

NEXPreS and the EVN

Arpad Szomoru¹

Joint Institute for VLBI in Europe

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

E-mail: szomoru@jive.nl

Paul Boven

Joint Institute for VLBI in Europe

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

E-mail: boven@jive.nl

This paper describes two important developments in the European VLBI Network (EVN). The first part deals with enabling higher bandwidth observations, both recorded and in real-time, while the second part focuses on the use of dynamic, international light paths for transport of VLBI data. Research leading to these results has received funding from the European Union's Seventh Framework Programme (FP7/2007- 2013) under grant agreement n° RI-261525 NEXPreS. This material reflects only the author's views, and the European Union is not liable for any use that may be made of the information contained therein.

11th European VLBI Network Symposium & Users Meeting

October 9-12, 2012

Bordeaux, France

¹ Speaker

1. Towards 4Gbps VLBI

1.1 Introduction

One of the aims of the NEXPreS project, funded by the European Commission, is to prepare the EVN for the higher bandwidths that are becoming available. With the rollout of digital base-band converters (DBBC) throughout the EVN and the steady increase of international network capacity, 4Gbps operations may soon become a reality, both recorded and in real-time. As part of one of the deliverables of the project, a first-ever 4Gbps EVN test observation was set up to take place in June 2012, involving the telescopes at Effelsberg, Onsala and Yebes, in collaboration with the NORDUnet, SURFnet, RedIRIS and GÉANT research networks.

1.2 Setting up a 4Gbps EVN demonstration

The specific aim of the demo was to exercise the new observational modes that will become the standard in the EVN in the next years. The first of these was the combination of high-speed recording (2 and 4 Gbps) on disk packs at the stations, while simultaneously transferring a subset of the data in the form of a 1-Gbps data stream to the correlator in Dwingeloo. In the future, this mode will become important for those stations that have 4Gbps-capable receiving and recording equipment but limited connectivity. The second mode was to be real-time correlation of the full data streams at 2 and 4 Gbps, with or without recording at the stations.

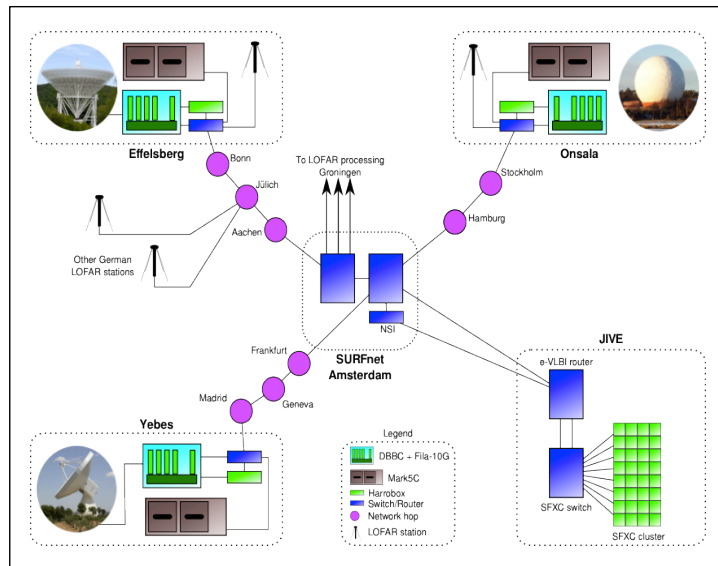


Figure 1: 4Gbps network and equipment

At the time it was decided to have this demo (January 2012), the Mark5Cs in the EVN had not been tested at 4Gbps, the Fila10G boards, needed to output DBBC data on a 10GE connection, could only produce Mark5B format data at 2Gbps, only the DBBC at Ef was equipped with a Fila10G, only three Fila10G boards were available in all of Europe (one of which was consequently sent to Chile for testing with APEX), the DBBC at Yebes needed

repairs, there was no Mark5C at Yebes and the new fibre to Yebes had not been lit yet. During the period January - June 2012 many telecons took place involving the Mark5C and DBBC designers, VLBI friends and operators at the telescopes and JIVE staff. New software and firmware was written, new modes of the DBBC tested, one Mark5C was shipped from JIVE to Yebes, powerful 1U server computers with multiple 10G interfaces purchased and installed at the stations, optical modules, cables and various bits of equipment shipped across Europe. The final missing Fila10G board was brought back from Chile as hand luggage to Germany, where it was promptly impounded by customs. Fortunately, it only took about a week to resolve that problem.

Although a lot of tests of equipment and network connectivity were done offline, tests using real observations were of course needed as well. As the EVN block session immediately preceded the demo, suitable gaps in the schedule had to be found, during which all three telescopes would be available for several hours. During these tests wide variety of problems were found and solved, ranging from pure software bugs to issues of configuration, scheduling, and even the maximum length of CX4 cables that could be used.

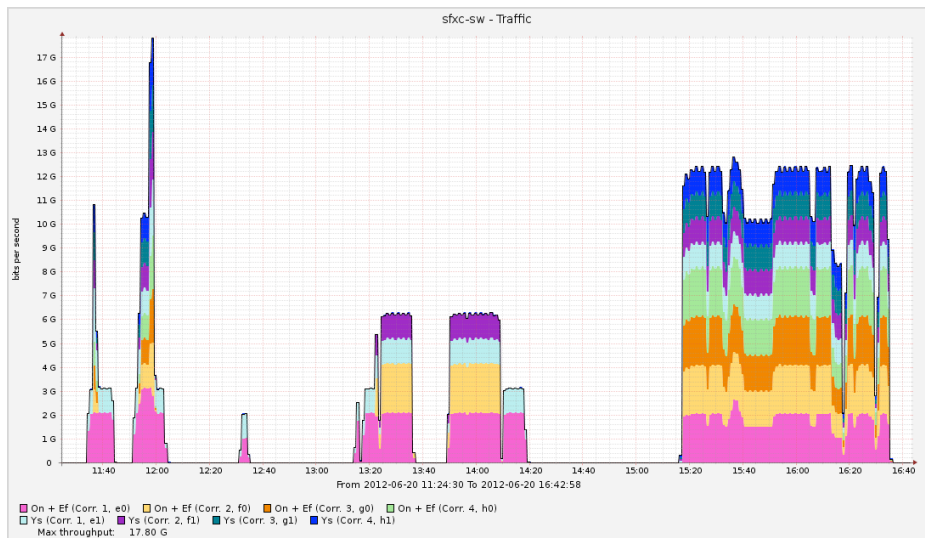


Figure 2: Data throughput at JIVE. The 18 Gbps peak is caused by internal flooding

1.3 Results

The demo took place on June 20, and started around 10 CET. A number of minor issues had to be dealt with immediately. After this, basically everything worked out of the box. Data were recorded at the stations at 4Gbps and simultaneously transferred to JIVE for real-time correlation on the SFXC software correlator, at 1, 2 and even 4Gbps. The networks performed flawlessly, barring some local problems with flooding. Unfortunately though, no fringes were found.

Most of the time that day was spent on trying to find fringes, ftp'ing bits of data from the Mark5C recorders to JIVE to explore different combinations, different time delays, but to no avail. The most likely culprits are the time synchronisation of the Fila10G, or the configuration of the DBBCs. Considering we were using practically brand-new equipment in never before

tried modes, this is not altogether surprising. But once this final problem is sorted, this mode will be ready for operational use, as we have shown that all parts of the equipment, from the stations to the correlator, can handle and process these data streams.

1.4 Conclusion

In spite of the lack of fringes, the demonstration was a technological success. And although the preparations of the demo at times caused a lot of stress for operators and engineers at the stations and JIVE, it was also an excellent opportunity for the EVN staff to showcase their technical competence and inventiveness as well as a demonstration of the strength of the international collaboration that forms the basis of the EVN.

2. International Bandwidth-on-Demand

2.1 Introduction

The EVN is able to perform e-VLBI observations thanks to the availability of network capacity at several National Research and Education Networks (NRENs), GÉANT, and the Global Lambda Integrated Facility (GLIF). As our traffic consists of constant rate UDP packets at high bandwidth, it doesn't coexist well with other network users, and we benefit from having most networking connections available as dedicated circuits, either SDH based lightpaths or other types of private international paths. The operational bandwidth for most e-EVN stations is 1024Mb/s of payload data, which already exceeds the capacity of a regular 1Gb/s Ethernet connection. This necessitates the use of pairs of 1Gb/s connections, or 10Gb/s Ethernet links and equipment [1].

The EVN is not a full time instrument, and its configuration changes depending on the needs of the observation that is being carried out. At the moment all our international links are statically configured (always on), an inefficient use of scarce international connectivity. An international path crosses multiple administrative domains, which makes building these international paths a fairly labour-intensive operation, which is why they are left in place between observations. As the EVN moves to higher bandwidths for higher sensitivity, with stations going to 4Gb/s capacity in the near future, and higher speeds already on the roadmap, we should find a more efficient way to provision the networking resources we need.

2.2 Bandwidth-on-Demand (BoD) and NSI

Several NRENs run a so called 'hybrid' network, which offers IP routing for its regular end-users, and dedicated paths (lightpaths etc.) for those e-Science users that generate so much traffic, that it would not be cost-effective to run it over the generic routed infrastructure. Hybrid networks can be both cheaper and greener, as the cost and energy usage of high speed routers is much higher than that of switches or Synchronous Digital Hierarchy (SDH) equipment. The administrative overhead of building dedicated paths within a network has prompted the development of user-interfaces for the provisioning of these paths. Exposing this interface to the end-user, and allowing them to create these paths as needed themselves, creates a new networking service within these NRENs: Bandwidth-on-Demand (BoD). Several NRENs now offer such a BoD service within their network, where an end-user can request a dedicated

networking path between two points within the NRENs network, and specify start and end-time, bandwidth and other parameters to get exactly the path that they want.

International BoD is much more challenging, as such a path needs to cross administrative boundaries, equipment from different vendors, and even different transport technologies. The Open Grid Forum (OGF), as part of their inter-domain resource management effort, has created the Network Services Infrastructure (NSI) [2] protocol. The Connection Service of this protocol (NSI-CS) describes the messages to be exchanged and a distributed state-machine, that allows network providers to accept and process requests for paths that cross their domain boundaries. Although the protocol is still under development, it is well accepted by the international community of NRENs, and several of them run testbeds with early versions of the protocol, allowing true international BoD to be tested. NEXPreS has early in the project decided to standardize on NSI for the BoD needs of the project.

2.3 The NEXPreS NSI client

Workpackage 6 "High Bandwidth on Demand" is a Service Activity within NEXPreS, tasked with building a BoD infrastructure for e-VLBI and LOFAR Long-Term Archive (LTA) data transport. As the first end-user involved in NSI based BoD, we have had to create our own NSI client that can schedule and request connection resources from network providers. The first release of the NEXPreS NSI client was demonstrated by building an international path from Onsala Space Observatory (OSO) in Sweden, to JIVE in the Netherlands.

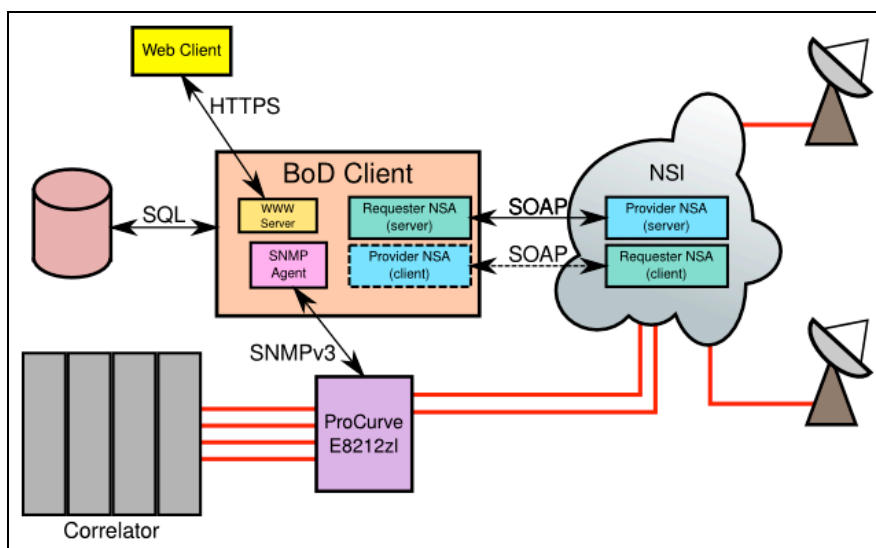


Figure 3: Block diagram of the NEXPreS NSI client

The NEXPreS NSI client is built using open-source tools such as Linux, Apache, PHP and MySQL. The source-code of the client is released under the Gnu Public License (GPL) and is available from the NEXPreS WIKI [3]. Since its initial release, development has been ongoing in the client. It now supports automated validation of a path by injecting synthetic traffic into it. Another important new feature is the ability to have the client determine which

path to take for an international connection ('Pathfinding'), and then requesting the required resources from each of the contributing network providers ('Aggregation').

With the cooperation of our partners (SURFnet, NORDUnet, PSNC) in NEXPreS and with GÉANT, we are now able to create paths of up to 10Gb/s to Onsala Space Observatory (Sweden) and Torun Observatory (Poland). A new connection to the GÉANT dynamic service is currently under construction and will provide on-demand connectivity towards Jodrell Bank (UK) and other European destinations.




Figure 4: NSI controlled 10Gb/s paths (dashed) for NEXPreS

The goal of WP6 is to integrate the NSI client with e-VLBI observations and LTA operations, where a planned observation or data transport will automatically request the required BoD paths to support the science task to be carried out. For e-VLBI, the NSI client will ingest the observation schedule, determine the required capacity and start- and end-time of the observation for each of the participating telescopes, request the necessary BoD resources, and test each of these paths well before the start of the observation.

2.4 Conclusion

Bandwidth-on-Demand is a useful new tool to support the networking requirements of e-sciences such as radio-astronomy. The adoption of the NSI standard across Europe, both in individual NRENs and within GÉANT, makes it possible to build on-demand paths over the distances that are required for VLBI. These dedicated but temporary paths offer a way to support the growing bandwidth needs of the e-EVN in an efficient and cost-effective way.

References

- [1] P. Boven, "e-VLBI Networking Tricks", Proceedings of the 8th International e-VLBI Workshop, 22-26 June 2009, Madrid, Spain http://pos.sissa.it/archive/conferences/082/036/EXPreS09_036.pdf
- [2] https://www.ogf.org/gf/group_info/view.php?group=nsi-wg
- [3] http://www.jive.nl/nexpres/doku.php?id=nexpres:nexpres_wp6