

# eLBA Network Requirements

## ATNF eVLBI Memo #4

### Draft

### Version 0.6

Chris Phillips

February 9, 2006

## 1 Introduction

The ATNF antennas (Parkes, ATCA and Mopra) are currently connected to ATNF headquarters in Sydney (Marsfield) with 1 Gbps network links. The main purpose is to transfer Very Long Baseline Interferometry (VLBI) “baseband” data for processing on software correlators. These currently are at Swinburne University of Technology in Melbourne and the University of Western Australia. Over the next 12-24 months it is hoped that telescopes at Tidbinbilla near Canberra and the University of Tasmania’s telescopes in Hobart will also be connected to broadband network links, and eventually the Ceduna antenna in South Australia.

The current goal is to be able to transfer data from the ATNF observatories (and maybe Hobart) to the data processors for realtime processing. As all the observatories will probably not have access to broadband links we will also need to support “remote disk” recording of data as well as a “store and forward” approach i.e. recording to local disks, moving them to the nearest broadband link and then copying over the network.

The longer term (4 year) plan is to be able to connect the observatories to Narrabri and to be able to use the Compact Array Broadband Backend (CABB) for data processing. This would require dedicated network connections of up to 40 Gbps per observatory.

This document discusses the network requirements for the next 3 years. Given the rapidly changing landscape of broadband connectivity, available data processing centres etc, the requirements may change significantly in the next year or so.

## 2 The LBA

The LBA consists of 6 antennas spread across Australia. The ATNF operates three antenna in NSW, the University of Tasmania operates an antenna in Tasmania and one in South Australia and there is some access to antennas at NASA’s Canberra Deep Space Communications Complex (CDSCC) at Tidbinbilla. See Table 1 for a summary of the LBA antennas.

| <b>Antenna</b> | <b>Size</b> | <b>Institution</b>           | <b>Closest population centre</b> |
|----------------|-------------|------------------------------|----------------------------------|
| Parkes         | 64 m        | ATNF                         | Parkes, NSW                      |
| ATCA           | 6×22 m      | ATNF                         | Narrabri, NSW                    |
| Mopra          | 22 m        | ATNF                         | Coonabarabran, NSW               |
| Tidbinbilla    | 70 m        | NASA                         | Canberra, ACT                    |
| Tidbinbilla    | 34 m        | NASA                         | Canberra, ACT                    |
| Hobart         | 26 m        | Tas Uni                      | Hobart, Tas                      |
| Ceduna         | 30 m        | Tas Uni                      | Ceduna, SA                       |
| Hartebeesthoek | 26 m        | National Research Foundation | Johannesburg, South Africa       |

Table 1: LBA antennas and common international collaborators.

Most of the time the antennas are operated individually and the data processed locally. About 20 days a year, VLBI time is scheduled simultaneously on all telescopes. During this time baseband data would be recorded and transferred over the network. This means that for most of the time network links required for this project would not be needed for VLBI. While a “store and slowly forward” approach is usable (say requiring a few days to transfer one day’s worth of data) it is not desirable as the only option.

The LBA also regularly observes with international observatories. Requirements on international links is beyond the scope of this document.

### 3 VLBI data processing

For VLBI, telescopes separated by hundreds or thousands of kilometres are operated simultaneously and the data collected from each telescope are combined in a supercomputer (using either custom built or more recently using software on general purpose computers). The data processing is used to synthesise a telescope the size of the longest separation between the individual antennas.

The basic data flow first involves digitising the analogue signal collected by the antenna. Usually only 2 bit quantisation is used as this give the most information per bit. The digitised data are then captured using a custom built digital input card which plugs into a PCI bus of a “generic” PC. The data then need to be transported to a central location to be processed. Data from all telescopes must be processed at a single location at the same time. This means that if, for example, one telescope does not have access to high speed network links and is recording the digitised data to disk then the data from all telescopes has to be saved and processed as a whole once the disks (or their contents) can be transported to the data processor.

In normal operation, the processor averages the data with an integration of a few seconds. Thus the output data rate is many orders of magnitude lower than the input data rate (few 10’s MB/sec). However for some special modes of operation the data must be divided into many (thousands) of frequency slices with very short integration times (10’s of millisecc), so the processed data rate can rival the unprocessed data rate.

## 4 Data Rates

The current VLBI recording back-ends installed at the ATNF observatories can produce 1 Gbps of user data ( $1024 \times 10^6$  bits/sec). Tidbinbilla, Hobart and Ceduna are currently limited to 512 Mbps, due to a limited number of available baseband digitisers. Data rates for observations will normally be 128, 256, 512 or 1024 Mbps. 768 Mbps is also feasible but unlikely.

## 5 Data Processor

As data from all telescopes must be combined simultaneously, the data must be sent to a central location. Because of this requirement VLBI does not lend itself naturally to distributed processing. However the data can be sliced in either time or frequency and sent to multiple processing centres if all data for a specific time or frequency are sent to the same place. In Australia there are two main processing centres which are currently being considered for eVLBI processing: Swinburne and the University of Western Australia.

### 5.1 Swinburne

Swinburne University of Technology is based in Hawthorn, Melbourne, Victoria. The University Information Technology Services (ITS) group operates a  $>300$  CPU cluster which runs a software correlator written by the Centre for Astrophysics and Supercomputing. This cluster has enough processing power for realtime data processing at the maximum data rate (1024/512 Mbps).

### 5.2 WASP, University of WA

The University of Western Australia Faculty of Engineering, Computing and Mathematics oversees the Western Australian Supercomputer Program (WASP). As part of this facility, a Cray XD1 supercomputer is dedicated for VLBI data processing. This computer consists of 12 64-bit 2.6 GHz AMD Opteron processors to deliver 58 GFLOPS. Cray have donated Six Xilinx Virtex-II Pro Field Programmable Gate Arrays (FPGA) and FPGA based VLBI processor routines to provide VLBI an accelerated parallel execution of the Swinburne software correlator. This machine has an aggregate memory bandwidth of 77 GB/s.

## 6 Modes of Operation

The long term goal for most eVLBI operations is to transport the baseband data in realtime from the observatories to the data processor without (or with minimal) buffering to disk. However until all observatories are connected to wideband network links, recording all data to disk will be required.

### 6.1 Store and Forward

Data would be recorded (at rates of up to 1 Gbps) on disks local to the antenna and later copied over the network to disk storage at the data processing centre. In this case the network bandwidth can be

lower than the recording data rate (although if the bandwidth is, say, an order of magnitude slower the transfer process will take too long to be viable).

For observatories such as Tidbinbilla and Ceduna, which currently have no wide bandwidth link, the disks would be brought to the nearest wideband network connection for transfer after the experiment.

For logistical reasons, data destined for UWA may be shipped on disk to Swinburne and the data then copied across the network.

## **6.2 Remote Recording**

When observing with a mix of observatories with and without wideband links, the data from the connected antennas would be sent over the network in realtime to remote data storage facilities. This may be data storage connected to the data processor but this is not necessary if there are fast links between the data storage facility and the data processor.

For this mode of operation to be viable, the available bandwidth on the network needs to exceed VLBI data rate. Initially it is envisaged that rates of 512 Mbps will be supported in this mode.

## **6.3 Realtime correlation**

For experiments only involving telescope connected to broadband links, data would be streamed in realtime directly to the data processor. Minimal or no disk buffering would be involved. Network availability would be similar to the remote recording mode, in that a guaranteed minimum data rate (preferably 512 Mbps) would be required.

## **6.4 Hybrid Transfer**

Currently the fastest links to the telescopes is 1 Gbps. This will not be fast enough to send 1024 Mbps of user data, due to protocol overheads etc. These links are also shared by all users at the observatories for day-to-day data transfer, email, web etc.

When running the recorders at 1 Gbps, the most probably mode of operation will be in a hybrid mode. A 512 Mbps data stream would be sent to the correlator or data storage facility in real time. The other 512 Mbps stream will be stored locally and either transferred after the experiment or transferred during the experiment at a data rate of a few hundred Mbps.

# **7 Data Storage Requirements**

Typical VLBI experiments last 12 hours to get the best image quality. Table 2 shows the disk requirements for a range of different data rates. The current scheduling approach for the LBA with approximately 3 one week sessions would consist of about ten 12hr experiments. Given access to large amount of disk storage we would expect about 3 experiments to run at maximum speed and the rest at 256 Mbps. Thus the total storage requirements is 126 TB (assuming Tidbinbilla only takes part in two experiments both at maximum speed and Ceduna takes part in 9 experiments).

| <b>Data Rate</b>                     | <b>Required Storage/telescope<sup>2</sup></b> | <b>Total storage<sup>3</sup></b> |
|--------------------------------------|---|----------------------------------|
| Maximum (1024/512 Mbps) <sup>1</sup> | 2.76/5.53 TB                                  | 24.88 TB                         |
| 512 Mbps                             | 2.76 TB                                       | 16.6 TB                          |
| 256 Mbps                             | 1.38 TB                                       | 8.3 TB                           |
| 128 Mbps                             | 0.69 TB                                       | 4.1 TB                           |

Table 2: Disk storage required for a 12 hour experiment at various data rates.

<sup>1</sup> 1024 Mbps at Parkes, Mopra and ATCA. 512 Mbps at Ceduna, Hobart and Tidbinbilla

<sup>2</sup> Assuming 1 TB is  $10^{12}$  bytes (as used by disk manufacturers)

<sup>3</sup> Assuming all six LBA antenna taking part. Tidbinbilla is only available for a few experiments each session and Ceduna can not operate all frequencies.

## 8 Telescope Connectivity

### 8.1 ATNF antennas: Parkes, Mopra, ATCA

The ATNF antennas are due to have 1 Gbps network connections in February 2006. This will be part of the AARNET regional network. These connections are private point-to-point connections to the ATNF headquarters in Marsfield, Sydney. Connection to the Internet is via the normal CSIRO Gbps connection.

### 8.2 Hobart

Hobart has received an ARC “LIEF” grant to connect the observatory to a 10 Gbps link. The contract for this link is in the final stages of negotiation. Once the Basslink pipeline is constructed it is hoped that the University will have access to a wideband connection to Melbourne.

### 8.3 Ceduna

There are currently no plans to connect Ceduna to broadband network links. Short term data transport is likely to involve either shipping data disks to the data processor or taking them to the University antenna in Hobart and using the broadband network connection to copy the data.

### 8.4 Tidbinbilla

There are currently no plans for a broadband network to Tidbinbilla. It is hoped that a source of funding can be found and Tidbinbilla connected to the AARNET regional network in a similar manner to the ATNF telescopes. Short term data transport is likely to involve either shipping data disks to the data processor or taking the disks to either Mt Stromlo or ANU and using the broadband network connection to copy the data.

## 9 CABB (Compact Array Broadband Backend)

The next generation data processor for the ATCA is currently being built and due to come on line in 2007. This data processor has the capability to process data from 8 telescopes each with data rate of 64 Gbps (ie 512 Gbps total). By 2009 we hope to use the CABB for realtime VLBI data processing. For VLBI usage the input data rate can be reduced from the maximum. The current plan would be to connect the ATCA to 20 Gbps connections from each LBA observatory. This would require up to 140 Gbps data rates coming into the ATCA from across Australia. Probably the input data rates would initially be significantly less than this and increase up to the maximum over time.

## 10 Time lines

| <b>Date</b> | <b>Network Requirement</b>   |
|-------------|--|
| 2006        | 512+ Mbps from Parkes, Mopra & ATCA to UWA (ie 3×512 Mbps to UWA)            |
| mid 2006    | 512+ Mbps from Parkes, Mopra & ATCA to Swinburne                             |
| late 2006   | 100-512 Mbps from Canberra (ANU or Mt Stromlo) to UWA and Swinburne          |
| 2007        | Additional 512+ Mbps connection to UWA and Swinburne from Hobart             |
| end 2007?   | 512+ Mbps connection from Tidbinbilla to UWA and Swinburne                   |
| 2008        | 1 Gbps from Parkes, Mopra, ATCA, Hobart and Tidbinbilla to UWA and Swinburne |
| 2009        | 1 Gbps connection from Ceduna to UWA and Swinburne                           |
| 2009/2010   | 10+ Gbps connection from all Observatories to Narrabri (CABB)                |