

Geodetic VLBI correlation in software

Steven John Tingay · W. Alef · D. Graham ·
A. T. Deller

Received: 13 October 2008 / Accepted: 25 May 2009 / Published online: 10 June 2009
© The Author(s) 2009. This article is published with open access at Springerlink.com

Abstract Correlation algorithms for geodetic very long baseline interferometry (VLBI) can now be effectively implemented on parallel computers of modest size. We have undertaken a detailed comparison of the output from a trusted geodetic correlator, one that has supported global geodetic VLBI observations for many years, with the output of a software correlator implemented on a small parallel computing cluster. We show that the correlator outputs agree closely, within expected error bounds, after accounting for the differences in the adopted geometric delay models, and therefore that use of the software correlator is feasible for geodetic VLBI processing, as a first step toward routine geodetic data processing. Recent developments in software correlation for geodesy are discussed, including the possibility of real-time processing options.

Keywords Geodesy · Very long baseline interferometry (VLBI) · Software correlation

1 Introduction

Very long baseline interferometry (VLBI) is a technique which allows the highly precise definition of a reference

frame (the International Celestial Reference Frame: ICRF) with respect to the most distant objects visible in the Universe, powerful sources of radio emission at the centres of galaxies. By virtue of their great distances, these objects, and the frame of reference they define, represent the closest realisation of an inertial reference frame possible (Ma et al. 1998). A terrestrial reference frame has been defined relative to the ICRF, the International Terrestrial Reference Frame (ITRF). Within this frame, regular VLBI observations are used to precisely monitor the positions of the radio telescopes making the VLBI observations (telescopes separated by hundreds or thousands of kilometres), tracking the motions of the telescopes relative to each other, due to geophysical effects. VLBI networks are therefore one of the standard geodetic tools.

The VLBI technique utilises methods of recording the signals at each individual telescope, now almost exclusively using computer hard disks. These recordings are physically brought together (disks are shipped) and played back into a processing device, known as a correlator, that combines the signals in such a way that delivers the fundamental observables of geodetic VLBI, the time delays between signals arriving at the individual telescopes as a function of time. VLBI is a very sensitive method of measuring the delay because each interferometer pair of telescopes in an array accurately measures the relative phase of the signals at the two telescopes, and the phase is related to the delay by the following equation:

$$\phi = 2\pi\nu\delta, \quad (1)$$

where ϕ is the phase difference in the electromagnetic signals between the two telescopes, ν is the frequency of the electromagnetic signals (typically in the range $1\text{--}10 \times 10^9$ Hz), and δ is the delay between the two telescopes in seconds. As ν is a very large number, small changes in δ produce a large and

S. J. Tingay (✉)
Curtin Institute of Radio Astronomy,
Curtin University of Technology,
GPO Box U1987, Perth, WA 6845, Australia
e-mail: s.tingay@curtin.edu.au

W. Alef · D. Graham
Max Planck Institut fuer Radioastronomie,
Auf dem Hügel 69, 53121 Bonn, Germany

A. T. Deller
Swinburne University of Technology,
Centre for Astrophysics and Supercomputing,
H39, P.O. Box 218, Hawthorn, VIC 3122, Australia

easily measurable change in ϕ . A simple method to measure the delay between two antennas is therefore to measure the change in phase across a range in frequency, since the delay is not a function of frequency and

$$\frac{d\phi}{d\nu} = 2\pi\delta. \quad (2)$$

Geodetic VLBI obtains very accurate measurements of delay by measuring the change in phase over as wide a range in frequency as possible.

Because the telescopes operate on the surface of the Earth (with the ionosphere and wet and dry components of the atmosphere intervening between telescope and radio source), which moves rapidly with respect to the celestial reference frame, the correlator uses a geometric model for the changing delays between telescopes to predict the delay (approximately) at which the correlated signal should appear. After application of this geometric model, the correlator measures the residual delay. The model delay plus the residual delay forms the total delay, from which geodetic analysis can proceed (Shapiro 1976).

Geodetic VLBI arrays have typically consisted of between 5 and 20 radio telescopes, recording data at rates up to 1 Gbps per telescope. Traditionally, the correlators that have processed these data have been large purpose-built machines, based on application-specific integrated circuits (ASICs) to perform the digital signal processing required to align the recorded data streams (according to the predicted geometric delay model) and perform the required signal processing.

This scale of computational problem can now be handled comfortably by modest parallel computers, in particular relatively inexpensive clusters of commodity computers containing standard processors. The correlation algorithm, instead of being hardwired into ASICs, can be implemented using standard high-level computing languages in an efficient and effective manner on such machines, giving rise to a software correlator.

Elsewhere is described a particular implementation of a software correlator, the Distributed FX (DiFX) correlator (Deller et al. 2006). DiFX is a software correlator that is currently in use at a number of facilities around the world, for astronomical VLBI. DiFX is also being developed as a correlator for geodetic VLBI and the purpose of this paper is to show that this is a feasible goal. We have taken data recorded for an experiment observed with the US National Radio Astronomy Observatory (NRAO) Very Long Baseline Array (VLBA) and correlated them using both a traditional VLBI correlator, the MarkIV correlator of the Max Planck Institut fuer Radioastronomie in Bonn, Germany, and DiFX. The MarkIV correlator has been used heavily for geodetic VLBI processing for many years as part of the International VLBI Service (Schlüter and Behrend 2007) and is one of a

very small handful of trusted geodetic correlators around the world.

Deller et al. (2006) contains a comparison of DiFX to the VLBA correlator (and the now decommissioned ATNF S2 correlator) that shows close agreement between these two processors. The purpose of the current paper is to demonstrate similar agreement between DiFX and the MarkIV correlator in Bonn. While DiFX and the VLBA correlators share the so-called FX architecture (Fourier transform of the time series voltages from each antenna [F] before cross-multiplication [X]), the MarkIV correlator uses a different hardware implementation, architecture and correlation algorithm, using the XF order of operation (as did the now defunct ATNF S2 correlator, albeit with a completely different hardware and algorithm implementation). Furthermore, subtleties in the application of delays at the different correlators and the different code bases used for data processing can potentially introduce differences between correlator outputs (including via post-correlation processing pipelines). It is therefore a highly valuable exercise to undertake as many independent comparisons of correlators as possible to identify and understand these potential subtle differences. In this way, we can show consistency between not just pairs of correlators, but we can, over time, establish groups of correlators across which results can be compared with confidence. This is an important exercise for science archives based on data from different correlators (i.e. astrometric, geodetic and astronomical archives).

Deller et al. (2006) presents the innovative approach afforded by DiFX in a strongly astronomical context. Here we place DiFX in a geodetic VLBI context, by comparison to the Bonn MarkIV correlator, a processor whose primary purpose has traditionally been geodetic in nature.

We show that the MarkIV and DiFX outputs agree closely, after taking account of differences in the geometric delay models adopted by the two correlators. This means that the two correlators produce the same total delay output and that the use of DiFX for geodetic VLBI is feasible. In Sect. 2 we describe the data used in the comparison, the correlation of the data on both the MarkIV correlator and in DiFX, and the methods used to perform the comparison. In Sect. 3 we present and discuss the results of the comparison. In Sect. 4 we conclude by discussing current and future developments in software correlation for geodesy, specifically with respect to the DiFX implementation, including the possibility of real-time geodetic correlation using DiFX in conjunction with high-speed optical fibre data networks between telescopes.

We emphasise that this demonstration is a first test of the DiFX system for geodetic VLBI and an extended period of geodetic usage and comparison with the output of trusted correlators, using the full standard geodetic VLBI data processing pipeline, is required as the next level of demonstration.

Table 1 Clock offsets and rates for the VLBA antennas

Antenna	Clock offset (μs)	Clock rate ($\mu\text{s/s}$)
FD	3.553	6.732×10^{-8}
PT	-0.644	9.167×10^{-9}
OV	-5.260	1.106×10^{-8}
KP	-3.108	-1.582×10^{-8}

In a subsequent paper we will report on the results of a full geodetic analysis based on data from DiFX and the MarkIV correlator in Bonn.

2 Description of the correlator comparison

2.1 The data

The data used for the correlator comparison was a subsection of the data collected for an experiment conducted with the NRAO Very Long Baseline Array (VLBA¹), BM261, observed on 2007 July 03. Data from four antennas during a 1-minute time range when a strong source was observed (3C454.3: 11:23:00 UT–11:24:00) were extracted from the full dataset for the purposes of the comparison: Fort Davis (FD), Pie Town (PT), Owens Valley (OV), and Kitt Peak (KP). The experiment was performed with a 4-IF (intermediate frequency) frequency sequence of: 1358.49, 1374.49, 1390.49, 1406.49 MHz. All IFs had both upper (USB) and lower sidebands (LSB) of 8 MHz each. Orthogonal circular polarisations were recorded for each IF (RCP and LCP). For the comparison the LSB RCP data for one IF were selected (1358.49 MHz).

The data at the VLBA antennas were recorded to Mark5 disk packs and transported to the MarkIV correlator at MPIfR. Pre-correlation clock models were derived for each antenna, consisting of a clock offset and rate, typically based on measurements at the antenna sites (using a comparison of the antenna clock with GPS, for example). The clock model provides a component of the overall delay model that is applied at the correlator. The clock offsets and rates used at each telescope are listed in Table 1.

2.2 Correlation with the MPIfR MarkIV correlator

The Mark IV correlator (Whitney et al. 2004) at the MPIfR is one of the major correlators for geodetic VLBI. It is used about 50% of its time for geodetic correlation. The Mark IV correlator is ideally suited to verify the DiFX software

correlator for geodetic VLBI. The MarkIV correlator uses the so-called XF architecture, meaning that the time series voltages from each telescope are cross-multiplied (X) before being Fourier transformed (F), as opposed to the architecture adopted for the DiFX correlator, which uses an FX order of operation.

For the correlation the VEX (VEX = ‘VLBI Experiment’) control file was created from the VEX observing schedule in a standard manner. Earth orientation parameters UT1 and polar motion as well as clock offsets and rates for the antennas were added and verified. The correlation was undertaken with 128 lags per 8 MHz sub-band and 2 s integration time. The data were later transformed to the cross-spectral domain using the HOPS (Haystack Observatory Postprocessing System) fringe fitting program ‘fourfit’. All phase-cal information was deleted from the MarkIV correlator output data in this process, as DiFX cannot yet extract the phase-cal signals. Phase-cal extraction is being added to the DiFX code base, along with other additions specifically targeted at geodetic VLBI, in preparation for a demonstration of full geodetic pipeline processing using DiFX. Phase-cal extraction is not, however, required to demonstrate the fundamental capability of DiFX to provide reliable data for geodesy.

The cross-spectra were read into the NRAO data-processing software AIPS using the task MK4IN (Alef and Graham 2002), which also extracts the geometric model used in the correlator and stores it along with the data in a so-called CL table. As MK4IN adds upper and lower sidebands together into one so-called AIPS IF (sub-band), they had to be separated again using the AIPS task UVCOP. The selected data were exported from AIPS using task FITAB.

2.3 Correlation with DiFX

The correlation of the data with DiFX was configured to be identical to the configuration of the MarkIV correlator, i.e. 128 frequency channels per 8 MHz band and 2 s integration time. A detailed description of the DiFX correlator is given in Deller et al. (2006). Due to a difference in the sign definition assigned to the clock offsets and rates, the DiFX correlation required that the values listed in Table 1 be negated. The correlated data were written to RPFITS format and read into AIPS using task ATLOD, before being converted to FITS format and exported from AIPS using task FITTP.

For this experiment, DiFX was run on a small ad-hoc cluster of ten intel-based machines at the Max Planck Institut fuer Radioastronomie with dual-core 32 bit 3.4 GHz processors, connected via 1 Gbps ethernet.

The actual comparison of the MarkIV and DiFX datasets was made by reading them into MIRIAD (Sault et al. 1995) and using plotting tasks such as UVPLT and UVSPEC within the convenient capabilities for scripting supported by MIRIAD.

¹ The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc.

2.4 Accounting for differences in the delay models

All correlators require a method for the generation of the geometric delay model, which is then applied to align the datastreams from the different telescopes before correlation of the signals. Most correlators used for astronomical and/or geodetic VLBI use the CALC software to generate delay models for this purpose (Gordon 2004). CALC has been developed over a large number of versions (current version is CALC 10) and has a large number of options that describe the type of terms that are included in the delay model calculation. Therefore, for local reasons at each correlator (historical and technical), CALC is run in a number of different ways, resulting in slightly different delay models being used at each correlator.

CALC generates delays between individual telescopes and the plane which is perpendicular to the source direction and which intersects the centre of the Earth (referred to as an antenna-based delay). Delays between two antennas (or baseline-based delays) are simply then the difference between antenna-based delays.

For geodetic analysis, the total antenna-based delay is the important input into the analysis, so as long as the delay model that is applied at the correlator is recorded with sufficient precision, the total delay can be reconstructed from an addition of the model delay generated from CALC and the residual delay that is measured from the correlated data. Small differences in the delay model at each correlator do not, therefore, affect the geodetic analysis. This is because the difference will appear with the opposite sign in the measured residual delay (assuming that both correlators are correctly functioning).

For the purposes of the correlator comparison here, we need to demonstrate that the total delays derived from the two correlators are identical. The approach we take is to simply compare the delay models applied at the two correlators, form the difference between the two models, and apply this difference to the residual delays measured from one of the correlated datasets. If the total delays agree between the two correlators, application of the delay model difference in this way will bring the residual delays (and therefore phases), as measured by the data, into alignment. This is the method as has been used to compare DiFX to the VLBA correlator and the (now decommissioned) ATNF S2 VLBI correlator (Deller et al. 2006).

At the MarkIV correlator, the delay model is generated at correlation time and includes the delay introduced by the clock offsets and rates between telescopes. The delay model is represented as a series of coefficients for a polynomial expansion that describes the variation of delay with time. For this experiment, the polynomial expansion was of the form

$$\delta(t) = a_0 + a_1t + (a_2/2)t^2. \quad (3)$$

The coefficients a_0 , a_1 , and a_2 were generated every 6 s, with $t = 0$ defined as the start of each 6-s period. From these coefficients, the delay model for the MarkIV correlator, for the time range of the comparison, for the four telescopes is plotted in Fig. 1a.

These coefficients are transmitted to the MarkIV correlator at correlation time and the appropriate delays are calculated for application to the datastreams. The coefficients are recorded with the correlated output, so that total delays can be reconstructed.

For DiFX, the delay model is generated prior to correlation time and the results held in a file that is accessed at correlation time. For this experiment, delays were calculated using CALC each second. At correlation time DiFX generates the applied delays by performing a cubic interpolation over three delay values centred on the time for which the delay is required. The total delay can therefore be reconstructed from the delay file plus the addition of the measured residual delay. The DiFX delay model for the four antennas, during the period of comparison, is also shown in Fig. 1a.

Differencing the DiFX and MarkIV delay models on a baseline basis gives quantities which, when converted to a phase difference and appropriately applied to the residual phases from one of the correlated datasets, should bring the two correlated datasets into agreement, if both correlators are performing correctly. Figure 1b shows the difference between the antenna-based delays, in the sense of DiFX–MarkIV. The differences in the baseline-based delays are then shown in Fig. 1c, in the same sense.

Figure 1(b, c) shows substantial differences in the delay models used by the DiFX and MarkIV correlators. For example, the maximum difference in the baseline-based delays in Fig. 1c, for the FD–KP baseline, is ~ 3 ns. For the experimental setup used in this experiment, with 8 MHz bandwidths, this represents a divergence in phase over a single 8 MHz IF of $\frac{3}{62.5}\pi \sim 10^\circ$, which is easily detectable in the correlator output. Furthermore, a change in the phase offset of $\sim 1,500^\circ$ (from Eq. 1) results from this difference between the two models. For our comparison we therefore need to correct both the slope and offset of the phase. Figure 1d shows the predicted phase offset correction required for each baseline in the comparison (modulo 2π).

3 Amplitude component of the visibility

The phase output of a correlator represents half the information content of the correlated data, and is the information of prime interest for geodetic analysis. The other half of the information is contained in the amplitude, which is of no direct use in a geodetic analysis but is of more interest for astronomical analyses. Since the amplitudes are derived from the same real and imaginary output as the phase, the

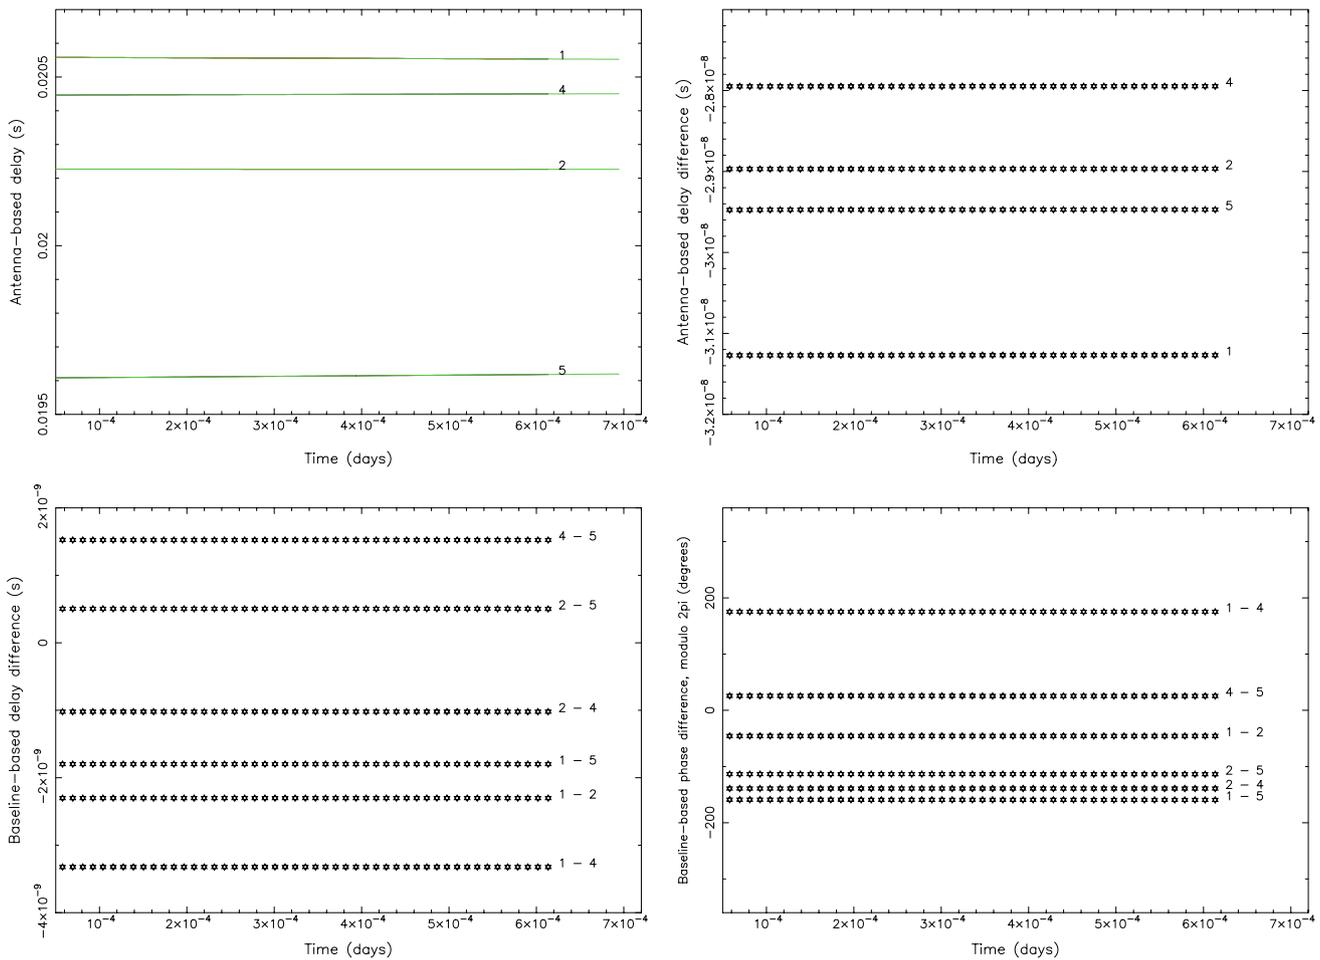


Fig. 1 **a** (Top left) shows the antenna-based delay model values from the two correlators [1=FD; 2=PT; 4=OV; 5=KP]. On this vertical scale, the delays from the two correlators cannot be distinguished. **b** (top right) shows the differences between the antenna-based delay models from the two correlators. **c** (bottom left) shows the differences in the

baseline-based delay models from the two correlators. **d** (bottom right) shows the baseline-based corrections to the phase, derived from the differences in the two delay models (modulo 2π). These phase corrections are used to correct the phase offsets between the two correlated datasets which exists as a consequence of the use of different delay models

amplitudes should in theory be identical in any correlator comparison. The phase output of a correlator is defined as $\frac{1}{2} \tan^{-1} \frac{\text{Im}}{\text{Re}}$, where Im is the imaginary component of the correlator output and Re is the real component. Likewise, the amplitude is defined as $\sqrt{\text{Re}^2 + \text{Im}^2}$, and measures the fraction of the power at the two telescopes that is correlated. For a full description of why and how a correlator produces complex output see, for example, Thompson et al. (1986).

Various subtle corrections to the amplitude output of correlators have to be performed, which depend on how the data were recorded, such as the van Vleck correction for finite precision digitisation of the analog signal at the radio telescopes (van Vleck and Middelton 1966). Historically different correlators have handled these corrections in different manners, with some correlators applying all corrections at correlation time, such as the ATNF S2 correlator, and some

applying corrections as part of post-correlation data reduction, such as the VLBA correlator (Romney 1995).

Different correlator architectures give rise to fundamentally different responses as a function of frequency across the recorded bandwidth. In an XF correlator a single spectral channel has a $\text{sinc}(v)$ response across the band (the sinc arises from the Fourier transform of the lag spectrum generated by an XF correlator). In an FX correlator, the response is $\text{sinc}^2(v)$ since the two data segments are Fourier transformed before being cross-multiplied. The differing responses are reflected in the spectra, especially if strong and narrow frequency width signals are present in the spectra, or if the band-pass filter has sharp cutoffs at the edges.

Finally, the way in which fractional sample error correction is handled in a correlator can affect the amplitude as a function of frequency. The DiFX correlator corrects every

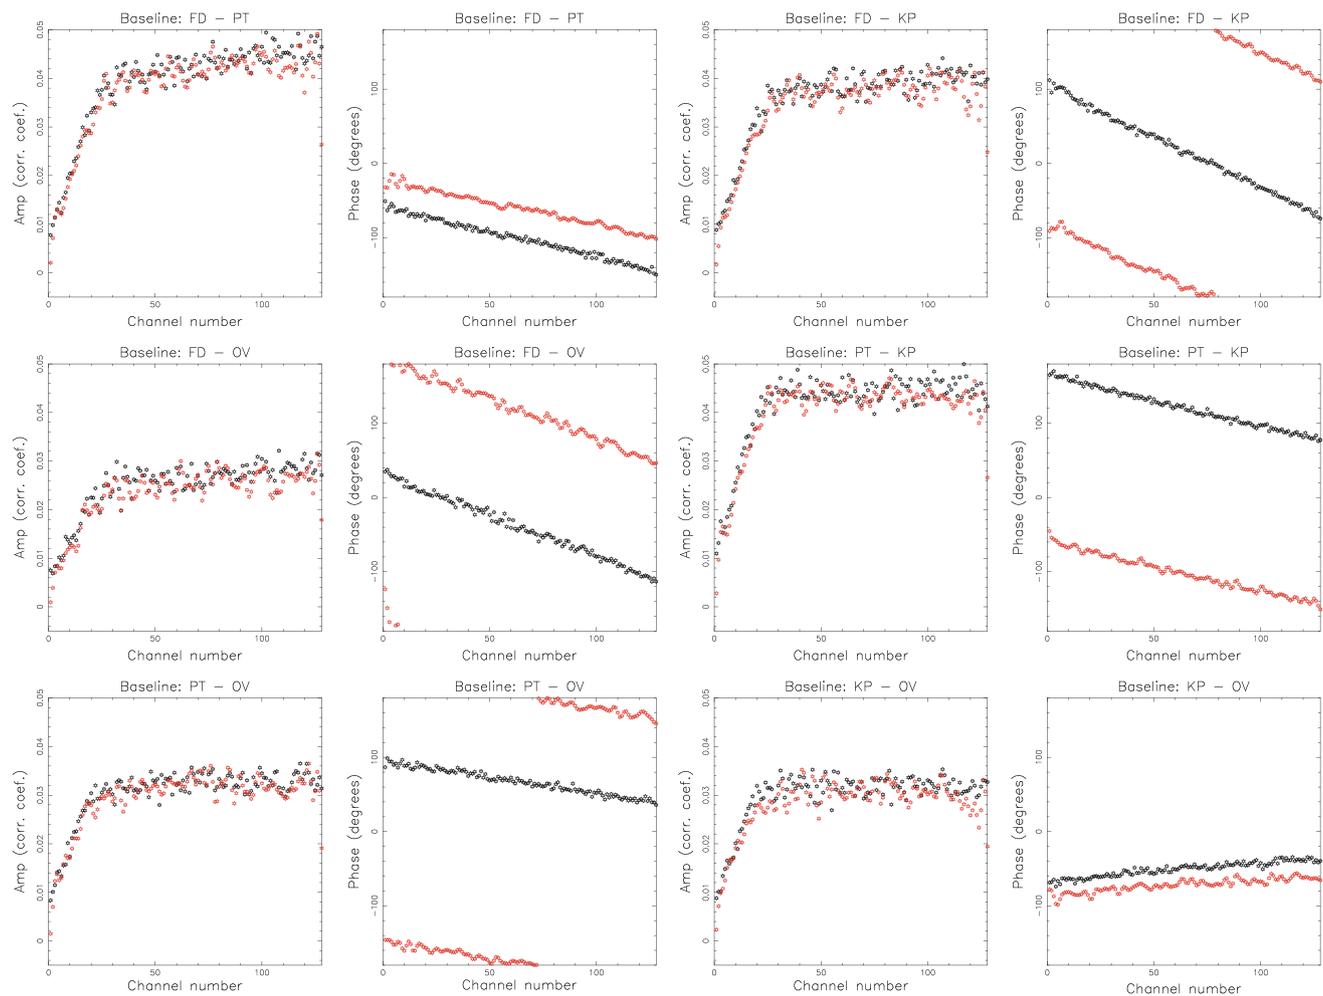


Fig. 2 The amplitude and phase output of the two correlators as a function of frequency for all baselines, shown in pairs of panels (amplitude first and phase second in each pair). *Red points* denote the output of the MarkIV correlator and *black points* denote the output of DiFX. On the FD–PT baseline (*first two panels*), the baseline-based delay differ-

individual cross-correlation for this error, before integration and therefore has essentially perfect response across the band. Correlators that do not implement this step usually suffer signal loss at the edges of the band, with only perfect response at the centre of the band. See [Deller et al. \(2006\)](#) for a full discussion of this effect as part of the DiFX comparison to the ATNF S2 correlator.

For the MarkIV and DiFX comparison, the amplitude corrections are sufficiently similar that we expect the amplitudes to agree to within a few percent. For the purposes of a geodetic correlator, this is more than adequate. Thus, we do not do a detailed comparison of the amplitudes beyond verifying this expectation. For astronomical VLBI, this level of difference is significant but still acceptable, since astronomical observations make reference to sources with known properties to calibrate both the amplitude and the phase in an absolute sense.

ence in the models is clearly detectable by eye in the divergence in the phase across the bandwidth. On all baselines the large phase offsets are obvious. The amplitudes on all baselines agree well, to within the few percent expected from a priori considerations

4 Discussion of results

Figure 2 shows the amplitude and phase output from each correlator as a function of frequency, before any corrections to the phase have been made. The data shown in Fig. 2 represent a single 2-s visibility sample from each correlator. The strength of the source observed means that a good signal to noise detection can be made in such a short integration. The phase shown is plotted between $-\pi$ and π , i.e. modulo 2π . In reality, many turns of phase difference exist between the two datasets.

Figure 3 shows the same data, after correction for the differences in the delay models as described in the previous section. The delay (phase) differences have been applied to the DiFX phases in this case. Despite the large differences between the baseline-based delays (more than 2 ns in some cases, equating to $\sim 1,000$ degrees of phase), extremely good

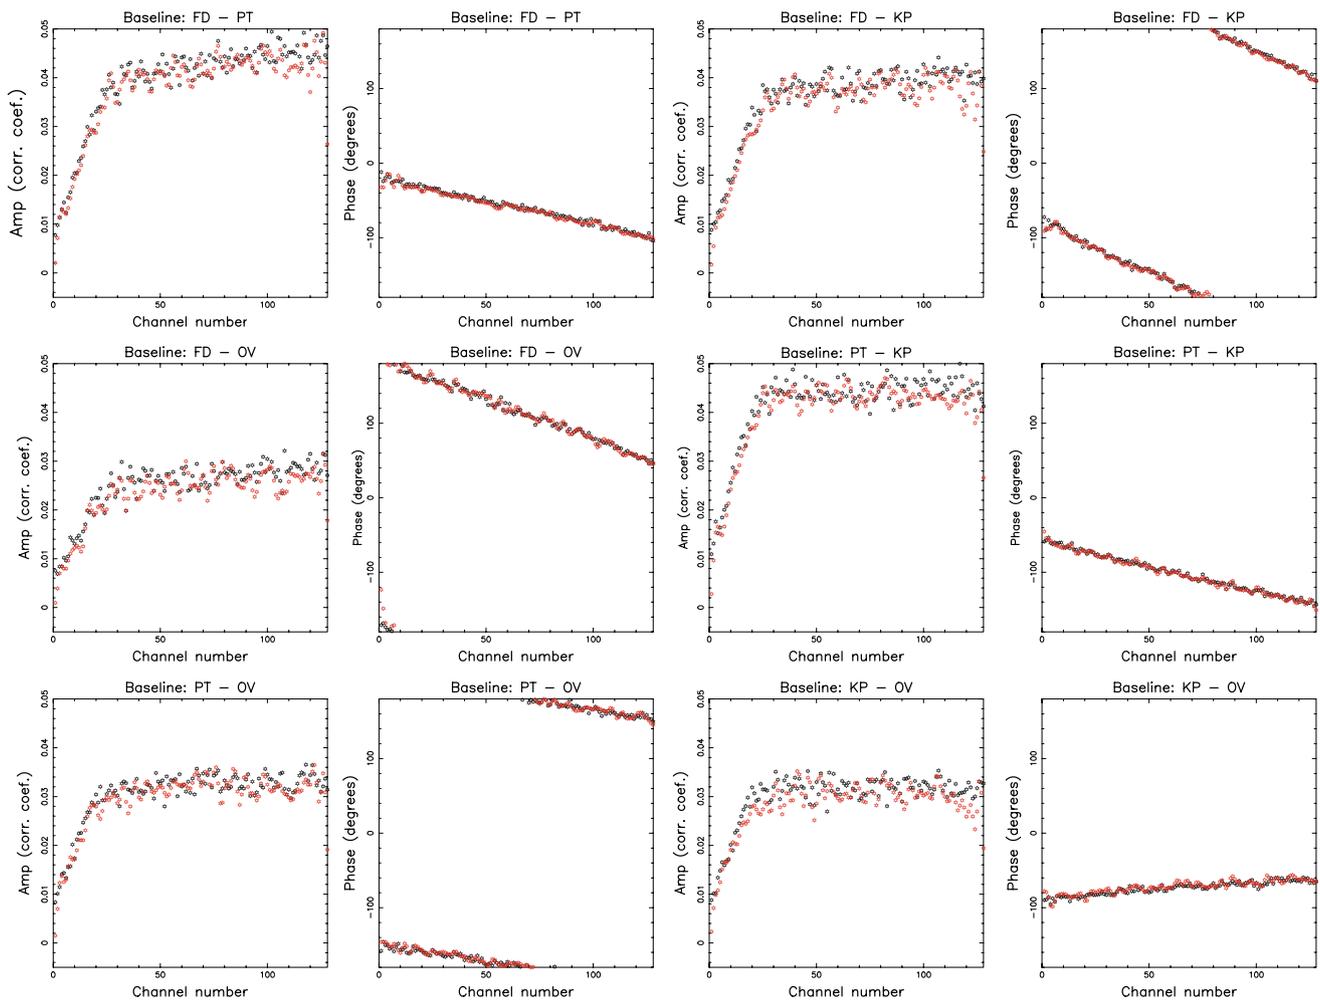


Fig. 3 The same data as shown in Fig. 2, but with phases corrected according to the differences in the correlator delay models derived in Sect. 2.4. Despite the application of large corrections, the final agree-

ment between the two correlators is very good. *Red points* denote the output of the MarkIV correlator and *black points* denote the output of DiFX

agreement has been reached between the phases from the two correlators, both in terms of the slope across the band and the phase offset. As expected, the amplitudes agree closely without correction and will not be considered further in this analysis, as they are not directly relevant to geodetic analyses.

To quantify this result, we difference the phases from the two correlators over the channel range 32–96, to avoid the edges of the band (excluding baselines FD–KP and PT–OV, which both experience phase wraps within this range), and examine these differences. For perfect agreement between the correlators, the phases would have a zero mean and a random scatter around the mean, reflecting the thermal noise errors on the measurements, i.e. with a coefficient of determination (r^2) close to zero from a linear regression on the differences. A comparison of the data on this basis is presented in Table 2.

Table 2 Comparison of phase differences as a function of frequency after correction for delay model differences between the correlators

Baseline	Mean of phase differences (°)	Standard deviation of phase differences (°)	r^2
FD–PT	–1.60	2.46	1×10^{-4}
FD–OV	0.56	4.22	4.5×10^{-3}
PT–KP	–0.74	2.56	1×10^{-6}
OV–KP	1.71	3.14	8×10^{-4}

In all cases, the mean phase difference is well within 1σ of zero, and the r^2 values indicate that far less than 1% of the variation in the data can be explained by linear regression. Thus, no significant differences exist between the phase outputs as a function of time, between the MarkIV and DiFX correlators.

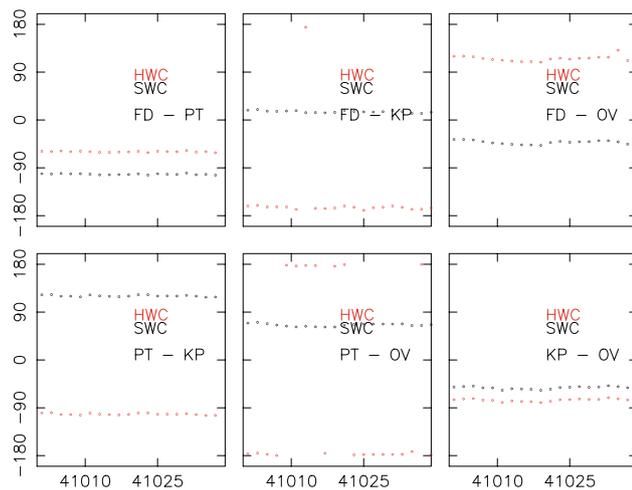


Fig. 4 The phase from the two correlators as a function of time, for all baselines. Variations in the phase with time are caused by variations in the path length through the atmosphere on the line of sight to each telescope and track in the correlator output. However, the same large phase offsets as seen in Fig. 2 can be seen here. *Red points* denote the output of the MarkIV correlator (HWC = Hardware Correlator) and *black points* denote the output of DiFX (SWC = Software Correlator)

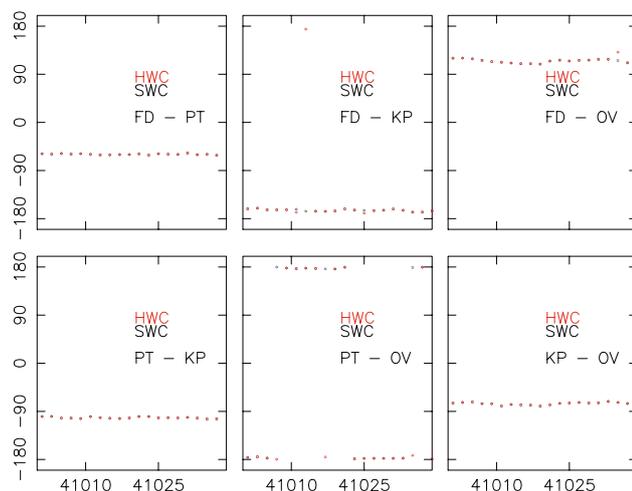


Fig. 5 The same data as shown in Fig. 4, but corrected for the large phase offsets caused by the differences in the delay models at the two correlators. Aside from occasional glitches in the data from the MarkIV correlator, the origin of which are unknown, agreement between the correlators is excellent. *Red points* denote the output of the MarkIV correlator (HWC = Hardware Correlator) and *black points* denote the output of DiFX (SWC = Software Correlator)

Likewise, Fig. 4 shows the phase output of the two correlators, again before any correction, as a function of time. In this case, each dataset has been vector averaged as a function of frequency over the channel range 32–96 (again to avoid the edge channels) before being plotted.

In Fig. 5, the data from Fig. 4 are shown again, after correction for the delay model differences in Figs. 2 and 3. Again, we analyse the differences in phase, as above, with the results

Table 3 Comparison of phase as a function of time after correction for delay model differences between the correlators

Baseline	Mean of phase differences (°)	Standard deviation of phase differences (°)	r^2
FD–PT	−0.234	1.700	4.7×10^{-2}
FD–OV	−0.663	1.676	1.03×10^{-1}
PT–KP	−0.188	1.620	1×10^{-3}
OV–KP	0.050	1.996	2.3×10^{-2}

summarised in Table 3 (for the same baselines as appear in Table 2). At some points in the time range, the data from the MarkIV correlator appear to deviate significantly, for single time samples. Aside from this effect with the MarkIV correlator, (which leads to slightly higher r^2 values for some baselines in Table 3) excellent agreement follows after corrections based on the delay model differences between the correlators.

Therefore, the conclusion from this comparison is that the total delays derived from the DiFX and MarkIV correlators (the sum of the applied correlator delay model and the residual delay information in the correlator output) are fully consistent. This means that it is feasible to use DiFX as the basis for a geodetic correlation system, due to this good agreement with a trusted geodetic correlator.

5 Current and future developments in software correlation for VLBI

Software correlation in radio astronomy, particularly in VLBI, is becoming more popular due to the increased performance and decreased cost of commodity compute platforms. For the typical scale of processing involved in VLBI using moderate numbers of antennas (up to 20) and for typical data rates per antenna (less than 1 Gbps), correlation using a moderate compute cluster is an attractive option. This is also the case for geodetic VLBI.

The DiFX software, in particular, is now being used as the basis for correlators for the Australian Long Baseline Array (LBA) and the NRAO Very Long Baseline Array (VLBA). The MarkIV correlator at Bonn (used as part of our comparison here) will soon be replaced by a cluster running DiFX. The VLBA and Bonn installations of DiFX will be used for both astronomy and geodesy.

DiFX will provide the correlation engine for a new dedicated geodetic array currently being built in Australia, the AuScope array.² The AuScope array will be complemented by a new antenna being built in New Zealand (Gulyaev and Natusch 2007).

² www.auscope.org.au.

One of the advantages of running a correlator such as DiFX on a cluster computing resource, is that it interfaces naturally to high-speed optical fibre networks, where available, allowing data to be streamed directly from the antennas to the correlator in real-time, for instantaneous processing. This mode of operation is known as e-VLBI and in Australia, the LBA has been set up as an e-VLBI array through the use of DiFX (Phillips et al. 2007). The advantages of such a system are the instant feedback of scientific results into the observational program decision making and the reduced need for a large stockpile of expensive hard disk storage.

For geodetic VLBI, these advantages hold, in particular the prospect of real-time correlation and instantaneous processing of the correlated data, for real-time geodetic analysis. Rapid turnaround length of day measurements or real-time sensitive monitoring of crustal movements during events such as earthquakes may be applications for such a system.

Further developments for geodetic usage are being added to the DiFX code base, as a matter of priority, in preparation for a demonstration of DiFX as part of the standard geodetic VLBI processing pipeline. These improvements include the addition of phase-cal extraction and record keeping to attach the full delay model to the correlated data in standard format.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

- Alef W, Graham DA (2002) The New Bonn Mk IV – AIPS data export path. In: Ros E, Porcas RW, Lobanov AP, Zensus JA (eds) Proceedings of the 6th EVN symposium. MPIfR, Bonn, pp 31–33
- Deller AT, Tingay SJ, Bailes M, West C (2006) DiFX: a software correlator for very long baseline interferometry using multiprocessor computing environments. PASP 119:318–336
- Gordon D (2004) CALC: the next upgrade. In: Vandenberg NR, Baver KD (eds) International VLBI service for Geodesy and Astrometry 2004 general meeting proceedings. Ottawa, Canada, February 9–11, 2004. NASA/CP-2004-212255, pp 265–265
- Gulyaev SA, Natusch TJ (2007) Collaboration and development of radio-astronomy in Australasia and the South-Pacific region: New Zealand perspectives. IAUS 5:137–140
- Ma C et al (1998) The International Celestial Reference Frame as realized by very long baseline interferometry. AJ 116:516–546
- Phillips CJ et al (2007) Detection of compact radio emission from circinus X-1 with the first southern hemisphere e-VLBI experiment. MNRAS 380:L11–L14
- Romney JD (1995) Theory of correlation in VLBI. In: Zensus JA, Diamond PJ, Napier PJ (eds) Very long baseline interferometry and the VLBA. Astronomical Society of the Pacific Conference Series, vol 82, pp 17–30
- Sault RJ, Teuben PJ, Wright MCH (1995) A retrospective view of MIRIAD. In: Shaw RA, Payne HE, Hayes JJE (eds) Astronomical data analysis software and systems IV. PASP Conf Series 77:433–440
- Schlüter W, Behrend D (2007) The international VLBI service for geodesy and astrometry (IVS): current capabilities and future prospects. J Geod 18(6–8):379–387
- Shapiro II (1976) Estimation of astrometric and geodetic parameters. Methods Exp Phys 12:261–276
- Thompson AR, Moran JM, Swenson GW (1986) Interferometry and synthesis in radio astronomy. Wiley-Interscience, New York
- van Vleck JH, Middleton D (1966) The spectrum of clipped noise. Proc IEEE 54:2–19
- Whitney AR, Cappallo R, Aldrich W, Anderson B, Bos A, Casse J, Goodman J, Parsley S, Pogrebenko S, Schilizzi R, Smythe D (2004) Mark 4 VLBI correlator: architecture and algorithms. Radio Sci 39:1007–1020