

LBADR: The LBA Data Recorder

Chris Phillips*

CSIRO ATNF

E-mail: Chris.Phillips@csiro.au

Tasso Tzioumis

CSIRO ATNF

E-mail: Tasso.Tzioumis@csiro.au

Steven Tingay

Curtin University of Technology

E-mail: s.tingay@ivec.org

Jamie Stevens

CSIRO ATNF

E-mail: Jamie.Stevens@csiro.au

Jim Lovell

University of Tasmania

E-mail: Jim.Lovell@utas.edu.au

Shaun Amy

CSIRO ATNF

E-mail: Shaun.Amy@csiro.au

Craig West

University of Massachusetts

E-mail: cwest@astro.umass.edu

Richard Dodson

University of Western Australia

E-mail: Richard.Dodson@uwa.edu.au

The LBA (Long-Baseline Array) is an ad-hoc network of radio telescopes within Australia and is the only VLBI array in the Southern Hemisphere. Since 2004 all experiments have been recorded using standard computer hard-disks, replacing the aging tape based S2 system. The recorder developed for this, the LBADR, comprises largely of commercial off-the-shelf components and is closely related to the PC-EVN system.

Science and Technology of Long Baseline Real-Time Interferometry:

The 8th International e-VLBI Workshop, EXPreS09

June 22 - 26 2009

Madrid, Spain

*Speaker.

1. Introduction

The Australian Long Baseline Array (LBA) is a "part time" VLBI array consisting of 3 telescopes operated by the CSIRO Australia Telescope National Facility (ATNF), 2 telescopes operated by the University of Tasmania plus telescopes from the NASA Deep Space Network facility at Tidbinbilla. Table 1 shows a summary of the telescopes in the LBA. Additionally, telescopes such as Hartebeesthoek, Kokee, TIGO & Shanghai have been included in LBA experiments when requested by proposers. In a typical year, around 25 days of LBA observing are scheduled, usually in 3-4 sessions. The LBA is an open access instrument with all time allocated in a competitive nature based on the science grading of a proposal. From 1994 until 2004 the LBA used the S2 VLBI system [1] which recorded on VHS magnetic tapes and the data were processed on a hardware correlator operated by the ATNF. Since 2004, all LBA operation switched to disk based recording, using the LBADR (LBA Data Recorder), and processed on the DiFX software correlator [2]. Currently all disk based correlation is run by Curtin University of Technology in Perth, Western Australia. The LBA also offers e-VLBI to telescopes with fast network connections, namely Parkes, ATCA, Mopra and Hobart.

When the LBA switched to disk based recording, the Mark5a disk recorder (which is used by the the majority of VLBI observatories in the rest of the world) was not an option as it was not compatible with the existing LBA digital backend.

2. The LBA VLBI system

The LBA VLBI system has three distinct parts: the VLBI front end electronics which digitize and filter the analog IF frequencies from the telescope, the VLBI recorder and the medium to which it is recorded.

2.1 LBA DAS

The LBA Data Acquisition System (DAS) is a sampler and digital filter and is used by all LBA stations. It was originally designed to work with the S2 recording system. A DAS receives two independent analog IFs (Intermediate Frequency, usually RCP and LCP) which are band limited to 64 MHz. The IF is digitized at 8 bits then run through a series of digital filters and resampled to 2 bits. A variety of filter types are available. The most common are either a single channel centered

Telescope	Diameter (m)	Frequency (GHz)	Maximum Data Rate (Mbps)
Parkes	64	1-26	1024
ATCA	6 x 22	1-100	1024
Mopra	22	1-115	1024
Hobart	26	1-26	512
Ceduna	30	2-26	512
Tidbinbilla	70 & 34	1-35	512

Table 1: The telescopes of the LBA

on the middle of the input 64 MHz, or a pair of upper and lower sideband channels (similar to the output from a single baseband converter in a standard VLBA and Mark4 acquisition rack). Filter widths of 0.5-32 MHz can be used (narrower filters are also available). A mode by-passing the digital filters and essentially using the DAS to generate the full 64 MHz band sampled at 2 bit is also available. In normal operation a single DAS is used to generate up to four 16 MHz channels sampled at 2 bits (giving a data rate of 256 Mbps) or two 64 MHz channels also sampled at 2 bits (512 Mbps). The three stations operated by the ATNF (Parkes, ATCA and Mopra) each have 2 LBA DAS units. This doubles the total bandwidth which can be recorded with these stations. The remainder of the LBA observatories each have a single DAS unit.

As the DAS was designed to be connected directly into the S2 recorder, it requires an interface box, the VSIC[4] designed by Metsähovi Radio Observatory (MRO). This external device changes the differential ECL (Emitter-coupled logic) signals to LVDS (low-voltage differential signaling) and changes the signal pinout to reflect the VSI-H standard (VLBI Standard Interface - Hardware)[5].

2.2 The LBADR

The LBADR is similar to the PC-EVN[3] designed by JIVE and MRO, although the specific realization of the recorder has significant differences. The recorder is a standard server grade rack-mountable PC, with the addition of a "VSIB" PCI card, designed by MRO[4], which makes it a VLBI recorder. The VSIB card is a simple 32bit/33 MHz PCI card with a VSI-H digital input. The digitized VLBI baseband data, along with timing signals, are received on up to 32 parallel data lines. The LBADR only ever uses up to 16 data lines clocked at 32 MHz giving a maximum data rate of 512 Mbps. The system has been tested, with synthesized data, up to 800 Mbps, close to the maximum theoretical data rate of the PCI bus.

The LBADR runs a standard Linux operating system with the addition of a simple kernel module which reads data from the VSIB PCI card into a large (512 MB) circular buffer in main memory. A user level C program then reads data from the circular buffer and writes the data to disk using a standard Linux filesystem or sends it over the network (for e-VLBI applications). This software is based on that written for the PC-EVN though it has been heavily modified to suit the requirements of the LBADR, including extensive data integrity checks and the addition of e-VLBI functionality.

The hardware chosen for the LBADR is a Tyan Thunder K8WE, a dual processor (AMD) server grade mother board, with dual gigabit Ethernet and a mix of independent PCI, PCI-X and PCI-Express add on slots. The motherboard was carefully chosen to avoid bandwidth "oversubscription" on the various buses and to be sure the system could sustain data rates of up to 512 Mbps on a single PC, recording to a variety of disks and the network. A generous number of PCI slots were required to support numerous add-on cards, such as IDE controller cards, fibre channel cards (see next section) and additional gigabit Ethernet ports.

2.3 Disk recording

The main disk recording medium for the LBADR is a Apple Xserve RAID system. This is an external fibre-channel attached hardware RAID solution. The device houses 14 IDE disks (up

to 750 GB each) in two sets of 7 disks. RAID5 is used to achieve some redundancy, at a cost of 14% of the disk capacity. The individual disks from the Xserve RAID can be removed and sent to the correlator which has Xserve RAID chassis to receive the disks. Because the Xserve RAID is a hardware RAID solution (and the disk enclosure means they cannot be physically inserted into a normal PC), the LBADR is also equipped with 4 removable IDE disks which are configured as a single software RAID0 volume. These disks can be either used as additional storage or allow data to be recorded which can be shipped to correlators which are not equipped with Xserve RAID hardware. An additional 2 "internal" SATA disks are permanently installed in each system. These contain the host operating system which occupies a small fraction of the disks. The rest is configured as an additional RAID0 volume which can be used for tests etc, though the data rate is limited to 256 Mbps. The design of the LBADR means it is simple to upgrade to a different commercially available disk technology. Any fibre channel attached storage could be used, or the fibre channel card could easily be replaced with a different interface (e.g. eSATA).

In all situations, data are recorded as files on a normal Linux file system. This makes data organization simple: multiple experiments recorded on a single disk are simply put into different directories, data can be selectively deleted and stored data can be inspected while new data are being recorded.

2.4 Compatibility

As the LBADR is essentially a normal PC with a digital input card it allows great flexibility. The LBADR normally uses its own data format (a 4 kB ASCII header followed by a few seconds of raw baseband data), which only DiFX can process. To allow compatibility with other correlators, an on-the-fly conversion to the Mark5B data format is possible. This essentially removes the LBADR headers from the data stream and inserts Mark5B headers, with appropriate time codes before writing to either disk or over the network. This data format has been tested on the Mark4 correlators at JIVE and Bonn. Support for the new VDIF data format has also been tested and will become the default data format when it is supported by DiFX and other correlators.

3. e-VLBI

The LBADR supports e-VLBI by simply copying the VLBI baseband data onto a network stream rather than directly to disk. Both the native LBADR format is supported along with real-time conversion to Mark5b format, using either the TCP or UDP transport protocols. For UDP the raw VLBI data stream is divided into chunks which fit within the UDP datagram size (MTU) with a leading 64 bit sequence number. This sequence number is used to reorder packets in time order and to account for missing packets (as UDP offers no guarantee of either delivery of the packets or in-order arrival). The LBADR hardware has four 1 Gbps network ports (2 on the motherboard and 2 on an add-on PCI card). This allows the recorder to be simultaneously connected to multiple networks allowing layer2 (Ethernet) trunking rather than relying on routers to distribute data between the recorders and correlator.

e-VLBI is offered to observers for routine observing. Data rates up to 1 Gbps are possible to the ATNF telescopes, but only 64-128 Mbps from Hobart due to the limited accessible bandwidth across the Bass Strait between Tasmania and mainland Australia. Linux clusters at the ATCA and




Figure 1: The ATNF network between Parkes, Mopra and ATCA. The links between the observatories and Sydney (ATNF headquarters) are 1 Gbps connections, presented as Ethernet. In Sydney the links connect to the AARNet3 backbone which includes 10 Gbps connections to Perth, WA and Seattle, USA.

Parkes are used to run DiFX for the realtime correlation of the data streams. Numerous e-VLBI experiments have also been run from Australian telescopes to the EVN Mark4 data processor at JIVE, the Netherlands. These used the UDP transport protocol and Mark5B data format. A mixture of routed IP networks and layer2 light paths have been used, with data rates up to 512 Mbps per telescope.

The ATNF internal networks are configured as shown in Figure 1. Basically each observatory has a 1 Gbps ethernet connection to the ATNF headquarters in Sydney and additional 1 Gbps circuits between Parkes & ATCA and Mopra & ATCA. These are private point-to-point connections only used by the ATNF but the bandwidth is shared between regular users (email, web, data transfer etc) and e-VLBI. As such e-VLBI data rates on a single link are never run at rates higher than 512 Mbps. To achieve 1 Gbps (1024 Mbps), the diverse nature of the point-to-point networks and multiple VLANs are used. Three separate e-VLBI VLANs exist, separate from the normal production network connections. These VLANs trunk the connections over different physical paths (e.g. direct Parkes-ATCA or Parkes-ATCA via Sydney). By careful choice of destination network address traffic from 1 Gbps e-VLBI from telescopes can be achieved by either splitting the data and sending 512 Mbps from each station to both the cluster at Parkes and the cluster at the ATCA, or all data can be sent to the ATCA at 1 Gbps.

3.1 Remote Recording

The LBADR supports remote disk recording. This is essentially a form of e-VLBI (the LBADR is just configured in e-VLBI mode) but rather than processing the data in real time, the data are received on a different computer and written to disk. As the Mopra telescope is operated remotely, local disk recording is problematic. All Mopra observing is run in the remote recording

form with the data sent over the ATNF networks and recorded at the ATCA, 115km away. Remote recording is also useful as it allows data to be recorded onto any PC at the observatory if disk media is limited (constrained by disk speed and capacity on the receiving PC).

3.2 Data Transport

Since the beginning of 2009, data from all LBA stations (except Ceduna) to be correlated at Curtin have been transported electronically over the available networks. This eliminates the cost of transportation as well as the risk of damage or loss. The data are copied to a "petabyte" store - a large hybrid disk/tape array hosted and provided by iVEC¹. The data are stored on this system until they are needed for correlation when they are copied onto the correlator cluster. The process of copying the data is handled by the Australian Research Collaboration Service, as part of a government funded service to scientists within Australia.

3.3 Control

The core of the LBADR software is a Linux kernel module and a command line C program. To allow for simpler use by operators, remote control etc, a persistent "daemon" is run on each recorder. This daemon receives XML like commands via a TCP socket. These control all aspects of the recorder configuration such as bandwidth, channels to record/discard, mode of operation (data format, e-VLBI or disk recording etc) and starting and stopping recording. The daemon can also be run on any PC which is available to receive data for remote recording. As such remote recording to any PC can be setup in a transparent fashion.

The standard control for the LBADR is a GUI written in Perl/Tk. This communicates directly with the control daemon and allows the operator to configure the system for the experiment and start the recorder and monitor the recording state (including sampler statistics). For LBA e-VLBI, the DiFX startup automatically starts the recorders. A simple "Mark5" control emulator has also been written which receives Mark5 commands and translates to appropriate commands for the LBADR.

4. Realtime Fringe Checking

A "realtime" fringe checker has been incorporated into the operation of the LBADR. This is an interactive suite of programs which on demand extracts a small portion of data from the recorder and copies them to a central location, where the data are correlated in software. Arbitrarily small time intervals can be selected - initially Hobart and Ceduna were only connected by dialup modems, so intervals as short as a few tens of milliseconds were used. As well as a short time interval, it is possible to just select a sub-selection of the data channels (to reduce the size of the data to be transferred). Because the LBADR allows data to be read and written simultaneously this system does not require fringe checks to be pre-scheduled (although they can be if desired). At any time during an experiment a fringe test can be run. Once the data are copied (which takes between a few seconds to few minutes depending on the speed of the various telescope connections and integration time used) the correlator is automatically run and results displayed on a webpage for

¹<http://www.ivec.org>

all observers to view. All experiments have a dedicated "fringe check" experiment before the main experiment, and this has significantly improved the reliability of the LBA.

5. Future Developments

Support for the VDIF data format is being developed and will become the "standard" recording mode, aiding in international compatibility.

The next major change will be support for "hybrid" recording. This will allow e-VLBI and disk recording simultaneously. The approach taken is to create multiple threads - one to read data off the VSIB card and multiple threads to write data to disk or over the network. An arbitrary number of threads is supported and each has the option of discarding a portion of the data. This allows for a lot of flexibility in recording, such as simultaneous e-VLBI and disk recording, simultaneous remote recording and e-VLBI or e-VLBI with distributed correlation (sending some of the data to one correlator and the rest to another). This development has been shown to work but requires further testing before it is used for all experiments. The expectation is that all disk based experiments will include e-VLBI for the telescopes which support it, potentially at a lower data rate than is recorded to disk. This will allow the astronomer to get a preview of the data as soon as the experiment has finished (or even before) and will allow realtime system verification.

6. Conclusion

The LBADR has been a great success for the LBA. It has been very reliable and is a very flexible system. The flexibility is largely because the bulk of the system is based on standard PC components with simple C software to join it all together. This makes it simple to add new functionality such as different data reliably and robustly.

References

- [1] Cannon, W. H. and Baer, D. and Feil, G. and Feir, B. and Newby, P. and Novikov, A. and Dewdney, P. and Carlson, B. and Petrachenko, W. T. and Popelar, J. and Mathieu, P. and Wietfeldt, R. D. 1997, *Vistas in Astronomy*, 41, 297
- [2] Deller, A. T., Tingay, S. J., Bailes, M., & West, C. 2007, *PASP*, 119, 318
- [3] Parsley, S., Pogrebenko, S., Mujunen, A., & Ritakari, J. 2003, *Astronomical Society of the Pacific Conference Series*, 306, 145
- [4] Ritakari, J., & Mujunen, A. 2004, *International VLBI Service for Geodesy and Astrometry 2004 General Meeting Proceedings*. Ottawa, Canada, 2004. Edited by Nancy R. Vandenberg and Karen D. Baver, 182
- [5] Whitney, A R. *VLBI Standard Hardware Interface Specification*, Revision 1.0, 2000.