Alan R. Whitney et al.: e-VLBI Development at Haystack Observatory, IVS 2006 General Meeting Proceedings, p.211–215
http://ivscc.gsfc.nasa.gov/publications/gm2006/Whitney3

e-VLBI Development at Haystack Observatory

Alan R. Whitney, Chester A. Ruszczyk

MIT Haystack Observatory

Contact author: Alan R. Whitney, e-mail: arw@haystack.mit.edu

Abstract

Haystack Observatory is engaged in a multi-faceted e-VLBI development program. Real-time intercontinental experiments using Mark 5A data systems and the Mark IV correlator at Haystack have been performed at sustained data rates up to 512 Mbps/station. Specialized algorithms and protocols for e-VLBI are being developed to take advantage of the unique characteristics of e-VLBI. Network performance monitoring tools have been developed. And research is being conducted using dynamically switched optical paths to provide wide dedicated pipes for e-VLBI data transmission. The biggest obstacles to widespread e-VLBI are limited connectivity of telescopes and time-varying end-to-end network performance.

1. Introduction

e-VLBI development has expanded rapidly over the past several years. At Haystack Observatory, a broad program of e-VLBI development is underway [1], including network strategy and protocol development, optically-switched network development, as well as both experimental and production e-VLBI data transfers. In this paper, we will briefly describe these areas of activity.

2. Advantages of e-VLBI

e-VLBI confers three strong advantages to VLBI measurements:

1. Bandwidth growth potential for higher sensitivity: Traditional record-and-ship VLBI is limited to a few Gbps, e-VLBI has the potential to increase data rates to 10-100 Gbps, dramatically improving sensitivity and higher-precision delay measurements.
2. Rapid processing turnaround: Processing of VLBI data can be reduced from weeks/months to hours, which is important for Earth-orientation measurements, particularly UT1, which support precision navigation.
3. Confidence in equipment and setup: e-VLBI allows test data to be correlated in advance of observing sessions to verify proper equipment setup and operation.

3. History of e-VLBI Experimental Developments

A brief accounting of e-VLBI milestones at Haystack Observatory may be instructive to understand the rate of progress in e-VLBI:

- 2002:
  - Near real-time VLBI data transferred at ~800 Mbps from Westford and GGAO

  - First intercontinental e-VLBI: Westford-Kashima at ~20 Mbps
  - Real-time e-VLBI data transfer from GGAO to Haystack observatory at 288 Mbps

- 2003:
  - First e-VLBI ‘Intensive’ UT1 result in under 24 hours - Westford-Kashima

- 2004:
  - Begin routine daily UT1 Intensive e-VLBI transfers from Wettzell
  - Real-time e-VLBI between Westford and GGAO at 32 Mbps
  - First intercontinental real-time e-VLBI experiment: Westford-Onsala at 32 Mbps
  - Real-time fringes at 128 Mbps between Westford and GGAO
  - Real-time fringes at 512 Mbps between Westford and GGAO
  - Begin regular transfers of K5 data from Kashima to Haystack

- 2005
  - Real-time fringes Westford-Onsala at 256 Mbps
  - Transmission of all CONT05 data from Tsukuba to Haystack (15 days at 256 Mbps); routine e-VLBI transmission of all data from Tsukuba and Kashima
  - Real-time e-VLBI demonstration at Super Computer 2005 conference: Three-station (Westford-GGAO-Onsala) real-time e-VLBI at 512 Mbps/station. Correlation results transmitted to Seattle showroom floor in real-time. See Figure 1.

- Early 2006
  - Begin commissioning of 100 Mbps e-VLBI link to Ny-Ålesund; expect to begin routine e-VLBI transmissions soon

This succession of events shows the evolution of e-VLBI from short one-time experiments to actual production transfers over a period of 3 years. Though this progress may seem impressive, there are still many problems to overcome before e-VLBI can become routine on a global basis.

Figure 1. Real-time e-VLBI demo at Super Computer 2005 conference, Nov 2005

3.1. Challenges of e-VLBI

The list of challenges confronting routine widespread global e-VLBI is significant:

- 'Last-mile' connectivity to telescopes
  - Most telescopes are deliberately placed in remote areas
  - Major initiatives in place in Europe, Japan and Australia to connect telescopes; U.S. is lagging
- End-to-end network performance is often well below advertised rates
- Performance of transport protocols
  - Untuned TCP stacks; fundamental limits of regular TCP
- Throughput limitations of COTS hardware
  - Disk to/from Network
  - Real-time VLBI data to/from network
- Complexity of e-VLBI experiments
  - e-VLBI experiments currently require significant network expertise to conduct
- Time-varying nature of network
  - R&E networks tend to be unstable, particularly towards end points

4. Antenna Connectivity

The following list shows the IP network connectivity of many of the worlds VLBI antennas and correlators (geodetic sites are underlined):

- ≥2 Gbps: JIVE Correlator (6 × 1 Gbps); Haystack (2.5 Gbps); Westford, MA (10 Gbps to Haystack; 1 Gbps to outside world); Kashima, Japan (2.5 Gbps); Usuda, Japan (2.5 Gbps); Nobeyama, Japan (2.5 Gbps); Koganei, Japan (2.5 Gbps); Tsukuba, Japan (2.5 Gbps)
- ≥1 Gbps: GGAO, MD (1 Gbps); Onsala, Sweden (1 Gbps); Torun, Poland (1 Gbps); Westerbork, The Netherlands (1 Gbps); Medicina (1 Gbps); Jodrell Bank (1 Gbps)
- Others: Arecibo, PR (155 Mbps); Wettzell, Germany (~30 Mbps); Kokee Park, HI (nominally ~30 Mbps, but currently disconnected); TIGO (~2 Mbps); Svetloe (speed unknown)

In progress:

- Hobart - agreement reached to install high-speed fiber; details not known
- Ny-Ålesund - work in progress to provide 100 Mbps link for e-VLBI data to Haystack Observatory
- Fortaleza - funds secured for fiber connection at 2.5 Gbps; early 2006
- Metsahovi - 1 Gbps in 2006
- Zelenchukskaya, Russia - details unknown

- Badary, Russia - details unknown

Considerable effort is now being expended to connect most of the remaining antennas and correlators, but connections can be quite expensive and are often limited by budgetary resources. One major unconnected resource is the VLBA array of ten 25 m antennas.

5. e-VLBI Standardization (VSI-E)

A reference implementation of the proposed VSI-E (VLBI Standard Interface - e-VLBI) [2] is now being tested in transfers between Kashima and Haystack. This implementation is intended to act as a demonstration model for VSI-E and is available to all interested parties. The VSI-E framework provides signaling, control, framing and statistics support and is an extension to the internet RTP standard. It also provides flexibility in that it allows users to choose the transport protocol that best suits the network environment (e.g. UDP, TCP or other variants). Once the reference implementation is fully checked out and any needed changes are made, attention will be turned to optimizing the code for high-speed operation.

6. Throughput Limitation of COTS Hardware

Though the Mark 5 data system [3] can record and playback data at 1 Gbps, it is surprisingly difficult to manage those same data rates to/from network connections. Routine e-VLBI operation at 512 Mbps is readily achievable with today's systems, but operation at 1024 Mbps is difficult. In laboratory tests at Haystack Observatory using high-end PC motherboards with dual bonded Gigabit Ethernet links, data transfer rates of up to ~1280 Mbps have been achieved. To date, however, these rates have not yet been achieved in any practical e-VLBI experiments.

A new generation of PC motherboards is now appearing which replace the PCI bus with the PCIExpress bus. This high-speed serial bus promises much higher speed operation and should allow routine operation of e-VLBI at rates exceeding 1 Gbps.

7. Intelligent e-VLBI Applications

Based on observed usage statistics of a typical R&E backbone, such as the NSF 10 Gbps Abilene backbone network in the U.S., it is clear that there is much unused capacity. e-VLBI data has some unique qualities that may allow e-VLBI to effectively use some or much of this unused 'secondary' network capacity. Special characteristics of e-VLBI that might be taken advantage of include:

- Ability to tolerate some loss of data (up to ~5% or so) in many cases
- Ability to tolerate delay and jitter in data transmission, which can normally be accommodated by large buffers at both transmitting and receiving end

With support from the National Science Foundation, Haystack Observatory is developing an 'Experiment Guided Adaptive Endpoint' (EGAE) strategy designed to meet this challenge. EGAE performs dynamic adaptation based on simple high-level description of experiment requirements and its surrounding environment, expressed in language appropriate to the scientific end-user. These requirements are then translated by an application-specific translator into specific network requirements. EGAE controllers at both the network data and control planes gather information and react to network conditions by altering the data flow between the DIM and DOM so as to be a 'friendly' user on the network, specifically lowering data rates quickly when higher priority 'primary' users demand network bandwidth. A prototype EGAE implementation is now in use in production e-VLBI transfers at Haystack.

8. Dynamically Optically-Switched Light Paths

Haystack Observatory is collaborating with the University of Maryland and others in the DRAGON (Dynamic Resource Allocation via GMPLS Optical Networks) project [4], under sponsorship of the National Science Foundation. The objective of the DRAGON Project is to create dynamic and deterministic end-to-end network transport services for high-end e-Science applications, and is concentrating on on-demand creation and teardown of optically-switched light paths. This type of network conditioning is made possible by two fairly recent technological breakthroughs:

1. Dense wave-division multiplex (DWDM) technology, which allows many (more than 100) individual wavelengths to be placed on a single fiber, each capable of carrying at least 10 Gbps of data.
2. Optical switches capable of switching individual wavelengths among an array of fibers.

These two capabilities, working together, allow the dynamic creation and teardown of all-optical high-bandwidth paths on user demand. Because the switching is all-optical, expensive routers and switches, which must convert from optical-to-electrical-to-optical, can be eliminated or minimized.

Haystack Observatory is working with the DRAGON project to develop software application interfaces applicable to e-VLBI and other applications that might benefit from this technique. Ultimately, the DRAGON project may allow e-VLBI data to/from Haystack Observatory to be routed entirely (or at least mostly) optically to/from the Abilene backbone network in the U.S. This work will be fully integrated with the EGAE project described above.

9. Summary

Though e-VLBI has been developing rapidly over the past few years, much more work needs to be done before it can become truly operational. This includes not only some of the areas we have described here, but must include others as well. For example, as more high-bandwidth applications compete for limited network bandwidth, bandwidth-reservation systems and direct optically-switched paths become attractive. Many groups in several countries are working on these problems, and e-VLBI will be truly successful only with the close international collaboration of the VLBI and global networking communities.

References

[1] Haystack e-VLBI program information available at *http://web.haystack.mit/e-vlbi/evlbi.html*

[2] VSI information available at "VLBI Standard Interface Specification - VSI-H," August 2000, Revision 1.0, available at *http://web.haystack.edu/vsi/index.html.*

[3] Mark 5 information available at *http://web.haystack.edu/mark5/Mark5.htm.*

[4] DRAGON project details at *http://dragon.maxgigapop.net/twiki/bin/view/DRAGON/webHome*