



LBADR: THE LBA DATA RECORDER

Chris Phillips
ATNF, CSIRO, Australia



Summary

The LBADR (LBA Data Recorder) is a VSI-H data recorder based on the Metsähovi Radio Observatory (MRO) "VSIB" digital input board. Other than the input card, the entire system is based on COTS PC components running Linux. Data is recorded at rates up to 512 Mbps on normal Linux filesystems. The system allows any type of commercial disk controllers to be used with minimum effort. Currently Apple Xserve RAID arrays (which use an external chassis connected via fibrechannel) are being used for the majority of experiments, while removable IDE disks running in a software raid configuration are used when data needs to be sent to institutes which do not have Xserve RAID systems. The PC architecture is a dual processor (AMD64) system with onboard dual gigabit Ethernet. This allows streaming of data at up to 512 Mbps for either remote recording or realtime correlation on the Swinburne software correlator.

Data Acquisition

Data acquisition depends on the ATNF DAS, which uses digital filtering to produce up to four 16 MHz bands (on 2 independent IF chains, usually Rcp and Lcp) with a data rate of 256 Mbps (2 bit sampling). If digital filtering is disabled, two 64 MHz bands can be produced, also with 2 bit sampling, with a data rate of 512 Mbps. It is also possible to produce two bands up to 16 MHz bandwidth with 10 bit sampling. All modes are supported by the LBADR recorder.

The ATNF DAS is not VSI-H compliant as it was designed to be used with the Canadian S2 recorders. To convert the DAS signals into VSI-H, a "VSIC", also produced by MRO is used. The VSI-H output from the VSIC plugs into the PCI based VSIB card.

The ATNF telescopes Parkes, Mopra and ATCA are each equipped with 2 DAS. These are have been routinely used together with a total recording rate of 512 Mbps and successful fringe test have been performed at 1 Gbps. The rest of the LBA observatories have a single DAS, so have a maximum recording data rate of 512 Mbps.

PC Architecture

Recorders based on the MRO VSIB card have minimal system requirements, any modern motherboard should be adequate, as long as the disk controllers are not electronically attached to the same 32bit/33 MHz PCI bus the VSIB card is plugged into. For the LBADR, the requirements needed to be met were: 1) record onto a mass storage device, 2) support 4 removable IDE drives running as a software RAID 0 device and 3) support network streaming of the data at 512 Mbps.

Apples Xserve RAID systems were chosen as the mass storage devices. These come in an external unit with 14 removable IDE disk drives (currently each are 500 GB) configured as RAID 5 for data protection. The Apple Serve RAID connects to the PC via a plug in PCI-X based fibre channel card.

Because the VSIB effectively consumes all PCI bandwidth of a standard 32 bit/33 MHz PCI bus, the PC motherboard needs to support multiple PCI bus. A Tyan "Thunder K8WE" motherboard was chosen as it supports two PCI-X 100 MHz slots, one PCI-X 133 MHz slot, two PCI-express slots and one standard 32bit/33 MHz PCI slot, each with dedicated bandwidth. This motherboard also comes with two on board gigabit Ethernet controllers which are electronically separate from the PCI expansion slots. As well as the PCI Fibre channel card we also use two Promise Technology IDE controller cards to control 4 standard IDE disks which are run as single software RAID 0 volume.



An Apple Xserve RAID chassis. This connects to the LBADR using fibrechannel.

Software Architecture


Data acquisition depends on a custom Linux "character device driver" (supplied by MRO) which can easily be compiled against any modern Linux kernel with a "bigphysarea" patch. This driver sets up DMA on the PCI VSIB card to write the received data into a large ring-buffer (typically 512 MB) in the PC ram. A simple user level C program then copies the data out of the ring buffer and either writes the data to disk, using a standard Linux filesystem, or across the network using TCP for remote recording or realtime processing on the Swinburne software correlator. The software has a flexible design, so enabling alternative transport protocols is quite straight forward.

A number of utility programs are installed on the recorder to allow data inspection. These include checking sampler statistics and computing the auto-correlation spectrum. The spectrum of a single 16 MHz channel can be monitored in realtime. In the future it is hoped optimised FFT libraries will allow more channels to be monitored simultaneously. These tests can be run while data is being recorded.

Realtime Fringe Checking

Because the LBADR uses a normal Linux filesystem, it is possible to copy a small fraction of the recorded data, while recording continues, to be processed on the Swinburne software correlator.

The "Real time Fringe Test" software (RtFC) is a distributed suite of Perl programs which are controlled interactively from a central location. The user requests correlation of a specified time and this data is copied to Swinburne and the software correlator automatically run, with the results posted on the web. The user can request an integration time of a fraction of a second up to 10 seconds. Either all recorded channels can be copied and processed or a flexible sub-selection of channels used (useful for observatories on slow network links). On the fast links to the ATNF observatories, the "photon to fringe" time is < 20 seconds. The system also works well on slow links, including dial-up modem. The process can be scripted using very simple command files but this is generally not used as the interactive nature is found to be more useful. Almost all LBA experiments (both tape and disk based) are now scheduled with a fringe test before the user experiment.



Sample fringe plot (Parkes-Mopra) from the real-time fringe finder. These plots are automatically generated and appear on a experiment specific webpage.

eVLBI

The LBADR can easily handle continuous streaming of data. In network test to University of Western Australia (Perth) the system could sustain 940 Mbps TCP data rates for many hours (memory-memory) using a 1500 byte MTU. Disk-disk transfers to Perth can be sustained at 840 Mbps.

Currently eVLBI (non-buffered) has been run up to 256 Mbps. This was limited by available computing power connected to the broadband links. (A spare/development LBADR recorder was actually used to run the Swinburne software correlator for this experiment, showing the flexibility of this system).