27-10-2006

Towards e-EVN: e-VLBI and the use of lightpaths.

Arpad Szomoru, JIVE
Bram Peeters, Niels den Otter, Erik-Jan Bos, SURFNet

In this document we will assess the connectivity needs of JIVE and the EVN as regards e-VLBI, and the way to achieve this. We will first give a short overview of the projects that depend on this connectivity.

EXPReS

EXPReS is an I3 proposal to the EC (Communication & Network Development Call) that will run for three years. It formally started on March 1, 2006.

Relevant are the two Service Activities, SA1 and SA2, and the Joint Research Activity, JRA1.

SA1 aims at establishing a real-time production e-VLBI facility, connecting up to 16 telescopes simultaneously at 1 Gbps to the EVN correlator at JIVE in real-time. The work packages of this service activity will concentrate on:

- Improving the real-time operational efficiency of the correlator
- Investigating data transport issues (e.g. protocols, use of light paths)
- Providing target-of-opportunity support
- Creating real-time monitoring functions
- Upgrading the correlator at JIVE
- Seamlessly incorporating e-MERLIN telescopes (UK) in the EVN array (through VSI interfaces).

SA2 mainly deals with the last mile connections to telescopes. An amount of 800 kEuro will serve as seed money; the actual cost of connecting all EVN telescopes will be far higher. Some major SA2 deliverables are:

- Inclusion of the (e)-MERLIN array as an element of the EXPReS e-VLBI facility, through the upgrade of the Jodrell Bank Observatory link to JIVE to 10 Gbps
- Establishment of a 10 Gbps link between Jodrell Bank Observatory (UK) and Onsala Space Observatory (Sweden), a test-bed for prototyping ultra-high-data rate VLBI studies (4 and 8 Gbps).
- Connection of (EVN) radio telescopes located outside of Europe to JIVE, in particular telescopes that already have clear connectivity ambitions or are

already in the process of establishing these – e.g. Arecibo in North America (Puerto Rico), Hartebeesthoek in South Africa, Shanghai Astronomical Observatory (Shanghai), Urumqi Astronomical Observatory (Nanshan, China) and TIGO in South America (Chile).

- Connection of Australian radio telescope facilities to JIVE in order to conduct e-VLBI tests (including spacecraft navigation demonstrations) and to promote the remote access (and real-time, local on-line data processing) of radio telescope facilities in Australia from astronomers located in Europe. This latter activity has close synergy with the way in which European astronomers may interact with the Square Kilometer Array (SKA), the next generation radio telescope, likely to be located in the Southern hemisphere.

JRA1, with the acronym **FABRIC** (Future Arrays of Broadband Radio-Telescope on Internet Computing) is a joint research activity, in which many institutes in Europe collaborate (led by JIVE). Its key objective is *the development of a prototype for high data rate e-VLBI that also takes advantage of distributed correlation*. The following components are the vital areas of development:

- Data acquisition platform and formatting for future e-VLBI:
    - Develop optimal buffering between the digital output of the radio telescope and transport protocol.
    - Allow for local quality control, calibration operations and recording of hard copies.
    - Establish (or identify) a transport protocol that optimizes transport, including copies for different routes (multicast).
    - Push data rates well beyond 1 Gbps, develop an interface capable of sustaining such rates between public and dedicated networks in order to integrate existing facilities (e.g. LOFAR).
- Distributed software correlator:
    - Develop code that will run on standard workstations using open source compilers, deployable on Grid nodes
    - Conduct a programme of research designed to improve data processor algorithms, allowing a variety of geometric models and enabling zooming functionality for spacecraft navigation and spectral line applications
    - Format output data such that it can be transparently merged into central archive, available to astronomers in real-time with automatic pipelining features
    - Explore the limits of Grid computing and evaluate its usefulness for future operational correlation
    - Demonstrate the feasibility of distributed correlation with a limited number of low-bandwidth data streams (e.g. 5 * 128 Mbps)

SCARIe

SCARIe (Software Correlator Architecture Research and Implementations for e-VLBI) has been funded by NWO to an amount of 500 kEuro under the STAR E-Science (STARE) program. It is a collaboration between JIVE, SARA and the UvA and has started on the first of October 2006. Its aims are related to those of FABRIC, but with a distinct Dutch emphasis. It should lead to a prototype software correlator using Grid based computing connections; DAS3 and StarPlane are strategically positioned to support this project. Specifically the following points will be addressed:

- Improving high-bandwidth connectivity
- Scheduling and running of a distributed Grid-based correlator, multicast, load balancing
- Selection and implementation of Grid middleware
- Storage of raw, correlated and intermediate data products
- Monitoring of the quality of correlation
- Demonstrating the concept using data from an existing data repository

Connectivity needs

For SA1, the production e-VLBI service, JIVE needs to be able to simultaneously hook up 16 telescopes to the correlator, at 1 Gbps. Throughout the project, a large number of tests will be conducted, between JIVE, the stations and computing platforms located at various universities.

As soon as telescopes get connected to GÉANT2 (with the help of SA2), they start participating in SA1. This will most likely happen first for a number of European EVN telescopes. However, SA2 also calls for connections to various places outside of Europe. Particularly important will be China and Australia. Construction on two brand-new radio telescopes in China, which will be part of the EVN network, has just finished. Four telescopes of the Australian VLBI network have been connected to their correlator through glass fibres. A very important demonstrator for EXPReS will be to connect a number of Australian and/or Chinese telescopes to the correlator at JIVE, correlating the data of what will be effectively a sub-array located on the other side of the world. This type of remote observing will also be a powerful conceptual demonstrator for the way the future Square Kilometer Array could be run. Technically we think such a demo will be feasible within a year (~mid 2007).

Also, and this has proven very important in the past not only because of the technical challenges but also for PR purposes, demonstrations will be conducted at major events (such as iGrid and SuperComputing) in collaboration with the American and Japanese e-VLBI networks, connecting telescopes unrelated to the EVN to JIVE and EVN telescopes to correlators abroad.

Bandwidth needs for FABRIC should be fairly modest, as mainly reduced data products will be transported to JIVE from various nodes across Europe. The Westerbork Synthesis Radio Telescope will have to be connected through JIVE however, and depending on the model, the raw data stream will have to be transported to a number of nodes across Europe (possibly time-sliced).

Finally, SCARIe will possibly generate 4 to 5 pre-recorded raw data streams (up to 1 Gbps each) from JIVE to the DAS3 network, while the data products that are returned to JIVE will be much more modest in size.

Current status

The following table summarizes the current situation and the expected upgrades. The first part of the table lists the telescopes that are already being used regularly in e-VLBI experiments.

Telescope	**Status**	**Comments**
JIVE	6 x 1 Gbps	16 x 1 Gbps upgrade 2007
WSRT (NL)	1 Gbps	-
Onsala (SE)	1 Gbps	10 Gbps upgrade 2007
Torun (PL)	1 Gbps	-
Jodrell Bank & Cambridge (UK)	2.5 Gbps	10 Gbps upgrade 2007 (inclusion of additional MERLIN telescopes)
Arecibo (USA)	155 Mbps	1 Gbps upgrade
Medicina (IT)	1 Gbps	-
Metsahovi (FI)	10 Gbps	-
Yebes (ES)	Not yet operational	10 Gbps planned 2006/7
VIRAC (LV)	Not yet operational	1 Gbps planned 2006/7 VLBI capable 2006 ?
Shanghai (CN)	Not yet operational	2.5 Gbps available. Uncertain connectivity between Europe & China
Miyun (CN) & Kunming (CN)	Not yet operational	0.1-1 Gbps planned Last mile: few km
Urumqi (CN)	Not yet operational	Fibre installed – 2 Mbps 0.1-1 Gbps planned

Hartebeesthoek (SA)	Not yet operational	1 Gbps planned as part of SKA ambitions Poor connectivity between South Africa & Europe
TIGO (Chile)	< 20 Mbps	Fibre installed, 64 Mbps possible?
Noto (IT)	Not yet operational	Last mile problem
Sardinia (IT)	Not yet operational	Last mile problem
Effelsberg (DE)	Not yet operational	Fibre being installed, 10 Gbps in 2007
LBA (AUS)	1 Gbps between telescopes and LBA correlator	Several Gbps needed between Australia and Europe

Future connectivity: IP vs. lightpaths

From the previous sections it can be concluded that the connectivity needs at JIVE will grow over the next years to 16 * 1 Gbps, to support an e-VLBI production service, and that about 6-7 * 1 Gbps will be needed to accommodate FABRIC, SCARIe and general network testing.

Several years ago it was recognized by SURFnet that a traditional IP-routed network would not suffice to combine the extremely large data streams that a few (mostly scientific) applications generate, with an undisrupted internet service to a large number of low-bandwidth users. Because of this, SURFnet6 was designed as a hybrid network, providing both normal IP-routed traffic and point-to-point lightpath connections. Likewise, the GÉANT2 network is also designed and now built as a hybrid network.

With the new possibility of providing direct lightpath connectivity a superior alternative to the IP routed network exists. Trying to achieve the above connectivity through the IP network would place a tremendous load on SURFnet and cause congestion on the GÉANT2 network.

Perhaps as important, apart from causing congestion e-VLBI would also suffer from competing traffic, which may become an important issue once the LHC comes online. In fact problems of this kind already regularly occur during e-VLBI sessions.

JIVE and the EVN should aim to make full use of this capability. Ideally, lightpath connections should be established between each telescope and JIVE. This will give the optimal quality of connection with guaranteed bandwidth and no competing traffic.

Practical considerations

SURFnet—JIVE

The current connectivity of JIVE will have to be upgraded. SURFnet and JIVE have agreed, in principle, on the topology illustrated in Figure 1.

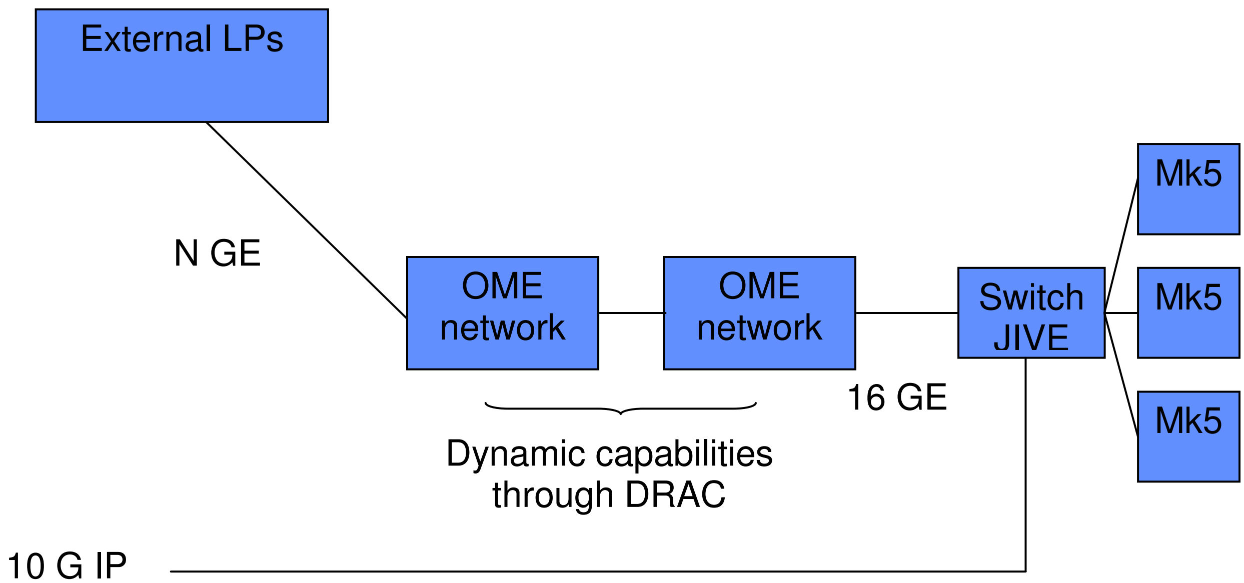

Figure 1: proposed topology of SURFnet connections to JIVE

Telescopes will be connected directly to a Jive switch at Dwingeloo, using lightpaths through the SURFnet network. Where and when appropriate the dynamic capabilities of DRAC can be used to dynamically select new or other telescopes coming in through the lightpath network. To provide connectivity to sites without lightpath capabilities and for >1 Gbps transfers, for example to software correlators, an IP connection will be available as well. This IP connection is rate-limited at 5 Gbps to safe-guard other traffic on the 10 Gbps uplink. Heavy traffic over IP connections should be off-loaded using lightpaths as much as possible.

Telescopes – NRENs – GÉANT2

A total of 15 NRENs participating in DANTE, that are part of the dark fibre cloud, have access to a GÉANT2 10 Gbps IP connection and a 10 Gbps connection aimed at providing lightpaths as part of their collaboration. In order to create dedicated 1 Gbps connections across the GÉANT network to SURFnet, the NRENs will have to make a formal request, together with SURFnet. This means the telescopes will have to negotiate this with their local NRENs.

It is possible some stations will be asked to pay for this service. As said, the 15 NRENs that are part of the dark fibre cloud have basic access to 10 Gbps lightpath connectivity (allowing for 9 x GE). Considering the price for an additional 10G connection to the GÉANT network is 20k€ per year, it is hard to imagine that when a fee is indeed required, a 1 Gbps lightpath will be beyond the means of an EVN station.

As far as the connection from telescopes to their local NREN is concerned, this too will have to be negotiated by the stations. In the cases where the connections are shared with other institutes, a strong argument can be made for the use of lightpaths. This may be harder to argue for stations with over-provisioned IP connections; however these will still put considerable stress on the NREN's network between the telescope's access point and the GÉANT IP connection and on the GÉANT and SURFnet network. This makes it clear that even in cases when the desired end-to-end lightpath connectivity can not be realized we should still aim for a situation in which transport across Europe is done via lightpaths, regardless of the way the data reaches GÉANT2.

Practical steps

- JIVE and SURFnet will start the upgrade of the connectivity between SARA and JIVE within the next few months. This will ensure that it will be possible to connect up to 16 telescopes simultaneously via dedicated lightpaths to the EVN correlator. The availability of a 10 Gbps lambda will provide ample capacity for the various research and development projects that JIVE and the EVN are involved in.

- The stations should contact their local NRENs and discuss to tackle the following points

    - Lightpath connections from the telescopes to the nearest GÉANT PoP or other parts of the GLIF are the only way to guarantee an un-congested dedicated connection to the correlator. The point should be argued with the NRENs that this will benefit their normal users, by shielding them from the disruptions caused by large e-VLBI data flows.
    - In case this is not possible (e.g. because of hardware limitations), an attempt should be made to ensure optimal connectivity through over-provisioning or any other means.
    - Whether a lightpath between telescopes and GÉANT exists or not, a request should be made to the NRENs for a lightpath connection to SURFnet. As stated, all NRENs participating in DANTE have access to (at least) 10 Gbps lightpath connectivity across Europe. It should be stressed that directing these large data flows through lightpaths instead of the routed connection will also be in their interest. The NRENs will have to make a formal request, together with SURFnet, for lightpath connections across GÉANT. SURFnet is more than willing to assist the stations in their negotiations with the NRENs, but obviously cannot initiate this process.