

The Swinburne distributed FX (DiFX) software correlator for astronomical and geodetic VLBI

Adam Deller*, Steven Tingay, Matthew Bailes & Craig West

Swinburne University, Hawthorn VIC Australia

*E-mail: adeller@astro.swin.edu.au, stingay@astro.swin.edu.au,
mbailes@astro.swin.edu.au, cwest@astro.swin.edu.au*

This paper describes the development of a general-purpose software correlator code known as Distributed FX (DiFX). The code was developed primarily to facilitate an upgrade of the Long Baseline Array (LBA) network for Very Long Baseline Interferometry (VLBI) in Australia, enabling an upgrade from the S2 tape-based recording system to a PC-EVN disk-based system. This upgrade has brought 4-8 fold increases in bandwidth and greatly expanded the available correlation parameter space. This expansion in correlator capability has led to a range of new science, both with the LBA and other VLBI arrays outside Australia, and these applications are briefly described.

*The 8th European VLBI Network Symposium
September 26-29 2006
Toruń, Poland*

*Speaker.

1. Introduction

1.1 Function of a correlator

Excellent overviews of correlator functionality are given in [7] and [9], and the operations will only be described briefly here. Essentially, a correlator manipulates the downconverted, sampled data from each telescope to produce the baseline visibilities which are the fundamental interferometric observable, containing information on the spatial frequency structure of the observed source. In order to produce visibilities as a function of frequency, the telescope data streams are appropriately delayed, channelised, compensated for Earth rotation, and cross-multiplied. This process of channelisation, usually via a Fourier transform (F), followed by a cross-multiplication (X), gives rise to the so-called FX correlator algorithm we have implemented in hardware.

1.2 Correlator implementations

Traditionally, correlators have been implemented on Application Specific Integrated Circuits (ASICs), although recently the use of Field Programmable Gate Arrays (FPGAs) has become widespread. In either case, custom design or configuration of the boards and data transport is required, and such implementations are referred to in this paper generically as ‘hardware correlators’. Examples of current VLBI hardware correlators include the LBA S2 correlator [11], the Very Long Baseline Array (VLBA) correlator [8], and the Joint Institute for VLBI in Europe (JIVE) correlator [2].

An alternative to customised hardware is to write software which carries out the correlation operations on non-application specific hardware, such as a Beowulf commodity cluster. A correlator implementation consisting of re-usable code which can run on more than one computing platform is referred to in this paper as a ‘software correlator’. An example of an existing VLBI software correlator is the code developed by the Communications Research Laboratory (CRL) in Japan [6]. Generally speaking, software correlations possess the advantage of much greater flexibility than hardware correlators, as data structures are allocated dynamically, rather than being constrained by the fixed hardware boards.

An intermediate approach could be considered, consisting of a software framework with minimal hardware-specific optimised plug-ins. Such a ‘hybrid correlator’ is planned as an extension to the DiFX project, but is not discussed here.

1.3 Development of DiFX

The DiFX code was initially developed as a more efficient replacement for the XF correlator code used for testing with the LBA from 2004 [10]. It can read and process LBA, MkV and K5 data formats. For LBA operations, it runs on the Swinburne supercomputer¹, which comprises over 200 CPUs connected by gigabit ethernet. It is written in C++, but all computationally intensive sections of the code are implemented using Intel Performance Primitives (IPP). IPP is a library of vector functions optimised for a range of Intel architectures, with a C interface.

The code can run on any number of nodes depending on available resources, and internode communication uses the Message Passing Interface (MPI) standard². On the Swinburne cluster,

¹<http://astronomy.swin.edu.au/supercomputing/>

²<http://www-unix.mcs.anl.gov/mpi/>




Figure 1: The high-level structure of the DiFX code

the mpich³ implementation of MPI is used.

The high level structure of the DiFX code is shown in Figure 1. The correlation is directed by a master node, which also performs long-term integration of visibilities and writes the results to disk in FITS format. Data is buffered in memory and coarsely delayed by datastream nodes, and transported in short (typically 1-10 millisecond) time ranges to core nodes where the actual correlation is performed. The partial visibility results from the core nodes are transported to the master node. Double buffering is used at each transfer stage to ensure that variable network latency does not reduce correlator performance. The spectral and time resolution of the correlator are completely configurable, unlike hardware correlators which have limits derived from the number of boards and output data rates, and sophisticated pulsar binning capabilities are built in.

The structure of the DiFX code is described in detail in [3]. The code and documentation can be downloaded at <http://astronomy.swin.edu.au/~adeller/software/difx/>.

2. Verification and performance

We have undertaken several comparisons with existing hardware correlators to validate the DiFX code. In this paper, we present a small subset of one comparison; further testing results and performance benchmarks are presented in [3].

Figure 2 shows the visibility amplitude and phase obtained from DiFX and from the VLBA hardware correlator for a two-minute window in a 2.3-GHz test experiment observing 0923+392, a strong and compact active galactic nucleus. Delay model differences between the correlators were subtracted from the phase plots. The agreement between the two correlators is excellent. A signal to noise analysis shows that the results agree within the RMS errors.

Benchmarking of the DiFX code has shown that the performance scales linearly with additional compute nodes to 40 CPUs, the largest cluster used for testing to date. Computational load scales linearly with station bandwidth, is relatively constant over moderate spectral resolutions, and scales near-linearly with number of stations (for typical VLBI array sizes), as shown in Figure 3. To illustrate typical requirements, correlating the LBA in real time at maximum data-rate (6

³<http://www-unix.mcs.anl.gov/mpi/mpich1/>

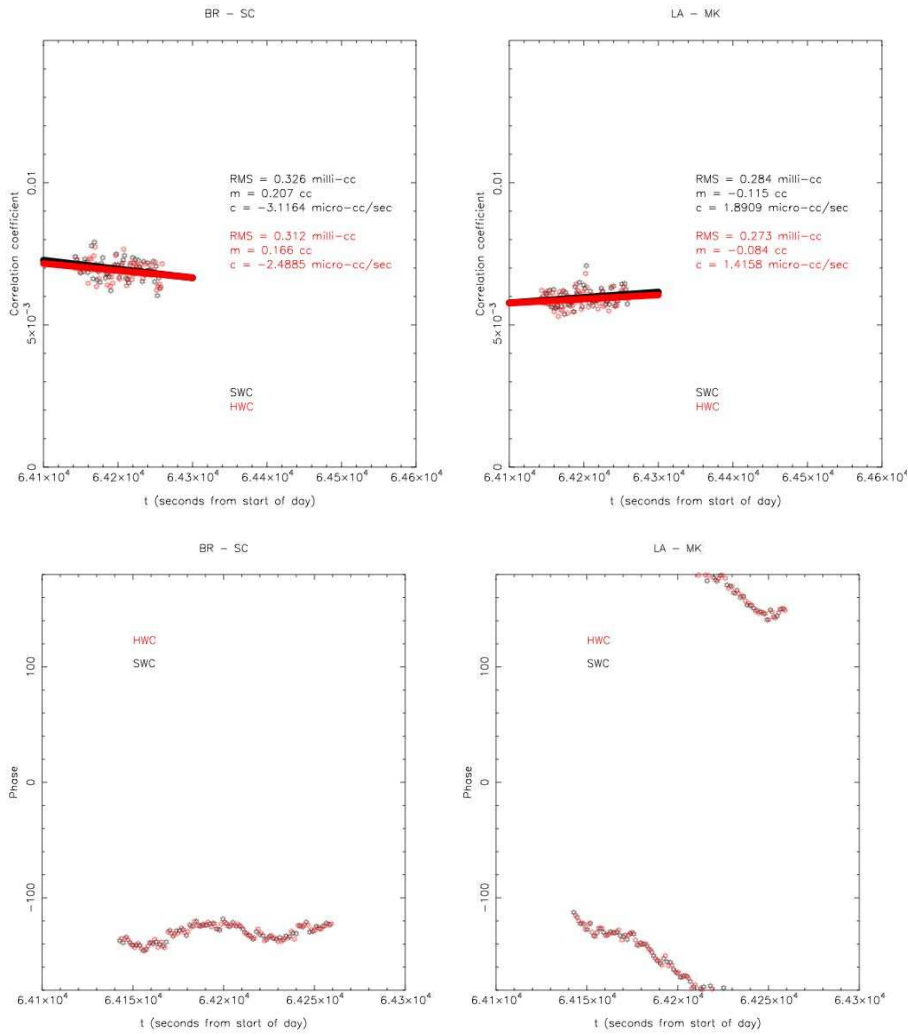


Figure 2: Amplitude (top) and phase (bottom) for two minutes of data on the Brewster-St Croix and Los Alamos-Mauna Kea baselines. The VLBA correlator visibilities are shown in red, and the DiFX visibilities in black.

antennas \times 256 MHz bandwidth, full stokes correlation, 256 spectral points per subband, 1 second integrations) would require 100-200 CPUs⁴.

3. Current applications

For the LBA, the disk-based recording system coupled with DiFX has allowed a major expansion to the array capabilities: $\sim 80\%$ of new proposals request use of the disk system. Some proposals require the increased sensitivity provided by the higher data rate disk system, while others require correlator parameters which the S2 system cannot provide. Examples of the latter

⁴3.0-GHz single core P4, 1 GB RAM




Figure 3: Scaling of computational requirements with array size and spectral resolution. Station-based processing costs dominate for arrays of less than 15 antennas.

include wide field imaging (wide bandwidths for faint sources with high time (≤ 1 second) and frequency (≤ 64 kHz) resolution), maser studies (< 250 Hz frequency resolution) and narrow duty cycle pulsar observations (duty cycle $< 1\%$).

Outside the LBA, DiFX has been used to correlate very high frequency resolution pulsar scintillation experiments [1] and trans-Tasman VLBI tests to an antenna in Auckland [5]. It has also been investigated as a possible upgrade to VLBA capabilities, as a geodetic correlator for existing and future geodetic arrays [4], and as a test correlator for part-time arrays. As a point of comparison, a simple calculation using the performance benchmarks from Section 3 shows that a small, modern 32 node cluster ($\times 4$ cores/node = 128 CPU cores total) could correlate $\sim 20\%$ of VLBA observing time at higher data rates (1 Gbps) and/or with correlation parameters that are currently unavailable to the VLBA correlator.

4. Conclusions

DiFX is a mature, general purpose software correlator which is currently being used successfully to extend the capabilities of the Australian LBA. It can be easily installed and used on a generic commodity computing cluster, and is compatible with all the major disk-based recording formats used in VLBI today. It provides a convenient way to exploit areas of correlator parameter space which are inaccessible to existing hardware correlators. The code has been thoroughly tested against the LBA and VLBA hardware correlators, and performance tests show that it is a competitive option compared to building a new hardware correlator for the small numbers of antennas typical for VLBI arrays.

References

- [1] Brisken, W. et al. 2007, in preparation
- [2] Casse, J. L. 1999, *New Ast. Rev.*, 43, 503
- [3] Deller, A. T., Tingay, S. J., Bailes, M. & West, C., 2007, *PASP*, 119, 318
- [4] Tingay, S. J. & Deller, A. T., 2007, in preparation
- [5] Gulyaev, S.; Natusch, T.; Addis, B.; Tingay, S.J.; Deller, A.T. 2005, *Journal of the Royal Astronomical Society of New Zealand*, 44, 12
- [6] Kondo, T. et al. 2003, in *New technologies in VLBI*, ASP Conf. Ser., Vol. 306. ed Y.C. Minh, San Francisco, CA: Astronomical Society of the Pacific
- [7] Romney, J.D. 1993, in *Proceedings from the 1993 NRAO Summer School*, NRAO Workshop No. 22, ASP Conf. Ser., Vol 82, eds J.A. Zensus, P.J. Diamond, P.J. Napier, 17
- [8] Romney, J.D. 1995, *BAAS*, 27, 815
- [9] Thompson, A.R., Moran, J.M. & Swenson, G.W. 1994, *Interferometry and Synthesis in Radio Astronomy*, Kreiger Publishing Company
- [10] West, C. 2004, M.Sc. thesis, Swinburne University of Technology
- [11] Wilson, W., Roberts, P., Davis, E. 1996, in *Proceedings of the 4th APT Workshop*, ed E. King, 16