

e-EVN progress

3 years of EXPReS

Arpad Szomoru¹

Joint Institute for VLBI in Europe

P.O. Box 2, 7990 AA Dwingeloo, The Netherlands

E-mail: szomoru@jive.nl

During the past years e-VLBI has undergone a remarkable development. At the start of the EXPReS² project in March 2006, the e-VLBI network was limited to a handful of European telescopes connected at data rates not exceeding 128 Mbps. Currently telescopes around the world routinely participate in global e-VLBI sessions at data rates of 512 Mbps and more. This means that both resolution and sensitivity of the e-EVN have reached the same level as that of the traditional EVN, using data recorded on magnetic media. In this paper I will give a short overview of the history of the development of e-VLBI in Europe and the hard- and software efforts in EXPReS.

*Science and Technology of Long Baseline Real-Time Interferometry:
The 8th International e-VLBI Workshop - EXPReS09
Madrid, Spain
June 22-26, 2009*

¹ Speaker

² EXPReS is an Integrated Infrastructure Initiative (I3), funded under the European Commission's Sixth Framework Programme (FP6), contract number 026642

1.Introduction

When the EVN switched from using tape recorders to PC-based disk-recording units [1], e-VLBI, transferring data via the Internet, immediately became a fascinating possibility. At the same time, national and international research networks such as the Dutch SURFnet and the Pan-European GÉANT network were quite keen on testing their expanding capabilities through demanding applications like e-VLBI.

In 2002 a European e-VLBI demonstration took place at the iGRID 2002 [2] conference, during which astronomical data were transferred from Manchester and JIVE to the venue in Amsterdam. This was followed by a proof-of-concept (PoC) project, aimed at connecting up to 4 telescopes in Europe in real-time to the EVN correlator located at JIVE, in the Netherlands.

The development of e-VLBI was greatly accelerated by the start of the EC-funded EXPReS project in March 2006. Now nearing its conclusion, this project has led to the establishment of an operational global e-VLBI network capable of doing scientific observations for extended periods of time at 512 Mbps and beyond, illustrated in Figure 1.



Figure 1. Current status of the e-EVN and EXPReS partners

Throughout the EXPReS project, a number of high-profile e-VLBI demonstrations were conducted at conferences and meetings around the world. While these demonstrations were primarily intended to showcase the achievements of the project to interested parties, such as the networking community, they also served as focal points in the development of e-VLBI, and in some cases greatly accelerated these developments.

The PoC and EXPReS projects and their motivations have been discussed in detail in [3]. In the following sections I will describe the various hard- and software developments that have taken place during the past years.

2. Networks and connectivity

The start of EXPReS coincided with the roll-out of the fully-hybrid SURFnet6 and GÉANT2 networks. In order to protect “normal” Internet users from special applications that generate extremely large data streams, these networks not only provide IP-switched connections but also point-to-point connections called lightpaths. Very early in the project, several telescopes in Europe were connected via 1-Gbps dedicated lightpaths across GÉANT2 to JIVE.

Inside Europe, the main obstacle in connecting telescopes was the so-called last mile problem, physically digging trenches and laying fibers to connect the often-remote sites to the Points-of-Presence (PoPs) of the local National Research and Educational Networks (NRENs). Outside Europe however other additional problems had to be dealt with.

2.1 European connections

By far the largest effort (mainly funded by MPIfR) went into connecting the biggest European EVN dish, the Effelsberg telescope. The inclusion of Effelsberg was of course an important milestone, as the resulting boost in sensitivity finally put the e-EVN on a competitive level with disk-based operations. Sharing the 10-Gbps connection from the telescope to the institute in Bonn and from there to the Netherlands with e-LOFAR also opened the possibility to go to a full 1024 Mbps. This is important, as 1024 Mbps of VLBI data simply does not fit through a standard 1-Gbps (1000-Mbps) interface. In much the same way it has become possible to connect Onsala at 1024 Mbps by sharing connectivity with the (future) Onsala LOFAR station.

At the start of the EXPReS project the Medicina, Torun, Onsala and Westerbork radio telescopes were connected to JIVE at 1 Gbps, with Jodrell Bank connected at 2 times 1 Gbps and Metsähovi at 10 Gbps. Since then, besides the already mentioned improved connections to Onsala and Effelsberg, Torun has upgraded its equipment to 10 Gbps. The Westerbork connection was upgraded to 2 times 1 Gbps, through the use of CWDM (Coarse Wave Division Multiplexing) equipment. By dividing the data in round-robin fashion over the two links, a full 1024 Mbps connectivity has been established between Westerbork and JIVE; this method has also been applied to the connection to Jodrell Bank.

2.2 Global connections

2.2.1 China

The Seshan telescope was connected at 1 Gbps to Shanghai Observatory early on in the project. From there, the data can either flow via the EC-funded trans-Siberian ORIENT connection (at 2.5 Gbps), or via a 622-Mbps lightpath HongKong – Canada – JIVE, provided by the Global Lambda Integrated Facility (GLIF) consortium. The existence of two NRENs in China has proven to be a somewhat complicating factor in using the trans-Siberian connection.

Urumqi telescope was connected at 622-Mbps for a brief period, to accommodate its participation in the 24-hour e-VLBI demo at the opening of the International Year of Astronomy in Paris. Of the other two new Chinese telescopes, Miyun and Kunming, only Miyun has participated in regular EVN observations; unfortunately both are connected at low data rates.

2.2.2 Australia

One of the EXPReS deliverables was a demonstration of the feasibility of correlating data from Australian LBA telescopes at JIVE in real-time. In order to make this possible, the Australian NREN AARNET set up 3 1-Gbps lightpaths, again via Canada to JIVE, specially for this demo. At this time one of these paths is still in place, and is regularly used.

2.2.3 South Africa

The Hartebeesthoek telescope was connected to Johannesburg at 1 Gbps in 2008, and a connection from there to London was rented for the duration of one month, in order to participate in a global e-VLBI demo. Unfortunately, later that year the telescope itself suffered a mechanical breakdown. When repaired (hopefully in the course of 2010), the connectivity to South Africa is expected to be much improved.

2.2.4 South America

TIGO is connected via their local NREN, the EC-sponsored RedClara backbone and GEANT to JIVE. As the capacity of the South-American backbone is limited to 155 Mbps, with additional local bottlenecks, data transfers have to be carefully coordinated with the local NREN and university.

2.2.5 Puerto Rico

Arecibo participated in the very first e-VLBI sessions, but over the years the capacity of their connection to mainland USA, shared with the University of Puerto Rico, decreased to the point of being unusable. However, since 2008 Arecibo shares a much improved connection to the mainland, where the signal is picked up by AtlanticWave and transferred via VLAN to SURFnet and JIVE. Restrictions exist, but 128 Mbps is always allowed, with special windows for 512-Mbps transfers.

2.3 Local connections

To accommodate the rapidly growing number of international connections the local network at JIVE was completely re-designed, and built around an HP router of sufficient capacity. The current situation is that SURFnet provides 8 1-Gbps lightpaths and 1 10-Gbps switched connection, capped at 5 Gbps. VLANs come in over the 10-Gbps connection, but are not counted towards the 5-Gbps limit (see Figure 2).

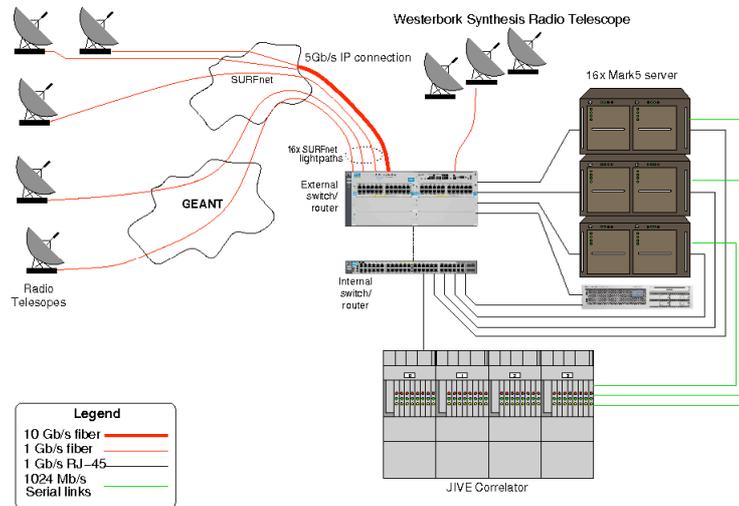


Figure 2. Network configuration at JIVE

3. Hardware upgrades

As real-time operations put quite different demands on the equipment used, a large number of changes took place, both at the stations and at the correlator. Some of these are:

1. New correlator control machines (dual interchangeable Solaris servers)
2. New data acquisition platform
3. Dedicated network monitoring machine
4. Mark5A-B upgrade kits, optical serial links, Correlator Interface Boards
5. Upgrade of all Mark5 motherboards, CPUs, hard disks, memory, power supplies (both at the correlator and at the e-VLBI capable stations)
6. 10 GE interfaces on router and Mark5 units
7. DWDM equipment on Westerbork – JIVE link

4. Software developments

4.1 Modifications of correlator control software

Switching from a batch-oriented recording-based mode of operating to real-time poses some interesting challenges, particularly so for a system designed to handle longitudinal tapes. For one thing, the stability and speed of the code become very important. Dealing with real-time demands forced us to take a critical look at the system. This in fact has led to some marked improvements in the behaviour of the correlator, both in e- and non-e operations, resulting among others to the longest unattended correlation jobs ever (> 12 hours).

4.2 Software tools

A large number of new tools were developed, for network and performance monitoring, error detection, station feedback, on-the-fly modifications of correlation and observing parameters, rapid fringe detection, on-the-fly fringe fitting, semi-automated post-processing, control interfaces for local and remote Mark5 units, control interfaces to set internet parameters, and others. These tools were needed to make it possible to

1. Detect errors at an early stage and provide feedback to the stations
2. Flexibly manage connections to telescopes spread around the globe
3. Produce results on a short timescale

One example is shown in Figure 3, showing the actual data throughput at the correlator.

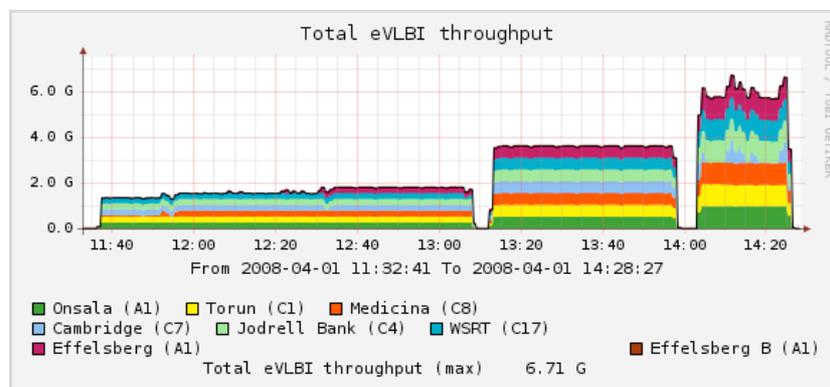


Figure 3. Data throughput measured at the EVN correlator

4.3 Mark5 software

During the first year of EXPReS, e-VLBI tests were done using the Transmission Control Protocol (TCP), the only one supported at that time by the Haystack-developed Mark5A software. The TCP protocol has been specifically designed to be fair to all users, cutting down on transfer speed as soon as packet loss (presumably caused by congestion) is detected. This is not necessarily the best strategy on dedicated point-to-point or over-provisioned links, and what is worse, the time that the protocol takes to recover from packet loss increases with the so-called Round-Trip-Time (RTT). While TCP proved reasonably successful on baselines within Europe (allowing 256 and very occasionally 512-Mbps transfers), the combination of TCP Reno (default with the then mandatory Linux 2.4 kernels) and intercontinental distances proved unworkable; first tests to Shanghai would not even reach 32 Mbps on a 622-Mbps link.

Modifying the existing Mark5A code to allow the use of the User Datagram Protocol (UDP), a so-called connectionless protocol, was only partly successful. In the end part of the Mark5 control code was completely re-written at JIVE, with rigorous thread control and the possibility to manipulate packets in various ways. This immediately led to an enormous

improvement in data transfer rates; 512 Mbps became the standard in European e-VLBI operations, causing a marked improvement in sensitivity, and long-distance links could be filled to capacity (as illustrated in Figure 4).

Further improvements in the control code included selective packet dropping (dropping packets at the sending end, and padding them at the receiving end), allowing better use of the available bandwidth at the cost of increased noise, channel dropping, which is a cleaner way of limiting data transfer, and simultaneous electronic transfer and recording on magnetic media at the stations.

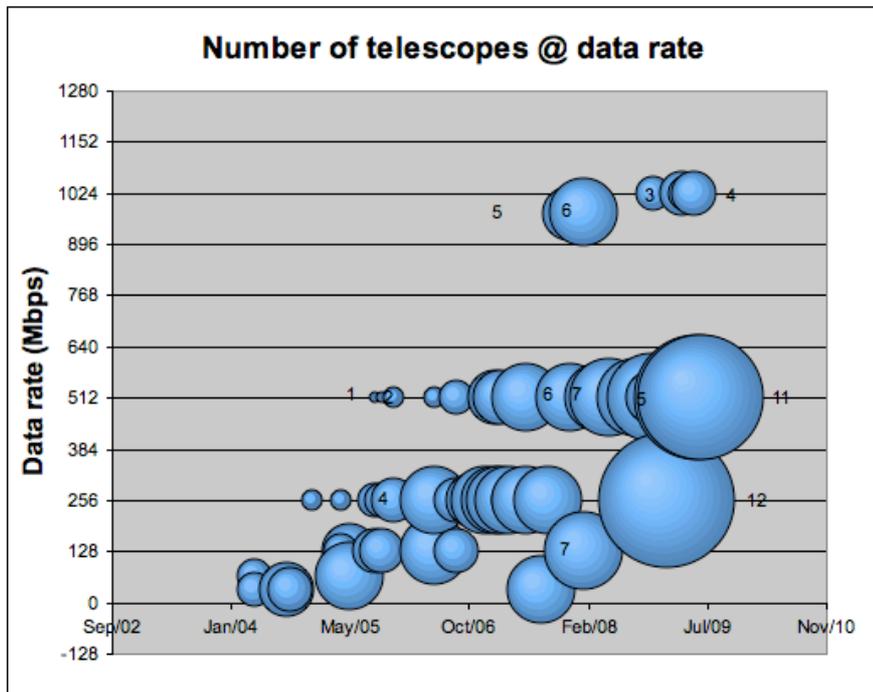


Figure 4. Number of telescopes transferring at a specific data rate versus time

5. Software correlator development

FABRIC, one of the Research Activities of the EXPRoS project, has the aim to exercise the deployment of the algorithm of VLBI correlation on Grid computing. This distributed correlator was based on the software correlator developed at JIVE for tracking the Huygens descent through the atmosphere of the Saturnian moon Titan, and builds on the expertise at JIVE and PSNC (Poznan National Supercomputing Center). A partial Grid-deployment has been tested, and the development of the correlator is going extremely well, with excellent agreement between the results of hard- and software correlation (Figure 5). Optimization efforts continue, and in a collaboration with the AutoBahn [4] project the use of dynamically allocated intercontinental lightpaths has been demonstrated.

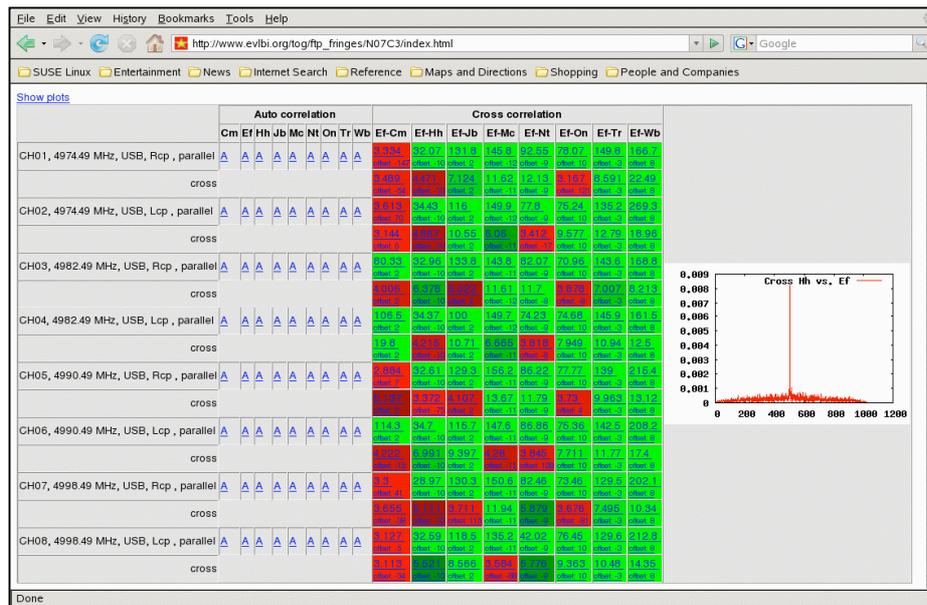


Figure 5. Web-based user interface for JIVE-developed SFXC software correlator

6. Conclusions

The last years have seen a rapid development of the e-EVN into an instrument that can compete with traditional disk-recorded VLBI, both in terms of resolution and of sensitivity. Further developments will optimize the use of available bandwidth, enable adaptive observing, possibly using sub-arrays of EVN telescopes, and add more short baselines through the inclusion of multiple Merline stations. With new telescopes coming online and the sensitivity enhancements that will become possible with increasing Internet bandwidth, the e-EVN is perfectly positioned to stay on the forefront of future technological and scientific developments.

References

- [1] A. R. Whitney, *New Technologies in VLBI*, in ASP Conference Series, Volume 306, p.393, editor Y.C. Minh, 2003
- [2] <http://www.startap.net/starlight/igrd2002/>
- [3] A. Szomoru, *Recent e-EVN developments*, in Proceedings of Science, PoS(8thEVN)053, 8th EVN Symposium, 26-29 September 2006, Torun, Poland
- [4] <http://www.geant2.net/server/show/nav.00d00a008002/targetBlock/5335/viewPage/1>, and references therein