up to tens of kilohertz) without significant decorrelation occurring. LO offset correction can be used in tandem with spectral selection to aid in correlating mismatched sub-bands, but at this time phase-calibration tones cannot be extracted from sub-bands that have had an LO correction applied. The lifting of this restriction will be the subject of future development.

In addition to reading from files and Mark5 modules, DiFX has the ability to accept baseband data from a network socket. In DiFX-1.x, this data transfer over a network was restricted to the use of a TCP (guaranteed transfer) transport protocol. While it was possible to obtain sufficiently good network performance over short and/or dedicated network links to perform high-speed eVLBI with DiFX-1.x (see, e.g., Phillips et al. 2007), the congestion control inherent in TCP makes it unsuitable for long, potentially lossy, network transfers. DiFX-2 retains the ability to perform correlations reading from a TCP socket, but adds the ability to connect data sources via UDP—a transfer protocol without congestion control that is better suited to maintaining high transfer speeds on a long, lossy, or shared transmission network.

DiFX-2 now supports the reading of complex sampled baseband data (in the VDIF format only). The use of complex sampled data (where the digital representation of the antenna voltage is stored in a complex representation at half the sampling rate of a real sample stream) offers potential advantages

over real sampling, including less processing required before storage and slightly reduced quantization losses at a given bit precision. Modern digital data-acquisition systems typically use complex sampled data internally, but in all current VLBI systems (and many non-VLBI systems) the complex sample stream is converted to a real sample stream before recording/correlation. Supporting complex sampled data will reduce the amount of work required for an ad hoc experiment utilizing a non-VLBI antenna with a system that does not convert to real sampling and will prepare for potential future VLBI systems utilizing complex sampling. As with real data streams, complex sampled data streams can be obtained from network connections, disks, or Mark5 modules and can be correlated against other complex data streams or real data streams.

Finally, DiFX-2 has extended the flexibility of the clock-model specification. Like most correlators, DiFX-1.x possessed the ability to compensate for a clock offset and linear rate of change at each station. DiFX-2 allows the specification of an arbitrary-order clock polynomial at each station, allowing more accurate correction for known clock variations. This is likely to be most applicable at very high frequencies.

3. PERFORMANCE IMPROVEMENTS

Testing of the performance of DiFX has been carried out on a variety of Intel-based clusters. Most have been comprised of recent Intel Xeon multicore CPUs—the results presented here were derived from the 10 node cluster installed at the NRAO Domenici Science Operation Center in Socorro, New Mexico, to replace the VLBA hardware correlator. Each node in this system is dual-CPU, where each CPU is an Intel Xeon quad core with 6 Mbytes of shared L2 cache, running at 2.5 GHz. Each node has 4 Gbytes of RAM. Including the 1 Gbit Ethernet switching infrastructure, the total cost of the cluster (in 2008) was approximately \$30,000. As shown subsequently, this system is capable of sustaining throughput of 512 Mbit s^{-1} for 10 stations, twice that of VLBA hardware correlator for typical experiments. As part of the ongoing VLBA sensitivity upgrade program¹⁶ (which aims to demonstrate a station data rate of 4 Gbit s^{-1} by 2011, along with routine operations at 2 Gbit s^{-1}) this cluster is being substantially enlarged, at a cost that is a small fraction of the cost of recording media required to operate at the higher data rates. DiFX is well poised to take advantage of ongoing improvements in CPU technology, such as the trend to many-core architectures and extended instruction sets, through its use of IPP for vector operations. IPP is regularly updated to make best usage of the latest Intel (and Intel-compatible) CPU architectures.

Two code changes in DiFX-2 dominate the performance improvements over earlier incarnations of DiFX. The first is a more efficient implementation of vector phase rotations, where

the phase change is constant from one element in the vector to the next. In DiFX, such operations include:

1. The (time domain) fringe rotation, which multiplies a constant oscillator frequency by a changing model delay (see Deller et al. 2007). Over short time periods, the delay change from one sample to the next is linearly approximated, which yields a constant phase increment.
2. The (frequency domain) fractional-sample correction, which multiplies a constant (for a given time window) delay by the frequency of the spectral point (see Deller et al. 2007). The linearly increasing frequency of the spectral points across the sub-band again yields a constant phase increment.
3. The repositioning of phase centers in a multicenter correlation (DiFX-2 only), which is conceptually identical to a fractional-sample delay, although the magnitude of the delay applied can be greater than one sample.

Phase rotations are applied to complex data using complex multiplications by a rotation vector of unit amplitude, whose real and imaginary components are computed by taking a sine and cosine (\sin / \cos) of a vector containing the desired phase changes. In general-purpose CPUs, the evaluation of trigonometric operations such as \sin / \cos is considerably more computationally expensive than complex arithmetic such as multiplication or addition. When the phase interval change between array elements is constant, it is no longer necessary to compute \sin / \cos values for every array element, as was implemented in DiFX-1.x. Instead, as shown in Figure 3, it is possible to compute the exact \sin / \cos values for only the first few elements of a vector (a subvector) and a series of offset elements. The subvector is then multiplied by each offset in turn, and the rotated subvector is placed in an appropriate position in the final vector. The number of trigonometric operations required could be further reduced by performing a similar

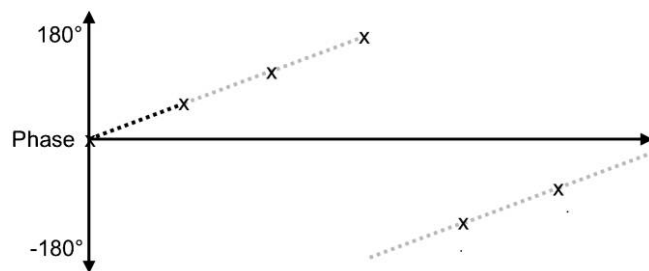


FIG. 3.—Illustration of the complex multiplication “filler” approach to minimize the trigonometric operations needed to produce a vector with linear phase change from one element to the next. The figure plots the phase in degrees of a complex vector of unit amplitude—the desired rotation vector. The complex values shown by the black points and crosses are calculated trigonometrically, and the values shown by gray points are then filled in with a series of complex multiplications. Each value shown by a cross is used in turn to rotate the vector of black points, and the resultant vector of complex numbers has an appropriate sequence of phases to be stored in the final rotation array between that cross and the next.

¹⁶ See <http://www.vlba.nrao.edu/memos/sensi/>.

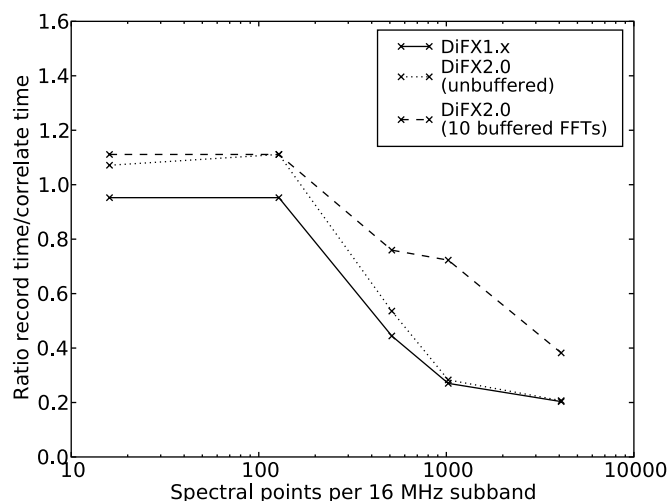


FIG. 4.—Throughput of DiFX-2 (with and without FFT buffering) compared with DiFX-1.x for a varying number of spectral points (quoted per 16 MHz subband). The test was run on the VLBA DiFX cluster (10 dual-quad-core nodes) on nine stations of data totaling $512 \text{ Mbit s}^{-1} \text{ station}^{-1}$ (eight sub-bands consisting of four frequencies recorded in dual polarization). Four polarization products were computed for each of the frequency bands. The vertical axis shows the ratio of record time to correlate time (the speedup factor of the correlation). For values above 1.0, the correlation is proceeding faster than the data were originally recorded. The DiFX-1.x performance (*solid line*) drops sharply when the visibility results exceed the available node cache, as does DiFX-2 with no FFT buffering (*dotted line*), but DiFX-2 with FFT buffering (10 FFTs buffered) (*dashed line*) is much less adversely affected.

decomposition of the subvector, but a cascaded approach such as this is not implemented, to avoid numerical precision issues and because the performance improvement from further decomposition is marginal for typical array sizes in DiFX.

With a suitable choice of the subvector length for this basic decomposition (typically the nearest factor of 2 to the square root of the vector length V) it is possible to reduce the required number of trigonometric computations by a factor of up to $\sqrt{V}/2$. In the DiFX-1.x implementation, trigonometric operations comprised almost one-third of the station-based computational load. In DiFX-2, where fringe rotation and fractional-sample vectors are typically of length 100–1000 elements, the resultant order-of-magnitude computational saving means that the cost of the trigonometric operations is reduced to a near-negligible value, at a small cost of extra complex arithmetic. For correlations with a small number of stations (≤ 10), the net increase in correlator throughput is 15–20%. (Alternatively, 15–20% less computational resources are required to obtain a prescribed throughput.)

The second efficiency improvement comes from an increase in flexibility in the order of traversing of the baseline-based cross-multiplications, allowing DiFX-2 to mitigate the effect of cache misses when the size of the visibility result vector grows large (many stations and/or many spectral channels). A cache miss occurs when data are no longer available for immediate use in the CPU cache (after being overwritten by some

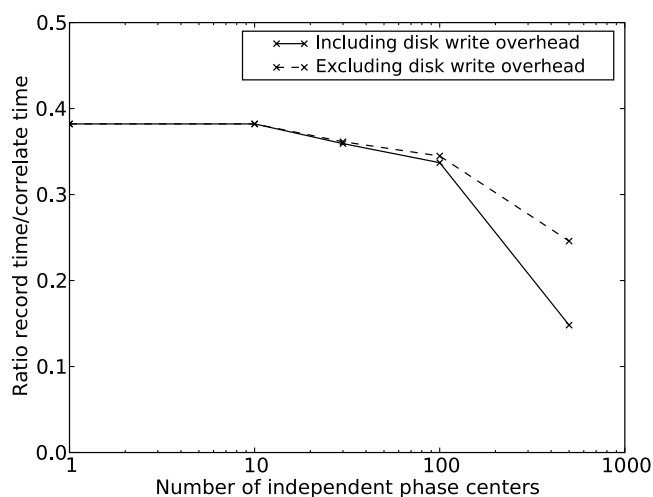


FIG. 5.—Throughput of DiFX-2 (10 buffered FFTs) for an increasingly large number of phase centers. The data set and correlator resources were identical to the previous benchmark, and as before, the vertical axis shows the ratio of record time to correlate time. Sub-bands were channelized with 4096 spectral channels, and the uv shifts were performed every 26 ms. All polarization products were computed. The solid line shows the observed throughput, while the dashed line removes the slowdown caused by writing the visibilities to a slow disk. Up to hundreds of phase centers can be correlated while imposing a near-negligible impact on correlator throughput, although disk write speed becomes a limitation with hundreds of fields. This could be mitigated with a faster RAID disk for storing output visibilities. Even combined with the slowdown due to the larger FFT size (as seen in Fig. 4), the correlation cost of producing 500 phase centers is only 4.5 times that of a normal single phase-center continuum correlation—yielding a speedup factor greater than 100.

more recently used data) and must instead be retrieved from main memory, which is a relatively slow operation. In DiFX-1.x, every polarization pair of every frequency sub-band for every baseline at a given time was cross-multiplied sequentially. When the visibility result array grows too large to remain in cache, this traverse leads to a cache miss for every accumulation and a dramatic slowdown.

In DiFX-2, the results of N channelizations (hereafter referred to as FFTs [fast Fourier transforms] in accordance with the implementation) are buffered for each data stream, and a single selection of spectral points for a given baseline are then cross-multiplied N times into the same area of the visibility result buffer. This allows the intermediate vectors at the data stream to remain in cache during the station-based processing for $N - 1$ passes and allows the visibility result vector to remain in cache for $N - 1$ passes. For moderate values of $N \sim 10$, a minimal amount of extra memory is required and the number of cache misses is greatly reduced.

Figure 4 shows the net improvement in throughput for DiFX-2 compared with DiFX-1.x, for varying numbers of spectral points (and hence varying visibility result lengths). The improvement is less marked for small numbers of spectral points, reflecting the 15–20% improvement solely from the

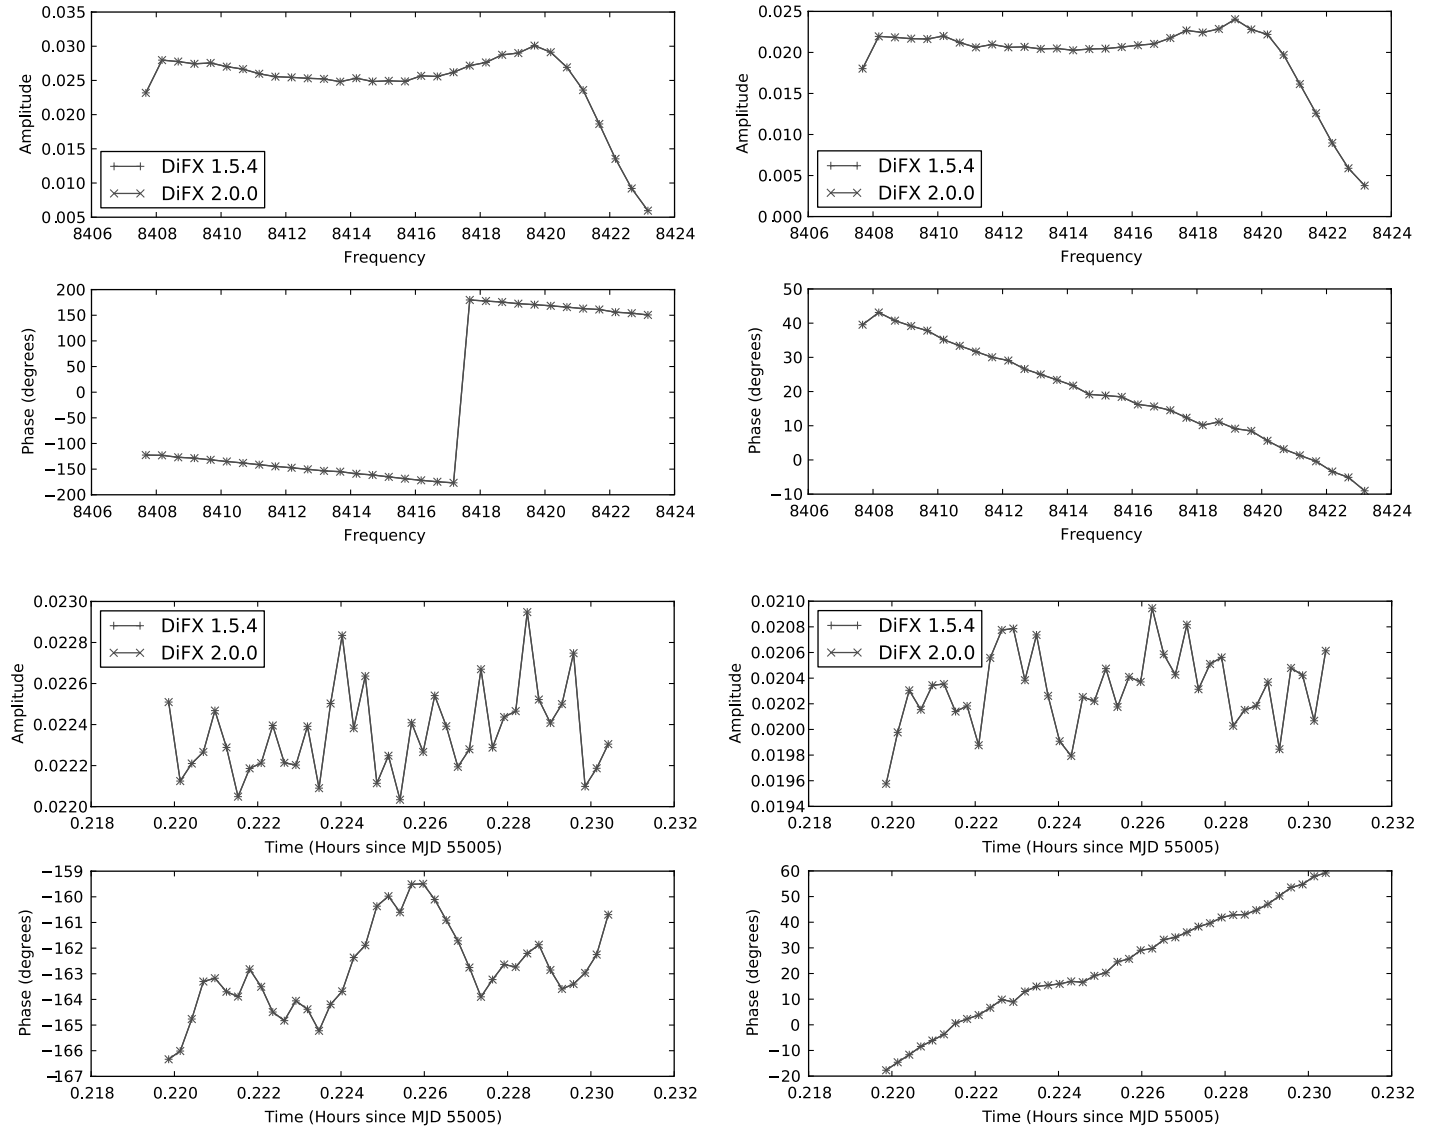


FIG. 6.—Comparison of the correlated output of DiFX-2.0.0 (*solid line*) and DiFX-1.5.4 (*dashed line*) for a single observing band for a pair of VLBA baselines over a 40 s scan. The target is the strong calibrator 4C39.25, and the observing band plotted spans 16 MHz from 8407.49–8423.49 MHz, in right circular polarization. *Top*: Baseline from Brewster to Fort Davis (2350 km). *Bottom*: Brewster to Saint Croix (5770 km). *Left*: Visibility amplitude and phase (averaged across the band) vs. time. *Right*: Visibility amplitude and phase (averaged for the scan duration) vs. frequency. No difference is apparent on this scale—formal analysis shows that the rms amplitude difference is 0.0007% and the rms phase difference is 0.0007°. See the electronic edition of the *PASP* for a color version of this figure.

more efficient trigonometric processing. As the visibility result buffer exceeds the cache size (each processing thread had access to approximately 1.7 Mbytes of L2 cache, which the visibility buffer exceeds in size for 512 spectral points) the older DiFX-1.x code suffers a marked drop in performance. DiFX-2 also experiences reduced throughput (primarily due to the increased FFT cost with a larger number of spectral points), but the reduction is much smaller.

Finally, the performance of DiFX-2 with multiple phase centers should be noted. Figure 4 shows that the high spectral resolution required to minimize the decorrelation suffered during

uv shifts carries its own penalty (dependent on the sub-band bandwidth, but typically 2048 or 4096 spectral points, yielding a computational load increase of 2–3 times over a standard 16 spectral point continuum observation). However, beyond this initial penalty, the cost of adding additional phase centers is very small. Figure 5 shows the variation of correlator throughput with number of phase centers for a fixed spectral resolution. For this test, spectral resolution of 2 kHz and temporal resolution of 26 ms were used—sufficient to shift to the edge of the VLBA primary beam at 1.4 GHz (15′) with <5% decorrelation due to time and bandwidth smearing.

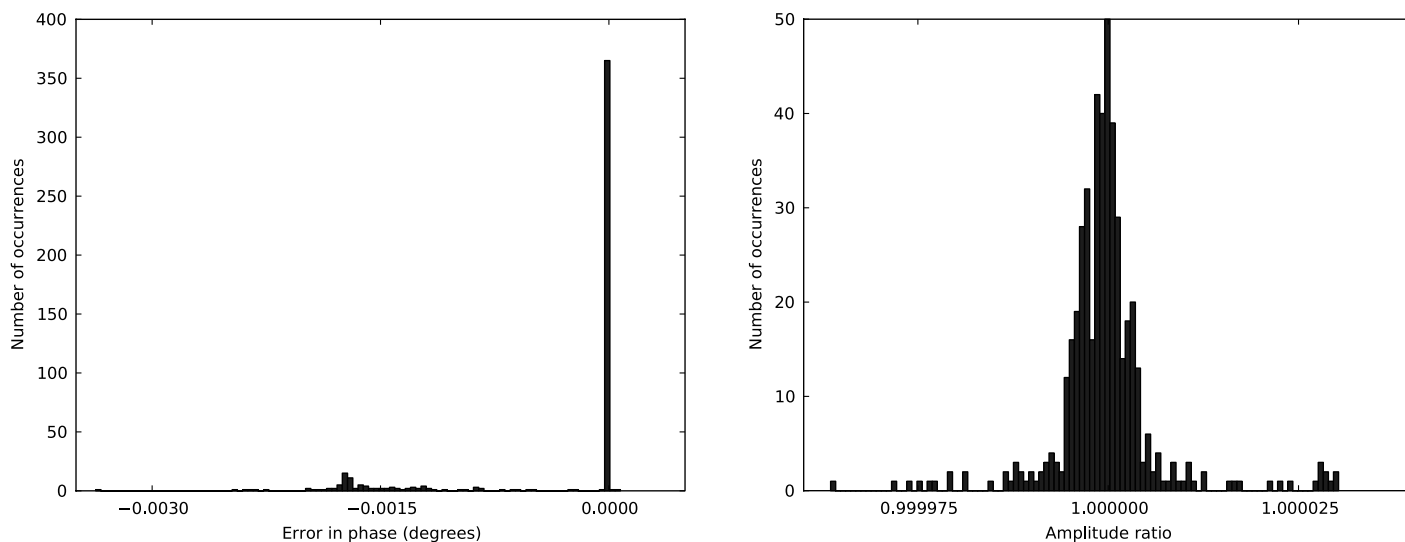


FIG. 7.—Comparison of the correlated output of DiFX-2.0.0 and DiFX-1.5.4 for a single observing band for a pair of VLBA baselines over a 40 s scan. The same baseband data were used as in Fig. 6. *Left*: Histogram showing phase difference between the two correlator outputs; 0.002° corresponds to a delay error of approximately 10^{-16} s. *Right*: Histogram showing fractional amplitude difference between the two correlator outputs. See the electronic edition of the *PASP* for a color version of this figure.

At very large numbers of phase centers, the visibility data rate to disk becomes large, even with heavy spectral and temporal averaging. In this test, 16 spectral points per band and 4 s averaging were used, yielding a data rate per phase center of 70 kbyte s^{-1} of recorded data. For 500 phase centers, this makes a total output data rate of 35 Mbyte s^{-1} of recorded data, which is not inconsiderable when the overhead of writing to an array of different output files is considered. This limitation can be overcome with minimal outlay by writing the output visibilities to commercially available high-speed low-latency disk arrays and by using a file system optimized for large numbers of file operations.

4. OPERATIONAL INFRASTRUCTURE

Considerable effort has been expended to improve the usability of DiFX in routine operations. Specifically, the configuration of correlator jobs has been simplified, as has the generation of the correlator model, and extensive monitoring and logging has been added. Some elements of the new infrastructure are specific to the VLBA installation, but can be easily customized in many cases to suit the needs of a different installation. All of the packages described subsequently were developed in the latter stages of the DiFX-1.x series and are available both in DiFX-1.5.4 and DiFX-2.

4.1. Correlation Configuration

The pathway for automatic configuration of the correlator control files has been considerably improved since DiFX-1.x. A new program, *veX-2difx*, can populate the entire set of nec-

essary correlator files based only on the “*vex*”¹⁷ observation description file, while default values for parameters such as integration time and spectral resolution can be overridden as desired. For the VLBA installation of DiFX-2, ancillary information not available at scheduling time (such as data module names, Earth orientation parameters, stations clocks, etc.) is provided automatically to *veX-2difx*, making use of the operational database available at the VLBA. For other arrays, these necessary inputs can be provided to *veX-2difx* by hand or using a similarly customized script.

4.2. Model Generation

The model generator used by DiFX-1.x, which was based on a customized implementation of *CALC*¹⁸ with limited support, has been surpassed by a more flexible client/server architecture. In this approach, the DiFX model file writing is handled by a standalone program that is completely divorced from the *CALC*-based delay calculations, which are performed with a standard installation of *CALC9* and communicated upon request. This *calcserv* program was already widely used at VLBI observatories, including the VLBA and MPIfR, and has now been adopted by the LBA. Accordingly, it has the advantage of easier integration with existing observatory setups. Concurrently with this change, DiFX-2 has been enhanced to directly read the polynomial-based delay models generated by the model client and stored in the FITS IM and MC tables.

¹⁷ See <http://www.vlbi.org/vex/>.

¹⁸ See <http://gemini.gsfc.nasa.gov/solve/>.

This ensures a perfect match between the recorded and applied geometric model. In contrast, DiFX-1.x read sampled delay files and used a low-order interpolation between the sampled points, which resulted in errors of the order of one-tenth of a femtosecond. These very small errors were discovered in the detailed comparisons presented in § 5. In addition, the use of a polynomial-based model representation saves disk space and memory, as it is more compact than the sampled-delay representation used by DiFX-1.x.

4.3. Correlation Monitor, Control, and Archiving

As part of the DiFX-2 development effort, a standardized message package (difaxmessage) was created. This uses multi-cast XML messages to broadcast the state of various resources involved with a correlation, as well of the progress of the correlation itself. Resource messages include the CPU and network utilization and, potentially, error messages describing equipment failure or misconfiguration. The messages are graduated in importance from “debug” to “fatal.” A configurable

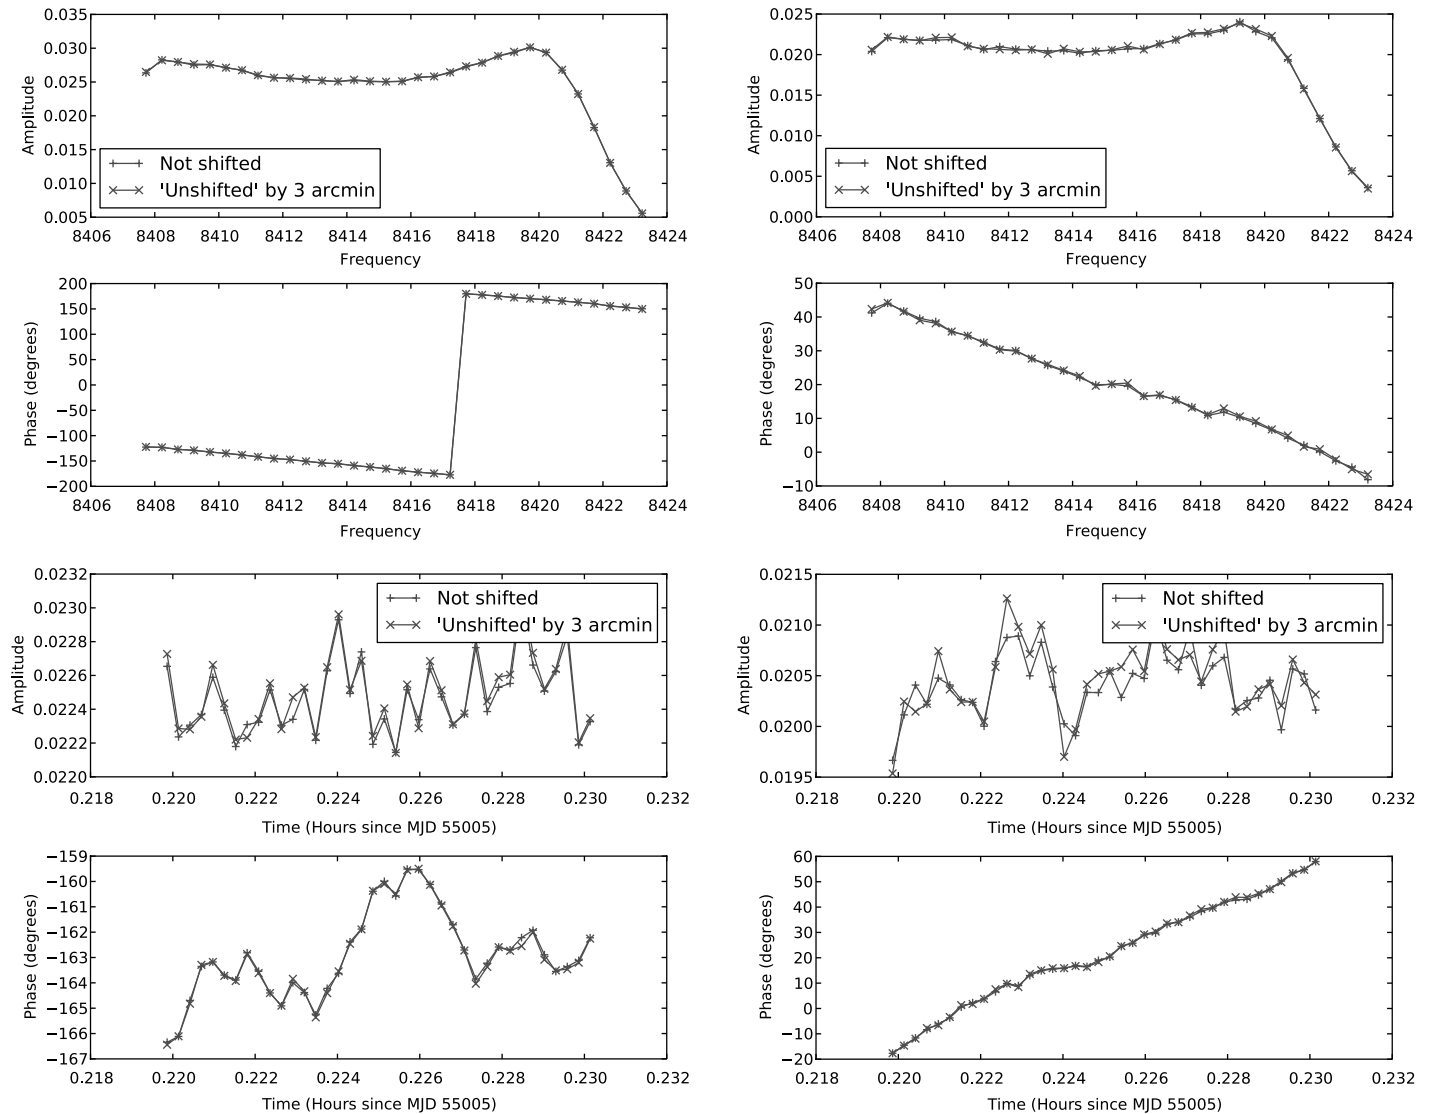


FIG. 8.—Comparison of the correlated output of an ordinary correlation (*solid line*) and one where the visibilities have been shifted back from an initial correlation center $3'$ away (*dashed line*) for a single observing band for a pair of VLBA baselines. The same baseband data and baselines were used as for Fig. 6. *Left:* Visibility amplitude and phase (averaged across the band) vs. time. *Right:* Visibility amplitude and phase (averaged for the scan duration) vs. frequency. The rms amplitude difference is 0.09%, and the rms phase difference is 0.014%. For each baseline the sensitivity loss can be roughly estimated using the ratio of the rms of the visibility amplitude over time. For the Brewster to Fort Davis baseline the sensitivity loss is $\sim 3\%$ (against an expectation of 2%), while for Brewster to Saint Croix the sensitivity loss is $\sim 25\%$ (expectation 23%). See the electronic edition of the *PASP* for a color version of this figure.

display and logging program for these messages has been created to meet the inspection needs of different users.

This message package is also used by other programs available in DiFX-2 to enable an immediate halt to correlation, the quarantining of resources, and other useful miscellaneous tasks. VLBA operations has developed a number of tools that are specific to the VLBA installation, including a graphical user interface that allows correlator jobs to be queued and monitored. These site-specific utility programs are also available as a starting point for adaptation to local needs.

5. VALIDATION TESTING

The initial release of DiFX was tested against three correlators in a selection of observing modes (Deller et al. 2007; Tingay et al. 2009). During its adoption by NRAO, DiFX-1.x was subjected to much more extensive testing and comparison against the VLBA hardware correlator. As with the earlier tests, the NRAO validation scheme for DiFX was primarily composed of point-by-point visibility comparisons, but many more recording modes (combinations of bandwidths, spectral resolution, and integration time) were compared and, unlike earlier validation correlations, also included a number of functional tests, where the final observable from astrometric or geodetic observations were compared between correlators. The VLBA DiFX test plan is described in detail by Romney et al. (2009).

By the time of the adoption of DiFX-2, the VLBA hardware correlator had already been retired, and so the primary validation of the DiFX-2 correlator was undertaken against the operating DiFX-1.x installation at NRAO. At the time of the tests, the specific versions in use were DiFX-1.5.4 (production) and DiFX-2.0.0 (testing). A representative series of comparison plots is shown in Figure 6, detailing the excellent agreement between the correlators. The comparison shown here utilized 40 s of data on the bright calibrator 4C39.25 (J0927+3902), with one 16 MHz wide recorded band spanning the frequency range 8407.49–8423.49 MHz. The integration time was 1 s. A 256 point FFT was used, with the resultant 128 spectral points averaged down to 32 in the FITS file. These 32 spectral points were further averaged across the whole band for a statistical analysis, which showed that the rms phase deviation between the two correlators was 0.0007° , and the rms amplitude deviation was 0.0007%.

Histograms of the errors in phase and fractional amplitude are shown in Figure 7. The histogram shows that the phase errors, while exceedingly small, do not have zero mean in this instance (the mean amplitude error, on the other hand, is less than one-tenth of the rms). This discrepancy is due to delay errors at the 0.1 fs level in DiFX-1.x, which used a second-order delay interpolator at 1 s timescales (DiFX-2 uses a fifth-order interpolator at 2 minute timescales, which is considerably more accurate). It is worth noting that errors of this magnitude cannot even be discerned in comparisons between DiFX and hardware correlators, such as that formerly used by the VLBA, due to the

coarser fringe rotation and internal precision used by the hardware correlator.

This same 40 s time range of data was also used to verify the correct functioning of the multiple-phase-center code. The correlation center for 4C39.25 was shifted by $2'$ in declination and $2'$ in right ascension, and a phase center was added at the true position of 4C39.25. The FFT size was increased by a factor of 4 to 1024, and the uv shifts were applied at a maximum interval of 40 ms. However, the final spectral and temporal resolutions were left unchanged at 0.5 MHz and 1 s, respectively. A $\sim 3'$ shift leads to differential baseline delays of up to $\sim 10 \mu\text{s}$ for the VLBA. For the correlation parameters chosen, this results in time and bandwidth decorrelations of 30% and 7%, respectively. Obtaining higher up-front time and frequency resolutions to greatly reduce the smearing would be straightforward, but these modest parameters were chosen to illustrate the correctness of the phase and amplitude compensation. Figures 8 and 9 repeat the visibility-to-visibility comparisons previously made for the DiFX-1.5.4 to DiFX-2.0.0 comparison.

Across all six baselines used for the test, the mean amplitude and phase error were 0.09% and 0.014° , respectively. In each case, this is less than one-tenth of the rms deviation observed in these quantities (0.9%, 0.46°). Taking the rms of the visibility amplitude over time (per baseline) as a proxy for sensitivity, the decorrelation due to time and bandwidth effects can be estimated. For each baseline, the decorrelation estimated in this manner is consistent with the predicted values to within several percent, and the average across all baselines agrees to 0.4%. Taken in conjunction with the perfect amplitude agreement (after the application of amplitude corrections for decorrelation), this shows that the uv shift is performing exactly as expected.

6. FUTURE WORK

The development of DiFX-2 will continue past the 2.0.0 version described in this article. Planned new functionality includes the ability to form one or more phased-array outputs, as an alternative (or in addition) to a normal cross-correlation. This would allow DiFX-2 to produce tied-array beams for high time resolution studies, such as pulsar analysis. Both a digital filter bank (with tunable integration length) and a reconstructed time series are planned to be selectable outputs from the phased array. A VLBI-capable phased-array system is necessary for the Large European Array for Pulsars (LEAP)¹⁹ project, and DiFX may be used for this application.

Other new functionality under investigation includes frequency-division multiplexing for improved performance with larger numbers of antennas, efficient support for numbers of spectral points per sub-band that are not a power of two, and expanded graphical correlation monitoring and display. As

¹⁹ See <http://epta.jb.man.ac.uk/leap.html>.

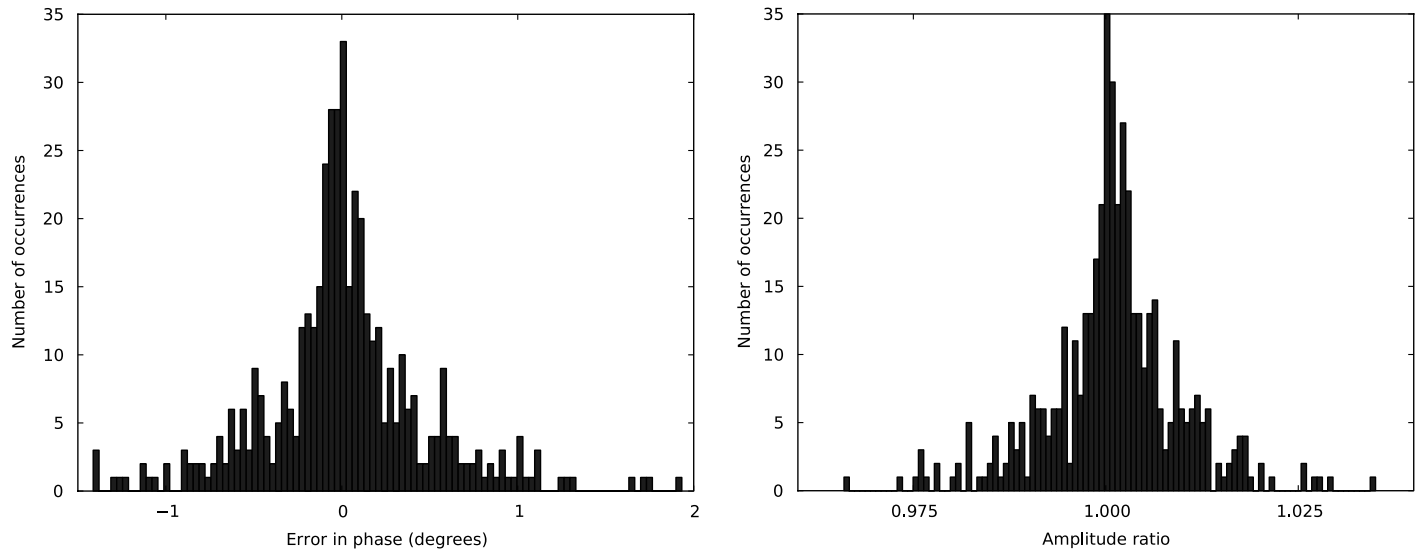


FIG. 9.—Comparison of the correlated output of an ordinary correlation and one where the visibilities have been shifted back from an initial correlation center $3'$ away for a single observing band for a pair of VLBA baselines. The same baseband data and baselines were used as for Fig. 7. *Left:* Histogram showing phase difference between the two correlator outputs. *Right:* Histogram showing fractional amplitude difference between the two correlator outputs. See the electronic edition of the *PASP* for a color version of this figure.

has been the case to date, development will be driven by the needs of the DiFX community, and it is likely that future applications will arise that are not presently envisaged.

7. CONCLUSIONS

A number of significant improvements have been made to the DiFX software correlator since its public release in 2007. These have encompassed improved robustness, greater performance, and the addition of several valuable new features. In particular, the DiFX-2 series now supports phase-calibration-tone extraction, multiple simultaneous phase-center correlation, and the production of high time resolution filter-bank and kurtosis data for use in transient searches and RFI mitigation. In certain areas of parameter space, such as deep VLBI surveys, these new features allow processing speed improvements in excess of a factor of 100. Collectively, all of these new features reinforce the advantages that software correlators possess over custom-designed hardware correlators. DiFX-2 has been adopted by

three major VLBI correlator facilities for production usage and by numerous other institutes and individuals for experimental use. DiFX was a key element in facilitating major bandwidth expansions at both the LBA and VLBA, and it offers the chance for expanded resource sharing and improved robustness in worldwide VLBI. Development of DiFX is expected to continue in the future, with the possibility of application to more existing and upcoming arrays.

A. T. D. is supported by a National Radio Astronomy Observatory Jansky Fellowship. The International Centre for Radio Astronomy Research is a Joint Venture between Curtin University of Technology and the University of Western Australia, funded by the State Government of Western Australia and the Joint Venture Partners. S. J. T. is a Western Australian Premiers Fellow, funded by the State Government of Western Australia. The authors gratefully acknowledge the assistance of Jan Wagner in the development of phase-calibration-tone extraction algorithms.

REFERENCES

- Chikada, Y., Ishiguro, M., Hirabayashi, H., Morimoto, M., Morita, K.-I., Kanzawa, T., Iwashita, H., Nakazima, K., et al. 1987, *Proc. IEEE*, 75, 1203
- Deller, A. T. 2010, *POS*, in press
- Deller, A. T., Tingay, S. J., Bailes, M., & West, C. 2007, *PASP*, 119, 318
- Lenc, E., Garrett, M. A., Wucknitz, O., Anderson, J. M., & Tingay, S. J. 2008, *ApJ*, 673, 78
- Macquart, J. P. R., Hall, P., & Clarke, N. 2010, in *Proc. ISKAF2010 Science Meeting (039; Trieste: SISSA)*, <http://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=112>
- Middelberg, E., Deller, A. T., Morgan, J. S., Rottmann, H., Alef, W., Tingay, S. J., Norris, R., Bach, U., et al. 2010, *A&A*, submitted
- Morgan, J. S., Alef, W., Briske, W. F., Deller, A. T., Mantovani, F., Middelberg, E., Nanni, M., & Tingay, S. J. 2010, *A&A*, submitted
- Nita, G. M., Gary, D. E., Liu, Z., Hurford, G. J., & White, S. M. 2007, *PASP*, 119, 805
- Phillips, C. J., Deller, A., Amy, S. W., Tingay, S. J., Tzioumis, A. K., Reynolds, J. E., Jauncey, D. L., Stevens, J., et al. 2007, *MNRAS*, 380, L11

- Romney, J. D. 1999, in ASP Conf. Ser. 180, *Synthesis Imaging in Radio Astronomy II*, ed. G. B. Taylor, C. L. Carilli, & R. A. Perley (San Francisco: ASP), 57
- Romney, J. D., Briske, W. F., Deller, A. T., Durand, S. J., & Walker, R. C. 2009, *VLBA Sensitivity Upgrade Memo Series*, (Memo 26; Charlottesville: NRAO)
- Thompson, A. R., Moran, J. M., & Swenson, G. W., Jr. 1994, *Interferometry and Synthesis in Radio Astronomy* (New York: Wiley)
- Tingay, S. J., Alef, W., Graham, D., & Deller, A. T. 2009, *J. Geod.*, 83, 1061
- Wayth, R., et al. 2011, *ApJ*, submitted
- Whitney, A. R. 2003, in ASP Conf. Ser. 306, *New Technologies in VLBI*, ed. Y. C. Minh (San Francisco: ASP), 123
- Whitney, A. R., Cappallo, R., Aldrich, W., Anderson, B., Bos, A., Casse, J., Goodman, J., Parsley, S., et al. 2004, *Radio Sci.*, 39, 1007
- Whitney, A., Kettenis, M., Phillips, C., & Sekido, M. 2009, in *Proc. 8th International e-VLBI Workshop*, (042; Trieste: SISSA), <http://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=82>