Stephen Parsley: eVLBI Research in the European VLBI Network, IVS 2004 General Meeting Proceedings, p.205–209
http://ivscc.gsfc.nasa.gov/publications/gm2004/parsley

# eVLBI Research in the European VLBI Network

*Stephen Parsley*

*Joint Institute for VLBI in Europe*
*e-mail:* parsley@jive.nl

**Abstract**

Multi-gigabit connections are now available in most European countries via National Research and Education Networks (NRENs) and across Europe via GÉANT. The EVN Data Processor at JIVE, and a growing number of telescopes, now have high bandwidth connections to these networks. With the support of DANTÉ and a consortium of NRENs, the European VLBI Network (EVN) is currently involved in a Proof-of-Concept project to connect up to six European telescopes directly to the EVN data processor. The project will explore the feasibility of achieving real-time eVLBI using shared IP across the production, research and education (R&E) networks.

## 1. Introduction

In January 2003 a meeting was convened in Amsterdam to discuss the technical details of a Proof-of-Concept project to connect up to six European telescopes directly to the EVN data processor at JIVE. The meeting was set up and chaired by DANTÉ, operators of GÉANT, the pan-European R&E network. This followed a series of meetings between DANTÉ and the EVN leading to approval for such a project from the NREN Policy Committee (NREN-PC). The meeting was attended by technical representatives from JIVE, six EVN telescopes, six European NRENs and other networking experts.

Information was exchanged and agreements were reached about the targets of the project and a time-plan was formulated. For each group the objectives agreed upon were as follows:

- For the EVN:
  - Feasibility of eVLBI: Costs, timescales, logistics.
  - Standards: Protocols, parameter tuning, procedures at telescope and correlator.
  - New Capabilities: Higher data rates, improved reliability, quicker response.
- For the NRENs:
  - To see significant network usage with multiple Gbit streams converging on JIVE.
  - Up to 5 telescope sites supported (6 including Westerbork).
  - Must be seen to enable new science.
  - Support will be best effort IP service limited resilience.

## 2. Project Status

The project has progressed largely according to plan. In the EVN, deployment of Mark 5 has provided a uniform interface to telescopes, correlator and the networks. Disk recording enables various eVLBI tests to be performed on and off-line without disruption to normal operations and, in some cases, ahead of the provision of a high bandwidth connection. The DWDM connection to JIVE has been upgraded to six Gigabit Ethernet lines and connections into the Netherlands were upgraded to 10 Gb/s well ahead of schedule. Progress with local connections to telescopes has been slower but toward the end of 2003 1 Gb/s connections to Westerbork and Onsala came on line. With Jodrell Bank's existing 150 Mb/s connection these three stations were able to take part in a series of eVLBI tests. Through these tests and some laboratory work at JIVE, good progress has been made with understanding the network interface and configuration issues, and the evaluation of optimum data transfer parameters.

## 3. Milestones

- May 2003: First use of FTP for VLBI session fringe checks.
- September 2003: Mark 5 - Mark 5 data transfer between Bologna and JIVE.
- October 2003: First light on Westerbork Gigabit connection.
- November 2003: International baseline, eVLBI fringes, only 15 minutes after observation.
- November 2003: Onsala Space Observatory connected at 1 Gb/s.
- January 2004: First European eVLBI image (see Figure 1).

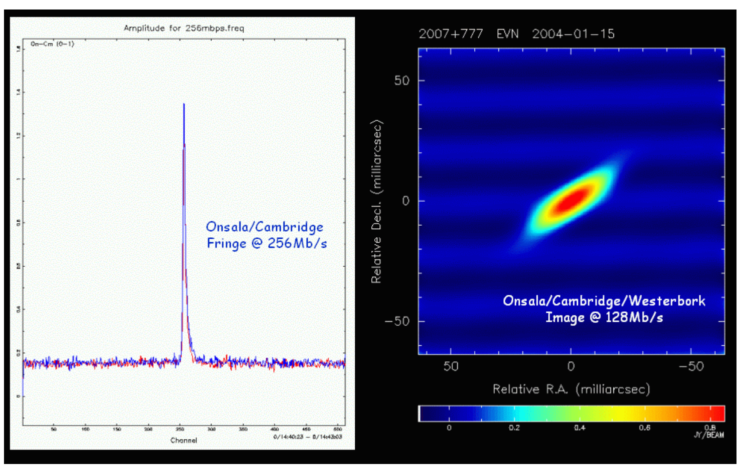


Figure 1. First European eVLBI Image

## 4. eVLBI Facilities at JIVE

Mark 5 is now used exclusively for eVLBI experiments at JIVE. A Cisco 15252 DWDM unit, connecting to one fibre pair from Amsterdam, provides six Gigabit Ethernet lines, using LX Optics, expandable to up to sixteen lines. A dedicated line from Westerbork has also been installed that is completely isolated from other network elements. A number of experimental environments are provided by these facilities with varying levels of isolation from external network elements and competing traffic.

## 5. Tests and Results

The built-in transfer processes of the Mark 5 (in2net, net2disk etc.) have been used in different combinations for tests involving various levels of disk buffering. Some tests were also performed using TCP and UDP based network performance monitoring programmes. The results of these tests are collected in Table 1. A number of questions arise from these results that are the subject of ongoing investigation:

Table 1. Test Results (J indicates jumbo-frames)

| | Memory-Memory | | Disk2net2Disk | | In2net2Disk | In2net2Out |
|---|---|---|---|---|---|---|
| | UDP | TCP | UDP | TCP | TCP | TCP |
| Bench - patch cable | | 960Mb/s (J) | 250Mb/s | 544Mb/s | | |
| JIVE-AMS-JIVE | 500Mb/s | 360Mb/s | 341Mb/s (J) | 456Mb/s (J) | | |
| Wb-JIVE | 867Mb/s | 680Mb/s | 249Mb/s (J) | 378Mb/s (J) | 256Mb/s | |
| Bologna-JIVE | 670Mb/s | 128Mb/s | | 307Mb/s | | |
| Onsala-JIVE | | | | 177Mb/s | 256Mb/s | |
| JIVE-Haystack | 612Mb/s | | | 71Mb/s | | |
| Jb (150Mb link)-JIVE | | | | | 64Mb/s | 64Mb/s |

- Why are data rates lower when disk access is involved? What can be done to improve this?
- JIVE-Amsterdam-JIVE and Westerbork-JIVE connections are private lines, with no competing traffic. Why are they slower than via a patch cable on the bench?
- When the underlying protocol for Mark 5 disk-disk transfers is changed from TCP to UDP the data rate goes down. How can this happen?
- Data rates across the European networks vary by a factor of two. Is this typical?

The data rates detailed in Table 1 where achieved after a lengthy investigation to find the optimum values for various parameters. Results of this investigation are given in Table 2. So far the underlying logic of these numbers is not well understood.

Table 2. Parameter Settings (Disk2net2Disk, TCP)

| recv/send (Bytes) | socbuf (Bytes) | data rate (Mb/s) |
|---|---|---|
| 268320 | 134160 | 331 |
| | 268320 | 421 |
| | 536640 | 425 |
| | 1073280 | 412 |
| 2097152 | 1048576 | 522 |
| | 2097152 | 509 |
| | 3145728 | 496 |
| | 3997696 | 453 |
| 4194304 | 1048576 | 521 |
| | 2097152 | 544 |
| | 3145728 | 506 |
| | 3997696 | 524 |

In an effort to move the enquiry forward JIVE is currently upgrading two Mark 5 units with Intel(R), SE7501BR2, server grade motherboards. This will confirm or otherwise that some of the performance limitations are a function of the underlying PC hardware. Installation of the new motherboard was moderately easy and operation as a normal Mark 5 unit has been confirmed.

## 6. Correlator Developments for Real-Time

In a real time system data will arrive at the correlator with some delay comprising a network delay and an instrumental delay caused by various bus/memory/peripheral manoeuvres within the interface units. Within Europe the network delay is in the range 10-30 ms. If other delays are known and stable, the correlator time reference can be synchronised, with a suitable delay, to UT. The existing input buffer will be able to absorb differential delays across the networks. To achieve this both hardware and software modifications are required at the correlator and this is the subject of one package of work at JIVE. If delays are unpredictable it may be necessary to provide a bigger buffer, and if this is done the synchronisation requirement can be relaxed. The Mark 5 has a 512 MByte buffer in the output path which, at the highest data rate, holds up to four seconds of data; more than enough for our needs. Tests are currently under way at JIVE to characterise the performance of Mark 5 in these direct-to-network modes. Of interest are the data rates achievable and the behaviour of the servo mechanism.

## 7. Forward Planning

In January 2004 NREN and EVN representatives met again to review progress one year on. The overall conclusion was that the project is largely on schedule. Developments for real-time operation are seen as the next priority. Plans are in place to achieve multiple telescope real-time operation by mid-2004. There are good prospects for a Gigabit connection to Torun on this timescale. Other telescopes, including Jodrell Bank and Medicina, do not expect to get a connection until the end of 2004 but DANTÉ and the NRENS are now willing to continue support for another year. In the coming three months attempts will be made to understand/remove Mark 5 data rate limitations and to develop the use of UDP for higher rates. Higher rate tests between Bologna and JIVE, using UDP, are also planned. Within six months further tests between Wb, On & Jb should be possible as link speeds improve and more NREN PoP JIVE tests will be staged to push larger data volumes across the European networks.

The meeting agreed that the original goal of six telescopes in real-time at 1 Gb/s was unrealistic. Many stations are unlikely to get more than one Gigabit Ethernet connection, and Gigabit Ethernet is actually limited to something less than 1 Gb/s. The existing VLBI system with Mark 5 is also restricted to fixed octaves— 64 Mb/s, 128 Mb/s, 256 Mb/s, 512 Mb/s, 1 Gb/s.

The revised goals agreed are real-time eVLBI at 512 Mb/s with at least three telescopes participating. Non real-time tests will be performed to maximise network loading up to the limits of Mark 5 and Gigabit Ethernet.

## 8. Conclusions and Acknowledgements

For the practitioners of VLBI the benefits of a real-time system are clear. Meanwhile the NRENs recognise VLBI as a genuine high-demand application, that can already source very large data flows from all corners of Europe. The big-science image of VLBI is also very welcome to the administrators and policy makers who need eye-catching publicity to maintain the profile of their product.

There are many practical difficulties however underlying these superficial enthusiasms. The astronomical community is, understandably, reluctant to commit heavily from very limited funds for local loop connections. During the normal VLBI session, eVLBI tests are an unwelcome distraction to already busy operations staff. Outside the sessions coordinated reservation of expensive telescope time can be difficult to achieve. At the same time the limited service afforded by the networks compounds any problem. No special tuning of the network for VLBI traffic is provided and technical problems have to take their place in the normal service queue, at a low level of priority.

For the above reasons the success of the project depends on the drive, effort and cooperation of individuals in both communities who are willing to work a little outside the normal parameters of their jobs. The experience of the first year of the project tells us that real progress only begins when real data starts to flow across real networks; and the key to this is to have telescopes connected directly to the network. When this begins to happen technical problems can be identified and engaged, and identifiable progress can be made. The following roll-call identifies some of the individuals and organisations who have made the progress achieved so far possible:


- European (R&E) networks: GARR, NORDUnet, SUNET, SURFnet, UKERNA, DANTÉ
- At the telescopes:
  - Michael Lindqvist, Michael Olberg: Onsala Space Observatory
  - Paul Burgess, Ralph Spencer: Jodrell Bank Observatory
  - Tony Foley, Hanno Holties: Westerbork Synthesis Radio Telescope
  - Giuseppe Maccaferri: Istituto di Radioastronomia, Medicina
- At JIVE/ASTRON:
  - Cormac Reynolds, Schedules
  - Arpad Szomoru, technical co-ordination
  - Sergei Pogrebenko, science advisor and software correlator
  - Klaas Stuurwold, network facilities and NREN liaison.
- Richard Hughes-Jones of UoM for network performance measurement software.
- The Haystack Mark 5 team for the Mark 5 itself and helpful support.
- Onsala and Westerbork Observatories management and staff for providing the high-bandwidth links to their telescopes that have got this project off the ground.


## References


[1] Hughes-Jones, R., Parsley, S., Spencer, R., High Data Rate Transmission in High Resolution Radio Astronomy vlbiGRID FGCS:- iGRID2002 conference special issue, 2002

[2] Parsley, S.M., VLBI Participation in Next-Generation Network Development, eVBLI Workshop, Haystack Observatory, 2002

[3] S.M. Parsley, P. Burgess, R. Hughes-Jones, A. Mujunen, S. Pogrebenko, J. Ritakari, R. Spencer, A. Szomoru, H. Verkouter, eEVN, ASP Conference Series, New Technologies in VLBI, 2002