e-VLBI: Creating a Global Radio-Telescope Array via High-Speed Networks

A. Whitney, Y. Koyama, S. Parsley, T. Tzioumis

[(1)] *MIT Haystack Observatory, Westford, Massachusetts 01886 USA <awhitney@haystack.mit.edu>*
[(2)] *NICT KSRC, 893-1 Hirai, Kashima, Ibaraki 314-8501, Japan <koyama@nict.go.jp>*
[(3)] *JIVE, Oude Hoogeveensedijk 4, 7991 PD Dwingeloo,The Netherlands,<stephen.parsley@astro.cf.ac.uk>*
[(4)] *CSIRO/ATNF, PO Box 76, Epping NSW 1710, Australia <tasso.tzioumis@csiro.au>*

ABSTRACT

Very Long Baseline Interferometry (VLBI) has traditionally been data-rate limited by data recording systems. The recent advent and continuing advancement of global high-speed optical networks provides the potential to transmit VLBI data directly to the correlator, dubbed 'e-VLBI', with higher data rates and faster processing turnaround than has ever been possible before. Within the past ~3 years, the pace of e-VLBI development has dramatically quickened. Among the recent accomplishments:

– Installation of high-speed network links to several correlators and a number of antennas worldwide
– International real-time e-VLBI experiments using antennas in U.S. and Europe
– Real-time e-VLBI demonstrations over shared networks at data rates to 512 Mbps
– Demonstration of a few-hour turnaround for creation of astronomical images and geodetic results
– Establishment of routine non-real-time transmission for some observations

INTRODUCTION

e-VLBI development has been rapid over the past few years. High-speed international networks, coupled with low-cost off-the-shelf equipment, have accelerated the pace of development and e-VLBI data transmission is becoming increasingly common and, in some cases, routine. The real promise of e-VLBI, however, lies in the potential for observing data rates which exceed those that can be economically recorded, to further increase the sensitivity of continuum VLBI observations. Gigabit/sec observations are routine today onto magnetic disk systems, but data rates in the range of 10Gbps to 100Gbps may be possible using e-VLBI in the future.

In this paper, we will briefly discuss the promise and challenges of e-VLBI, as well as the direction of research and implementation efforts on four continents.

e-VLBI CHALLENGES

Though, in principle, e-VLBI over high-speed networks may seem straightforward, the list of challenges confronting routine widespread global e-VLBI is significant, including:

- Poor 'last-mile' connectivity to telescopes
- End-to-end network performance that is often well below advertised rates
- Poor performance of transport protocols, particularly over long distances or with data loss
- Throughput limitations of COTS hardware
- Adding another layer of complexity to already complex business of VLBI
- Time-varying nature of network; R&E networks tend to be unstable, particularly towards end points

Many of these issues are now being addressed by e-VLBI research and development groups around the world.

e-VLBI STANDARDIZATION

Since VLBI experiments are conducted with heterogeneous sets of data-acquisition equipment in different parts of the world, one of the major challenges for e-VLBI is to define e-VLBI data format and protocol standards that will allow these systems to be interoperable. In 1999 a VLBI Technical Working Group (TWG) was established to develop interface standards for VLBI. The TWG was created as a joint effort under the sponsorship of the Global VLBI Working Group, representing astronomy interests, and IVS, representing geodetic interests. The charge to the TWG was to define a standard interface to and from a VLBI Data Transmission System (DTS) that allows heterogeneous DTSs to be interfaced to both data-acquisition and correlator systems with a minimum of effort. The result of this effort is the VLBI Standard Interface (VSI) specification, which is now being widely adopted [http://web.haystack.edu/vsi/index.html].

The VSI specification was originally confined to the local user hardware and software interfaces of the DTS, resulting in the VSI-H and VSI-S specifications, respectively. With the advent of e-VLBI, however, it became clear that a specification for electronic data transmission was necessary to insure efficient electronic transfer of VLBI data between the heterogeneous DTSs; thus was born the VSI-E specification. Though the VSI-E specification is not yet complete, significant progress has been made. The well known and mature RTP protocol has been chosen as the basis for VSI-E. Among the reasons are: 1) RTP allows the reuse of many standard monitoring and analysis tools, 2) RTP is seen as internet-friendly by the network community, 3) RTP is designed to have minimum overhead for in-band data, 4) RTP is easily scalable, and 5) RTP allows a choice of underlying transmission protocol (TCP, UDP, UDT, FAST, etc)

In addition, the RTP data transmission is controlled by an independent low-data-rate Real-Time Control Protocol (RTCP) stream that provides many additional features useful to e-VLBI, such as: 1) dissemination of session information, 2) monitoring of network and end system performance, 3) adaptation to varying network capability and performance, 4) message sequencing and reordering, 5) multi-cast support for distribution of statistics and eventual possible implementation of distributed correlator systems

At this writing, a reference VSI-E implementation has been created by Haystack Observatory and is now being evaluated by VLBI institutions worldwide. The widespread implementation of VSI-E should be complete within the next year.

INTERNATIONAL e-VLBI DEVELOPMENT PROGRAMS

e-VLBI development programs are currently active in the U.S., Japan, Europe and Australia. We will briefly summarize the activities in each of these areas.

e-VLBI Development in U.S.

Haystack Observatory is engaged in a multi-faceted e-VLBI development program. Many e-VLBI experiments have been carried out over the past 3 years. In October 2002, a near-real-time e-VLBI experiment was conducted at 788Mbps using the Westford (MA) and GGAO (MD) antennas. Since that time, many national and international experiments have been conducted, including full real-time (i.e. no disk buffering between antenna and correlator) at 512 Mbps between U.S. antennas. Routine non-real-time e-VLBI transfers from Japan to Haystack now take place on a weekly basis. Connectivity of Haystack to national and international networks is currently 2.4Gbps, with the expectation that it will be increased further to 10Gbps within the next year.

Work continues to develop intelligent applications for e-VLBI that take advantage of unique characteristics of VLBI data, for instance, the ability of VLBI to absorb some loss of data, to maximize data transfer rates through shared networks. The challenge is to structure e-VLBI network usage so that it can efficiently make use of excess network capacity without interfering with 'normal' users. The 'Experiment Guided Adaptive Endpoint' (EGAE) is designed to meet this challenge. With funding from the U.S. National Science Foundation, Haystack Observatory is developing the EGAE strategy to makes use of "scavenged," low priority bandwidth to transport e-VLBI data. A simple prototype EGAE implementation is now being used at Haystack Observatory for routine weekly non-real-time production e-VLBI data transfers of ~300GB each.

Sometimes painful experience gained over the past few years have made it clear that end-to-end performance over R&E networks is sometimes not stable. Indeed, it is not possible to guarantee that network conditions will stay the same from one day to the next or one week to the next. As a result of this experience, the need for continuous or semi-continuous network monitoring has become increasingly important. For this purpose, an automated periodic performance monitoring system and accompanying 'Network State DataBase' (NSDB) has been constructed that allows instant web access and easy interpretation of all the data collected by the monitoring system. With a few simple clicks, the detailed performance statistics of every monitored link can be viewed, aiding greatly in diagnosis and restoration of services. This type of monitoring needs to be extended to a truly global basis if e-VLBI is to become a routine procedure.

Haystack Observatory is collaborating with the University of Maryland and others in the DRAGON (Dynamic Resource Allocation via GMPLS Optical Networks) project, under sponsorship of the National Science Foundation, to create dynamic and deterministic end-to-end network transport services for high-end e-Science applications. Using standard dense wave-division multiplex (DWDM) technology and optical switches capable of switching individual wavelengths among an array of fibers, dedicated wavelength paths can be constructed and de-constructed dynamically on user demand. 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, this work will be fully integrated with the EGAE project described above.

e-VLBI Development in Japan

In Japan, the first operational real-time VLBI network was realized in the Key Stone Project (KSP) of Communications Research Laboratory which is now National Institute of Information and Communications Technology (NICT). Four 11-m VLBI stations were constructed around Tokyo to monitor the crustal deformation of the area by daily or sub-daily geodetic VLBI sessions. The regular observations with the four stations started in September 1996 with the tape-based recording systems. Later, the four stations were connected by dedicated OC-48 fibers with cooperation of the Telecommunication Network Laboratory Group of Nippon Telegraph and Telephone Corporation. In June 1997, the observations were switched from tape-based recording systems to the real-time system where the observed data at each station with the data rate of 256Mbps were transferred using Asynchronous Transfer Mode protocol to the real-time correlator at Koganei station. Real-time VLBI observations were conducted daily or sub-daily until December 2001. With the KSP VLBI system, various possibilities of the real-time VLBI were demonstrated. Continuous or quite frequent VLBI observations without human operations from the observations to the data processing were made possible. Especially, the usefulness and the power of the automated real-time VLBI system was clearly demonstrated by the dense measurements during the dramatic crustal deformation event associated with volcanic activities of the Miyake-jima volcano started in June 2000.

The enhanced real-time VLBI system has been developed by National Astronomical Observatory (NAO) with cooperation with other Japanese research institutes. The system is capable to correlate single channel 1Gbps data stream and supports maximum data rate of 2Gbps at each station by using two data channels. Currently, 64m station at Usuda, 45m station at Nobeyama, 32m station at Tsukuba, 11m station at Gifu, and 34m station at Kashima are connected by high speed research networks such as Super-SINET, Galaxy, and JGNII.

In NICT, developments of the IP based e-VLBI system started after the developments of the KSP system. The system, named K5 VLBI system, utilizes the conventional PC systems running on the FreeBSD or Linux operating system and newly developed PCI interface boards. The observed data at VLBI observing stations can be recorded as usual data files in the local hard-disks of the PC systems. The developments of the software programs which enables the real-time data transfer from the observing stations to the correlation site over the IP-based high speed network is also underway. Software correlation program has also been developed to enable distributed correlation for the real-time or near-real-time e-VLBI operations. The K5 system was used along with the Mark-5 system to demonstrate rapid turn around of the UT1 estimation using the Kashima-Westford baseline in June 2004. From this session, it was succeeded to estimate UT1-UTC within about 4.5 hours after the one-hour session.

e-VLBI Development in Europe

Current e-VLBI research in Europe is centered on a three-year, proof-of-concept project that aims to connect up to six telescopes to the JIVE data processor in real-time. The first tangible result of this effort was the installation, by SURFnet, of a DWDM link between JIVE/ASTRON and the Amsterdam Internet Exchange in September 2002. Within a few days of first light on this line, iGRID2002 provided an opportunity to showcase e-VLBI. Recorded data were transferred from the UK to JIVE and correlated to form a single baseline interferometer between the Jodrell Bank and Westerbork telescopes. The feasibility of performing e-VLBI using the European production R&E networks was thus established.

One of the key targets of this project is to pursue real-time e-VLBI, the rationale being that the network problems only become interesting when large data flows from multiple sources converge on a single location. To achieve this, participating telescopes need high bandwidth connections to their local NRENs. So far three European telescopes are connected internationally at a full 1Gbps: Onsala (Sweden), Torun (Poland) and Jodrell Bank (UK). Cambridge (UK) and Medicina (IT) should soon be connected at 1Gbps, as well as several other EVN telescope in Europe, China and South Africa. Data from these stations are routed to GÉANT via their respective NREN and, via SURFnet, onward to JIVE in the Netherlands. Data from Westerbork, close to JIVE, arrive via a dedicated 1Gbps fiber link. Arecibo, in Puerto Rico, has a 155Mb/s connection and can be used for data rates of 64Mb/s or less.

A number of successful e-VLBI demonstrations have been conducted. In September 2002, non-real-time e-VLBI data transfers at 600 Mbps were demonstrated. Starting in May 2003, ftp e-VLBI data transfers followed by software correlation have been used on a regular basis to verify station operation prior to scheduled EVN observing sessions. Real-time correlation without disk buffering has been demonstrated at data rates of up to 256Mb/s. In March 2005 the number of EVN stations participating simultaneously was increased to 6, creating a fully functioning 6-station e-VLBI network with real-time correlation at JIVE.

The DWDM line connecting the JIVE data processor to the European networks has capacity for virtually unlimited expansion. The six-telescope array created has begun a series of e-VLBI science demonstrations and, within the EVN; e-VLBI will soon achieve the status of an available observing mode.

In March 2005 a proposal called "EXPReS" was approved by the European Commission Research Infrastructure Call. "EXPReS" will collect into one coordinated project all actions and developments needed to create a real-time e-VLBI network, transparently connecting together an array of up to 16 telescopes operating reliably, at data rates of up to 1Gbps per telescope. The project will also include research activity that will look forward to the next generation e-VLBI system, enabled by evolving network capacity and emerging GRID technologies.

e-VLBI Development in Australia

Australia’s Long Baseline Array (LBA) consists of ATNF telescopes at Narrabri (ATCA), Mopra and Parkes, NASA telescopes at Tidbinbilla, and University of Tasmania telescopes at Hobart and Ceduna. Australia’s location in the southern hemisphere adds special value to global VLBI experiments, and is thus an important area of the world to connect with e-VLBI.

The LBA has recently been completely outfitted with PCEVN disk-based recorders based on the VSI-B digital input card. Each PCEVN recorder will record at 256 Mbps, capturing two 16 MHz bands for each of 2 IFs. Parkes, Mopra and ATCA are each outfitted with 2 PCEVN systems for a total maximum date rate of 512 Mbps. The data are recorded to normal Linux file systems. Real-time fringe testing is routinely done by the LBA by transferring small amounts of this data (~0.5-1.0 seconds) via ftp to Swinburne University for software correlation. In this way, test fringes are usually achieved within a couple of minutes. Experimentation with full real-time correlation using software correlators is also planned. In the longer term, a new hardware correlator is being developed for ATCA, which should be able to hand data rates up to 64Gbps.

In January 2005, four LBA stations collected data for the Huygen’s probe wind experiment. As the Huygen’s probe descended into the atmosphere of Saturn’s moon Titan, data were collected at 512 Mbps at these stations. These data were then transferred via high-speed ftp to JIVE in The Netherlands for correlation.

Work is presently in progress to connect a number of the LBA telescopes via high-speed optical fibers using the AARNet3 network as the backbone. Fiber tails from the AARNet3 network to the antenna sites are planned for 2005. Initial connectivity will be at 1Gbps per station, with future plans to extend to 10Gbps per station, possibly using lambda switching. Australia is well connected to the U.S. and Japan, and hence to the rest of world, at speeds in excess of 10Gbps, giving Australia excellent global connectivity

Summary

e-VLBI has the potential to dramatically enhance the power of VLBI, particularly as global networks become capable of handling data streams of many Gbps from each station, pushing VLBI sensitivities to levels not possible with traditional recording technology. VLBI technology is moving rapidly to take advantage of this new capability, and the nations of the world are working and cooperating to bring these new capabilities to reality in the near future.