Proceedings of the 6th European VLBI Network Symposium
Ros, E., Porcas, R.W., Lobanov, A.P., & Zensus, J.A. (eds.)
June 25th-28th 2002, Bonn, Germany

Status of the EVN MkIV Data Processor at JIVE

H.J. van Langevelde*,**

Joint Institute for VLBI in Europe, Postbus 2, NL-7990 AA, Dwingeloo, the Netherlands

Abstract. Nowadays most EVN experiments are being processed at JIVE. One of the main considerations is the release of tapes in a timely fashion back into the EVN. Processing of continuum experiments is now straightforward, and we routinely process cross-polarisation and phase-referencing experiments. Monitoring the playback performance and checking the data for a number of known problems is still a labour-intensive process. Dealing with oversampled recordings has become operational in 2001. Recently significant enhancements have been made to the read-out capacity of the correlator.

1. Introduction

The EVN MkIV data processor at JIVE has been operational for three years. Most of the EVN experiments and a considerable fraction of the Global VLBI experiments are being processed in Dwingeloo. Still, there are a number of areas where development is continuing. It is important to keep a balance between maintaining a reasonable throughput, fixing bugs and developing new capabilities. The main priority is always the quality of the data handed to the user. This requires detailed scrutiny of the correlator output. At the same time the correlator has to keep up with the incoming flux of recordings and make tapes available for new sessions. By carefully prioritising efforts, progress is made on introducing new capabilities without jeopardising the established data quality or throughput. Recently a number of improvements have been made that are important for spectral line or wide field of view processing.

2. Technical Background

The EVN MkIV 16 station data processor was built within an international consortium. It is located at the Joint Institute for VLBI in Europe, in Dwingeloo, and is hosted by ASTRON (Fig. 1). The correlator was designed to be able to deal with 1 Gb/s data rates produced by the MkIV data acquisition system. It can also accept VLBA and MkIII data. The data playback units are in principle capable of reading 64 tracks per tape at 16 Mb/s/track (see Schilizzi et al., 2002, for a description of the data processor).

There are 16 Play Back Units of P&G make at JIVE. To assure read-back quality they are monitored in a monthly cycle by playing back a standard recording. In 2001 all units were upgraded with special hardened stainless steel parts. In the process the complete tape path, including the reel motor positions, had to be re-adjusted and re-calibrated. This upgrade led to noticeable improvements in the playback quality.

The processing power of the correlator is delivered by 1024 custom chips, organised in an XF architecture (Whitney 2000). The processor can deliver a total of 262144 complex lags. The data are read out on standard Ethernet hardware and stored on Unix workstations in a raw, internal format. For quality control and also for testing, the data are converted into aips++ data structures. For astronomical processing the data can then be exported in FITS format.

Although the correlator at JIVE has much of its hardware in common with the other (geodetic) MkIV correlators, its software was developed completely independently. A rigorous version control system ensures we can switch easily between the operational version of the software and development and test versions. Typically two to four times a year a new production version of the software is released. At that point detailed comparisons are made between data from the new and old versions.

The processing of a specific experiment runs from a MkIV VEX file. This file is read in at the correlation stage and drives the entire data processor. Most of this information is directly obtained from the observing schedule, also in VEX format, but additional information about tape labels, earth orientation and clocks is supplied from other sources. The station GPS files and logs are also used to generate the correlation VEX file.

3. Operations

The correlator produced its first fringes on July 21st, 1997, and started production correlation exactly two years later. The first paper produced from "JIVE data" was a spectral line experiment on HI absorption (Van Langevelde et al. 2000). By the start of 2001, the EVN MkIV data processor at JIVE was processing almost all EVN projects; previously some (spectral line) projects had to be correlated at Socorro.

In the first three years 83 user projects were processed. In addition, every year approximately 20 network monitor-

* On behalf of the correlator group.

** The author acknowledges partial support from the EC ICN RadioNET (Contract No. HPRI-CT-1999-40003).

Fig. 1. The EVN MkIV Data Processor at JIVE

ing experiments and tests are processed. The processor is manned for 80 hrs/week, and currently one or two projects per week get finished, despite a fair fraction of time being spent on testing new capabilities and correlation in support of the Network. In the future the throughput will go up; speed-up correlation, in which the tapes play faster than they were recorded, was tested successfully in June 2002.

The first Global VLBI experiments came to JIVE in 2001, after methods were implemented to deal with experiments involving more than 16 telescopes observing simultaneously. The first (and probably last) MkIII experiment was successfully completed as well.

In November 2001 the EVN switched to the exclusive use of thin tapes, the main reason being to improve recording and playback quality, and indeed the use of a single tape type seems to have had a positive effect. However, as the thin tapes are expensive (and rare), there is great concern to recycle the tapes timely. In 2001 the EVN, through JIVE, ordered an additional 150 thin tapes. These were re-reeled on a dedicated playback unit and added to the tape pool. From 2002 on we are keeping up with the tape requirements: for the first 2002 session approximately 180 tapes (thin only) were shipped out. We anticipate that future sessions could require up to 250 tapes.

3.1. Data Quality Issues

The data from the correlator is subject to quite detailed quality control. There are a number of known issues that can affect the data quality and sometimes require re-correlation. In addition, it is inevitable that yet unknown bugs will be uncovered, as new areas in parameter space are explored.

The most serious issue in this respect is always playback quality. Although station recorders and JIVE playback units are regularly monitored, it does happen that certain tapes do not play back properly on some drives. If such cases are encountered we move tapes around until reasonable playback is achieved. Sometimes such cases are only uncovered after correlation, and we do re-correlate such experiments to ensure that the data quality is satisfactory.

Next to playback issues, byte slips are the main operational problem. Imperfections in the decoders occasionally cause the correlation function to be shifted by eight lags. We have been able to reduce the occurrence of these effects by moving suspect hardware out of the production correlator. It was discovered that the occurrence of byte slips could be quickly inferred from inspecting the autocorrelation function, as the sampler statistics change when a byte-slip occurs. When byte-slips are detected, re-correlation of the affected pass is performed. A fix of the embedded data handling logic is being tested.

It has been noted that in some cases the autocorrelation functions come out normalised different from unity, although the functions have the right shape. It is clear that this originates from the sampler statistics of the 2-bit data, as produced at the stations. The current weighting

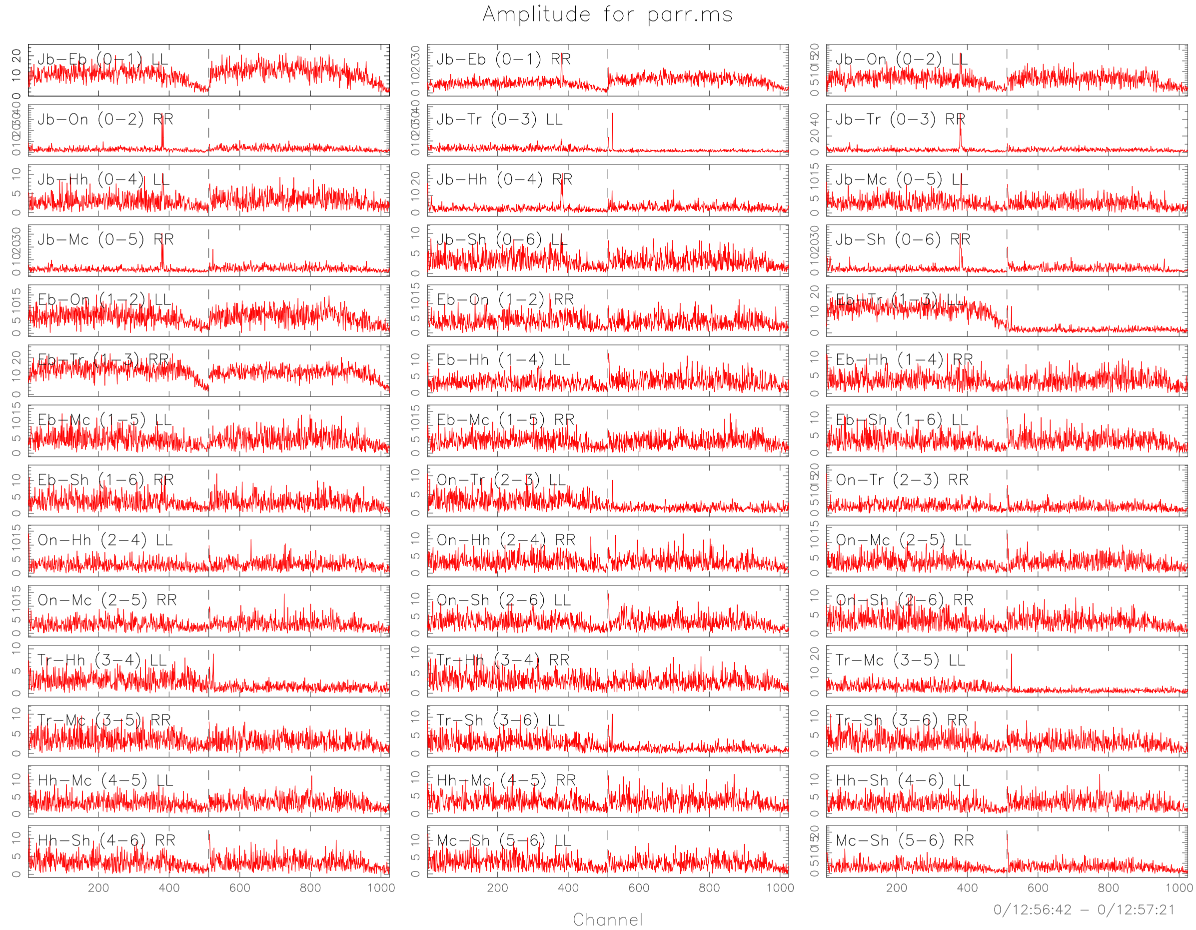

Fig. 2. First results of the improved correlator readout: a 7 station Methanol experiment with 2×2 products, 512 spectral points per product, and 1s integration time.

scheme is very sensitive to non-optimal sampler settings. This will be made more robust in the future.

A fraction of the user projects we receive use previously untested combinations of parameters. It is therefore not surprising that small bug fixes are sometimes required before production correlation. In 2001 we repaired 1-bit correlation, the processing of 16 MHz bands and a variety of mode switching problems.

Currently up to six weeks are required to check out an experiment. Several projects have been initiated to streamline the quality-control process. Detecting the problem cases at an earlier stage and automating the data generation will increase throughput.

3.2. Procedures

Before correlation of user products begins, the PI's are contacted about the desired correlator parameters. At this point the astronomer can set the spectral resolution, integration time and source coordinates.

In the process of setting up the correlation, the latest telescope coordinates are inserted. New positions for EVN stations were reported by Charlot et al., (2001, Charlot, Campbell, these proceedings), and these have been shown to improve the quality of phase reference experiments considerably.

In November 2001 we started adopting a policy of automatic experiment release. Tapes can be erased after a 4 week grace period, during which the astronomer has the opportunity to inspect the quality of the correlation through plots produced by the correlator staff or by analysing the data.

3.3. User Product

The data is shipped in FITS format, predominantly on DAT tapes. The information content of the data such that the processing resembles closely the VLBA path through AIPS. However, there are a small number of areas where EVN/JIVE specific (or rather VLBA a-specific) changes to AIPS are needed; notably the Doppler tracking for MkIV correlators is different from the VLBA or MkIII implementation.

As soon as the data are released to the PI, the experiments are also processed in the JIVE pipe-line. The data are calibrated and fringe-fitted. For sources that the PI has indicated as calibrators, preliminary maps are produced. The results of the pipe-line can be inspected through a web interface and the calibration tables can be down-loaded.

		Ntels	bands/ pols	freq points	min. t_int [s]
Current	EVN line exp	8	2	1024	1
	Global line exp	16	2	256	1
	Wide field	16	4	32	0.25
Future	EVN line exp	8	4	4096	0.125
	Global line exp	16	2	2048	0.125
	Wide field	16	8	256	0.06125

Table 1. Table of correlator modes

A project has started to make an archive of user data available through the Internet.

4. Capabilities

4.1. Modes and Formats

The EVN MkIV data processor was designed to deal with MkIII/MkIV/VLBA formats. As the implementation of barrel-rolling has been given low priority, this feature is switched off for experiments to be processed in Dwingeloo. The processing of two head recordings has been shown to work, and all Network Monitoring Experiments use 512 Mb/s modes. These modes are available to users as well. A decision was made to process all 2 head recordings in two passes, rather than to mount two heads on all playback drives. As the second head will also wear when playing back single head recordings, this is the most economic approach.

The processing of oversampled data, necessary for spectral line VLBI, in order to get to bandwidths of e.g. 500 kHz proved to be relatively tricky. After many tests and some changes to the processor software, these observing modes were made operational at the end of 2001. This has caused delay to a number of spectral line projects.

4.2. Current Capacity

There are two related parameters that determine the correlator capacity. The first is the total number of lags the correlator has available. Together with the number of telescopes and the number of sub-bands, this determines the maximum number of spectral points in each cross/auto spectrum. The second constraint is the read-out bandwidth of the correlator, which determines the minimum integration time that can be used for a specific product of baselines and spectral points.

In 2001 the so-called DSP project was completed. By adding processing power to the correlator boards, they could then be used simultaneously, making all 262144 complex lags available. In concrete terms, this implies that for an eight (or fewer) station experiment 2048 spectral points are available on each baseline (Table 1). So when observing, for example, 2 OH transitions in dual circular (no cross-polarisation) 512 spectral points are available for each of the cross-spectra. For more stations (9 – 16) the number is reduced by a factor of four, allowing 512 spectral points per baseline. In some exceptional cases more resolution can be obtained by processing sub-bands in separate passes, but this should be requested in the original proposal. The number of available frequency points will improve in the future by introduction of recirculation (below).

For the above cases, where the entire correlator is producing data, the read-out speed is such that the minimum integration time is one second. A major milestone was passed May 2002, when three related enhancement projects were finished (Fig. 2). Before that time, the minimum dump time was eight seconds. The improvement is the result of an increase in read-out capacity from 320 kB/s to 1.5 MB/s, and a reduction of the correlator block size. Further advancements will come from the PCInt project (below). Integration times below 1 second are available for modes that do not use the full correlator.

4.3. Future Enhancements

The number of correlator lags, and therefore the number of spectral points, will be increased when recirculation is implemented, possibly by the end of 2002. For bandwidths below 16 MHz the processing power of the chips can be time-shared between up to 8 products, increasing the capacity accordingly.

This will also require an enhancement of the read-out capacity, addressed by the PCInt project. It involves the introduction of single-board computers into the correlator. Eventually the data-rate will be increased to 160 MB/s. In addition to spectral line projects, the PCInt will also be very useful for wide field imaging (Table 1). A project to correlate multiple field centers simultaneously is therefore of reduced priority.

Initial tests were done of the pulsar gating hardware. Although the absolute timing is not available yet, consistent results were obtained that show that the gating hardware can be used to enhance signal to noise on pulsed radiation. Correlations were performed to prove the concept of multiple field centre processing. There are currently no detailed plans to enable Space VLBI or the production of geodetic data, although these capabilities were designed in.

References

Charlot, P., Campbell, et al., Proceedings of the $15^{\rm th}$ Working Meeting on European VLBI for Geodesy and Astrometry, ed. D. Behrend, & A. Rius, 194, `http://www.ieec.fcr.es/hosted/15wmevga/proceedings/charlot`

Schilizzi et al., 2002, Experimental Astronomy *in press*

van Langevelde H.J., Pihlström Y.M., Conway J.E., Jaffe W., Schilizzi R.T., 2000, A&A, 354 L45

Whitney, A.L., in Proceedings of the First IVS General Meeting ed. N.R. Vandenberg & K.D. Baver, NASA/CP-2000-209893, 187