The Mark 5 VLBI Data System and e-VLBI Development

Alan R. Whitney
MIT Haystack Observatory
Westford, MA, USA
awhitney@haystack.mit.edu

Abstract

The 1 Gbps Mark 5 VLBI data system, based on magnetic-disk technology, is being developed at MIT Haystack Observatory as the next-generation VLBI. Incorporating primarily low-cost PC-based components, the Mark 5 system supports data rates up to 1024 Mbps recording to an array of inexpensive removable IDE/ATA disks. Prototype Mark 5 units are now in routine use at several stations and correlators, with ~25 units operating as of early 2003. Mark 5A units are now available from a commercial source.

e-VLBI, the electronic transmission of VLBI data, is coming of age more rapidly that anyone of us might have predicted only a few years ago. High-speed international networks, coupled with low-cost off-the-shelf equipment, make for a potentially very useful tool for VLBI. We will discuss the work at Haystack Observatory aimed at making global e-VLBI a practical reality. In addition, we will describe a new research project being undertaken in collaboration with the MIT Laboratory for Computer Science to develop a new network protocol tailored to using e-VLBI on shared networks; this work is intended to make good use of large amounts of 'secondary' bandwidth available on almost all shared networks, but in a way that has minimal interference with 'primary' users.

Mark 5 VLBI Data System

The Mark 5 system is being developed as the first Gbps VLBI data system based on magnetic-disk technology. Incorporating primarily low-cost PC-based components, the Mark 5 system supports data rates up to 1024 Mbps, recording to an array of inexpensive removable IDE/ATA disks. It is expected that disk-based VLBI systems will replace current magnetic-tape systems over the next few years.

The goals of the Mark 5 system are:

- Low cost
- Based primarily on unmodified COTS components
- Modular, easily upgradeable
- Robust operation, low maintenance cost
- Easy transportability
- Conformance to VSI specification
- Compatibility with existing VLBI systems during transition
- Flexibility to support e-VLBI
- Minimum of 1 Gbps data rate
- 24-hour unattended operation at 1 Gbps

All but the last goal are clearly achievable with today's technology; 24-hour unattended operation at 1 Gbps is expected to arrive naturally within ~2-3 years with continued development in disk technology.

<u>Why Disks?</u>

Though both magnetic-disk technology and magnetic-tape technology have made great strides over the past few years, the pace of magnetic-disk development has been no less than spectacular, far exceeding even disk-industry projections. Figure 1 displays a comparison of disk and tape prices over the past several years, showing that disk prices (on a $/GB basis) continue downward in a still-accelerating trend. Current (early 2003) consumer IDE disk costs are ~$1.25US/GB and falling; current Mark4/VLBA tape prices are ~$2US/GB and remaining steady. By ~2005-2006, industry projections suggest the price of disks will fall to ~$0.5/GB. Similarly, current single-disk capacities are ~200 GB and rising; by ~2005-2006, single-disk capacities are expected to rise to 500-1000 GB! A single Mark 5 system with sixteen 700 GB disk drives will record continuously 1024 Mbps for 24-hours unattended.

In addition to falling prices and increasing capacity, disks have several other advantages:

- Readily available inexpensive consumer product;
  continually improving in price/performance with standard electrical interface
- Self contained drive mechanism, so host system can be inexpensive
- Technology improvements independent of electrical interface
- Rapid random access to any data
- Essentially instant synchronization on playback to correlator (no media-wasting early starts needed)
- No headstacks to wear out or replace – ever!

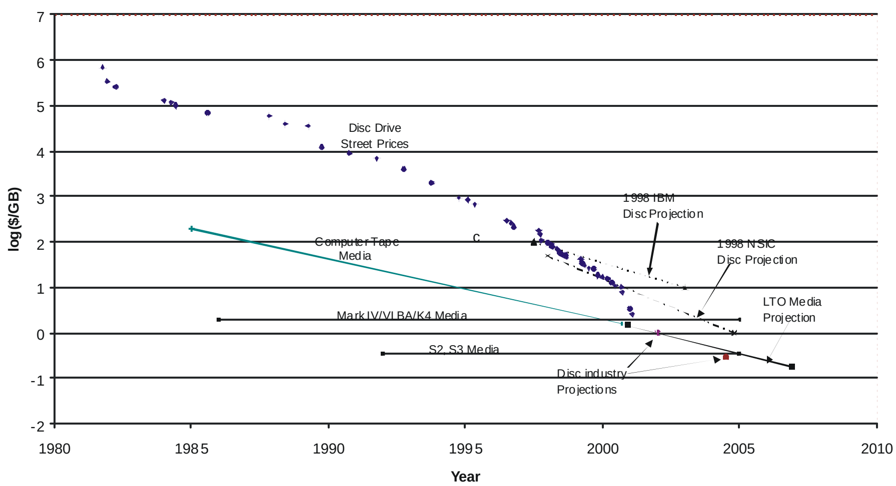

Figure 1: Disk and tape prices vs. time

<u>Mark 5 Development Program</u>

Based on the success of a 512 Mbps Mark 5 demonstration unit [Ref 1] in early 2001, Haystack Observatory is developing an operational 1 Gbps Mark 5 system with support from BKG, KVN, MPI, NASA, JIVE, NRAO and USNO.

The Mark 5 system is being developed in two stages:

1. <u>Mark 5A</u>: The Mark 5A system is intended as a direct replacement for a Mark 4 or VLBA magnetic-tape transport at either a station or a correlators. It records 8, 16, 32 or 64 tracks from a Mark4/VLBA formatter, and plays back in the same Mark4/VLBA format. As such, the Mark 5A is a direct replacement for a Mark4 tape unit at 1024 Mbps and VLBA tape unit at 512 Mbps. The Mark 5A system is in operation at several antennas and correlators, with ~25 Mark 5A prototype systems expected to be deployed by early 2003.
2. <u>Mark 5B</u>: The Mark 5B is VSI-compliant [Refs 2, 3] system with capability up to 1024 Mbps; no external formatter is necessary. The system will also support several backwards-compatibility modes with existing Mark4/VLBA correlator systems. The Mark 5B is expected to be deployed in late 2003.

Figure 2 shows a photograph of the prototype Mark 5A system with its two removable '8-pack' disk modules. A Mark 5A system may be upgraded to a Mark 5B system simply by replacing a Input/Output PCI board in the host PC.

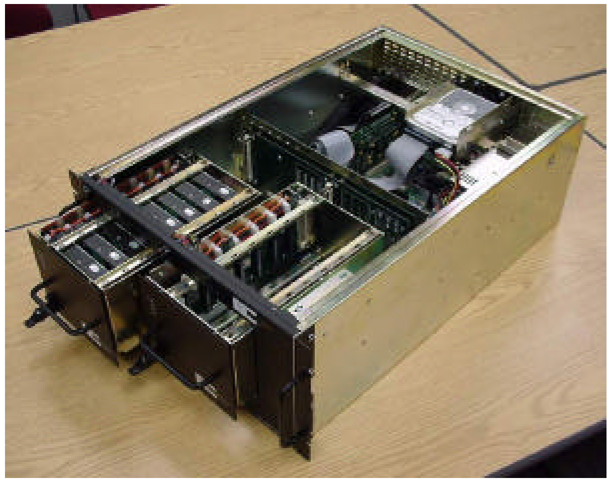

Figure 2: Prototype Mark 5A VLBI data system

Triangle of Connectivity

The Mark 5 is based on a standard PC platform and using a combination of COTS and custom-designed interface cards.

The heart of the system is a 'StreamStor' disk interface card from Conduant Corporation that is specially designed for high-speed real-time data-collection and playback. The StreamStor card supports three physical interfaces in a 'triangle of connectivity' as shown is Figure 3:

1. Data Port/FPDP: This port is present as a 32-bit card-top bus which supports the industry-standard 'Front-Panel Data Port' interface specification [Ref 4]. This is a two-way port through which high-speed real-time data may be either input or output. All 32-bits of the FPDP bus are always active.
2. Disk array: This port supports up to 16 standard IDE disks for reading or writing, arranged in 8 master/slave pairs.
3. PCI bus: This is the standard connection to the host PC platform; the StreamStor card supports a 64-bit/66MHz bus, though it is backwards compatible with standard 32-bit/33MHz buses.

The 'triangle of connectivity' shows that data may be moved in either direction between any two of the three ports. The StreamStor card supports a maximum sustained data transfer rate of up to ~1200 Mbps between any two ports, though only one connection path may be exercised at a time and the maximum data rate for VLBI usage is anticipated to be ~1024 Mbps. The path exercised for traditional VLBI observations is between the FPDP bus and the disk array; note that in this mode, the VLBI data never touch the PCI bus, so the speed of the PC platform is largely irrelevant. Of course, the path between the disk array and the PCI bus allows the PC to read and verify VLBI data written to the disks via the FPDP port. The direct connection of FPDP bus to PCI bus will be used in upcoming e-VLBI experiments where data are transferred directly to a high-speed network and are not recorded locally. An on-board 512 MB buffer provides the necessary 'elasticity' between the three connection nodes to support full real-time operation.

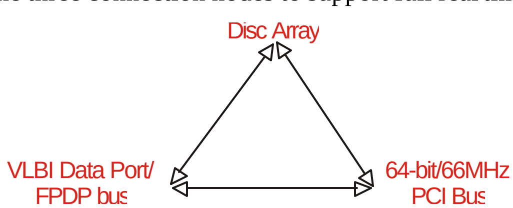

Figure 3: 'Triangle of connectivity' of StreamStor interface card

Mark 5A

The Mark 5A system is designed to be a *direct plug-compatible replacement for a Mark 4 or a VLBA tape recorder at either a field station or a correlator*, except that:

1. The Mark 5A can record and playback 8, 16, 32 or 64 'tracks' of data from a Mark 4 or VLBA formatter[1].
2. The software control of the Mark 5A is based on VSI-S rules and syntax, and is vastly simpler than controlling a Mark 4/VLBA tape transport.

The I/O Panel on the rear of the Mark 5A is shown in Figure 4. It contains the same set of connectors as a Mark 4 or VLBA tape drive so that the Mark 5A is very easily substituted for a tape drive.

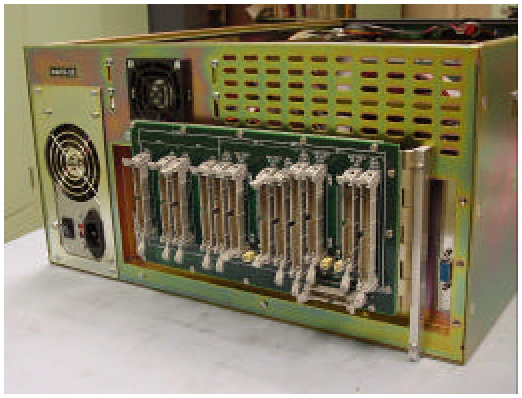

Figure 4: Mark 5A I/O panel at rear of unit

The Mark 4 formatter can output up to 64 tracks of data at 18 Mbps/track, which includes the parity overhead, for a total of 1152 Mbps. However, it is not necessary to record the parity bits in the low-error-rate environment of the disks; therefore, the Mark 5A Input/Output Interface Board does do some special processing to remove the parity bits for recording and restore them on playback. A simplified block diagram is shown in Figure 5.

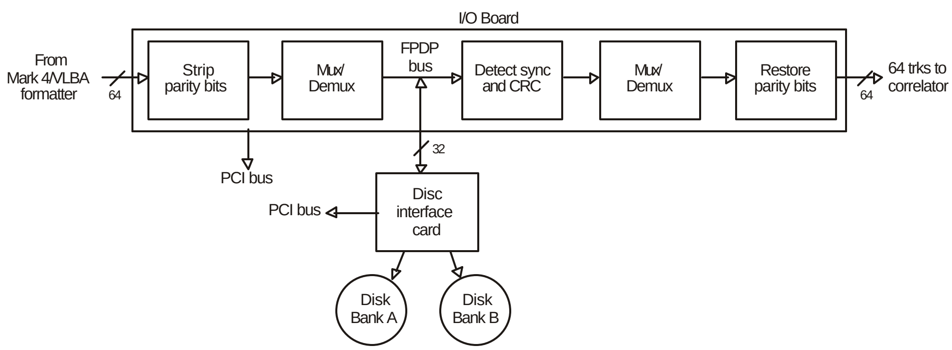

Figure 5: Mark 5A simplified block diagram

Input Section of I/O board (in block diagram in Figure 5):

1. All parity bits are removed.
2. In the 8 and 16-track modes, the data are demultiplexed by a factor of 4 or 2, respectively, and the resulting 32 parallel bit streams are sent directly to the FPDP bus for recording.
3. In 32-track mode (equivalent to Mark 4/VLBA headstack 1 only), the resulting 32 parallel bit streams are sent directly to the FPDP bus for recording.
4. In 64-track mode (equivalent to Mark 4/VLBA headstack 1 and 2), adjacent even and odd track-pairs are interleaved bit-by-bit before being sent to the FPDP bus. Maximum data rate at this point will be 1024 Mbps for Mark 4 and 512 Mbps for VLBA[2].

---

[1] When recording 8, 16, 32 or 64 'tracks', 100% disk utilization is achieved; other numbers of tracks can be recorded, but with <100% disk utilization. Both Mark 4 and VLBA formatters can be configured to multiplex a single channel to 1, 2 or 4 tracks with 1-bit samples or to 2, 4 or 8 tracks with 2-bit samples, so that normally disk utilization is near 100%.

For compatibility with standard VLBA and Mark 4 formatter modes [Ref 5], 8-track mode records VLBA-equivalent even track numbers in the range 2-16 inclusive. 16-track mode records VLBA-equivalent even track numbers in the range 2-32 inclusive. 32-track mode records VLBA-equivalent tracks 2-33 inclusive from headstack 1. 64-track mode records tracks 2-33 from both headstacks 1 and 2.

The Input Section also includes a 'straight-through' mode in which 32 input bit streams are recorded directly on the disk with no processing or multiplexing; this is a useful test mode and may also be used for other applications.

The Input Section also can generate several test patterns for system testing and diagnosis, including a standard VSI-H test pattern.

Output Section of I/O board (block diagram in Figure 5):

The Output Section must reverse the actions of the Input Section to exactly reproduce the original 'track' data streams. In all but the 'straight-through' mode, the board must first undo any multiplexing/de-multipexing to reconstruct parity-free 'track' data streams. Embedded Mark4/VLBA sync words and CRC characters are then detected to determine data framing, parity is restored, and the fully reconstructed track data are sent to the output.

For normal VLBI operation, the Output Section is instructed whether the data are VLBI or 'straight-through'. If VLBI mode, the board automatically recognizes whether data is in 8, 16, 32 or 64-track mode and reconstructs track data accordingly.

Data Clocking

The Input Section is driven by data and clock from the Mark 4 or VLBA formatter at a maximum frequency of 18 MHz. In turn, the Input Section drives the FPDP bus with data and clock at an average frequency 8/9 (32-track case) or 16/9 (64-track case) times the formatter clock frequency.
The Output Section has two clock modes:

1. In Record/Bypass mode, the Input Section passes the 'track' clock over the FPDP bus to the Output Section.
2. In Playback mode, the Output Section may clock the output data either according to an on-board digitally-synthesized clock (0-40MHz, ~23mHz resolution) or from an external user-supplied clock in the range 0-40 MHz. The playback data rate is completely independent of the record data rate, though the normal maximum 'track' playback data rate is ~18 Mbps/track for compatibility with existing correlators. The Output Section will read disk data at whatever rate is necessary to sustain the requested 'track' data rate.

Data format on Disk

The format of the data recorded to disk is such that barrel-rolling and/or data modulation are not necessary. The StreamStor card accumulates 32-bit 'words' (corresponding to the 32-bit wide FPDP bus) into 64 kB 'chunks and writes these chunks sequentially in a 'round-robin' fashion to the disk array. Neither 'barrel-rolling' or 'data-modulation' need to be applied by the formatter since these functions are specifically designed to overcome limitations with magnetic tape.

Operating modes

For normal operation with disks, the Mark 5 system has three basic operating modes:

1. Idle – When the system is 'idle' (i.e. not recording or playing back), the data entering the Input Section passes through the FPDP bus and is re-created (with a delay of a few clocks cycles) at the Output Section. This is also known as 'pass-through' mode and allows multiple Mark 5 systems to be daisy-chained together to increase overall data capacity, as shown in Figure 7..
2. Record – This mode is identical to 'idle' mode' except that the disks record the data on the FPDP bus.

---

[2] Actual maximum data rate for VLBA is slightly higher (516.096 Mbps to be exact) due to non-data-replacement nature of the format.

3. Playback – Data pre-recorded on disk is played back to the FPDP bus and then to the Output Section. In this mode, the playback rate is controlled by a clock provided by the Output Section (either on-board or external) and is independent of the record clock rate.

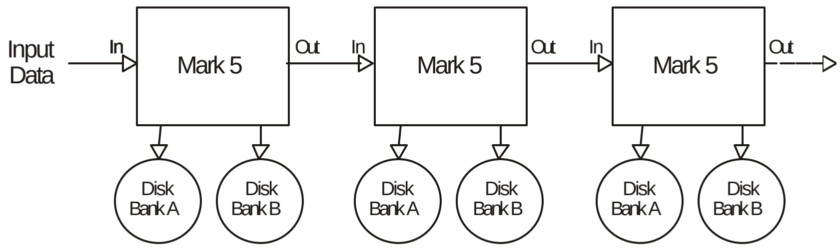

Figure 6: Recording with cascaded Mark 5 systems

When recording, the Mark 5 *always* adds a new scan (recording) after the end of the previous scan, just like tape. No deletion or erasure of individual scans is allowed; an 'erase' command deletes all scans.

Additional features

Several additional features exist in the Mark 5 system to ensure data integrity and continuity:

1. The Mark 5A chassis holds two '8-pack' modules of disks in two slots in the Mark 5A chasses, labeled as 'Disk Bank A' and 'Disk Bank B'. Each module can support up to 1024 Mbps and each module can operate independently (but not simultaneously). A 'module switching' mode is being implemented so that recording is automatically switched from one module to another when the first is filled; this allows continuous data to be taken over module boundaries. Modules are 'hot-swappable'.
2. If a disk becomes slow or fails during recording, the recording load is dynamically adjusted among the disks so that no data are lost.
3. If a disk fails or is missing during playback, the Mark 5 will fill any data gaps with a user-specified data pattern that can be detected and cause the data to be invalidated at the correlators.
4. '8-pack' disk modules can be managed just like tapes with little or no modification to present management procedures. Each '8-pack' disk module contains a permanent electronically-readable Volume Serial Number (VSN) just like Mark 4/VLBA tape. Current tape labeling schemes can be directly transferred to 8-pack disk modules.

Software Control

Control of the Mark 5A system is through a software program on the host PC machine operating under a Linux OS. The control syntax and rules are taken from VSI-S, though the commands are mostly not standard VSI-S commands. Control be from a local or remote terminal or over programmatically over a network. Though there are many modes and many test and status commands and queries, the fundamental operation of the Mark 5A is quite simple. The Mark 5A system is supported by the NASA/GSFC Field System for control at stations.

Deployment and Availability

As of early 2003 approximately 25 Mark 5 systems are operating at various stations and correlators in the world. Haystack Observatory has transferred the Mark 5 technology to Conduant Corporation (Longmont, CO) and may be purchased for ~$17,500 each without 8-pack modules or disks.

Mark 5B

The Mark 5B will be fully VSI-H and VSI-S compatible, allowing more compatibility among various VSI-compatible VLBI data systems. A new Mark 5B Input/Output Board will be designed to replace the Mark 5A Input/Output Board. Among the features of Mark 5B will be:

1. Internal data formatting eliminates need for external formatter.
2. Simple input and output interface, each on a single 80-pin connector.
3. 32-bit wide input and output data channel.

4. Any 1, 2, 4, 8, 16 or 32 input bits streams may be selected for recording with 100% disk-capacity efficiency.
5. Easy synchronization at a correlator with an external second tick.

The Mark 5B system is expected to be available in late 2003.

Compatibility Considerations

The Mark 5 system is being designed for extensive forward and backwards compatibility with existing VLBI systems. For example, data may be recorded with a VSI-compatible interface and re-played into any Mark4/VLBA correlator. Conversely, data may be recorded from a Mark4/VLBA system and re-played into any VSI-compatible correlator. In addition, it is expected that existing interfaces to S2 recorders can be easily adapted to record on Mark 5B, which can then be re-played into either a VSI-compatible or Mark4/VLBA correlator.

This inter-compatibility among various systems will allow a much broader and flexible use of existing VLBI facilities throughout the world.

Possible Future Enhancements

Among the future enhancements being studied for Mark 5:
1. Support for interchangeability between 8-pack modules containing parallel-ATA disk drives and 8-pack modules containing serial-ATA disk drives.
2. Support for an inexpensive 'expansion' chassis that will support additional 8-pack disk modules.
3. Full compatibility with standard Linux file systems.

e-VLBI

Electronic transmission of VLBI data from antenna to correlator has been an obvious but difficult goal of VLBI practitioners since the origin of the VLBI technique in the 1960's. A pioneering experiment in 1977 linked the signals from two antennas over a real-time satellite link [Ref 7]. Transmission of small amounts of data (~1 Mb/station) over ordinary telephone lines to a software correlator was successfully accomplished at Haystack in 1979 [Ref 8]. Since that time, the dominant activity in e-VLBI has been in Japan, with the Keystone project in ~1995 [Ref 9] linking four antennas in real-time at 256 Mbps and, more recently, dedicated Gbps networks. Furthermore, high-speed radio links have been used to transmit data from orbiting antennas, notably the TDRSS satellite, in 1986 [Ref 10] and the Japanese dedicated orbiting VLBI satellite HALCA in 1997 [Ref 11]; both of these satellites transmitted data to the ground for recording, but did not transmit data in real-time or near-real-time for correlation.

Until very recently, however, it has not been possible to take significant advantage of the rapidly developing public networks, most notably the World Wide Web, which links nearly all worldwide VLBI antenna sites with network connections of varying speed and quality, but which is rapidly expanding in bandwidth and quality.

Advantages of e-VLBI

The potential advantages for scientific productivity and technical operations of e-VLBI over traditional VLBI are:
1. *Faster turnaround of results*: Typical time to completion of processing of traditional VLBI is delayed until tapes or disks can be shipped, often from distant locations, and is usually measured in at least days and sometimes weeks. This makes it almost impossible to use 'fast-response' observations of important transient events as a means of guiding further critical observations. This is especially important as transient events such as extragalactic supernova or gamma-ray-burst events are becoming increasingly important in understanding various aspects of our universe.
2. *Higher sensitivity*: The potential to extend e-VLBI to multi-Gbps data rates will allow an increase in the sensitivity of observations beyond those possible with traditional recorded media. For most observations, VLBI sensitivity increases as the square root of the data rate; larger antennae and quieter receivers are generally the only other methods to increase sensitivity. However, larger antennae are hugely expensive and many modern receivers are already near theoretical quietness limits.
3. *Lower costs*: e-VLBI will eliminate the need for expensive tape or disk pools while at the same time allowing full automation of observations, all towards the goal of lowering cost.
4. *Quick diagnostics and tests*: Some aspects of VLBI equipment are very difficult to test and diagnose without actually taking data with another station and processing it. Unfortunately, this characteristic has too often led to

bad or poor observations. With e-VLBI, a small amount of data may be transmitted and processed to certify the proper operation of the equipment at each station.

Both faster turnaround and higher sensitivity will open doors to better science while lower costs and easy diagnostics will lead to more science impact per dollar.

e-VLBI using the Mark 5 system

As shown in Figure 3, the Mark 5 system is built on a 'triangle of connectivity' that allows data to move between flexibility between the VLBI data port, the disk array and the PCI bus at data rates of 1 Gbps. When the system is being used for e-VLBI, these paths provide flexibility to support several operating modes:

Full data buffering – data are recorded on disk at each antenna, then transmitted post-observation to a processing center, where they are again recorded on disk. When data from all stations are assembled, correlation takes place in the normal manner. Full data buffering is required whenever the link between station and processing center is significantly slow compared to the recorded data rate.

Direct data transfer – data collected at a station are transmitted over a real-time link to processing center, where they are recorded on disk. Correlation processing is done when data from all stations are assembled. This mode requires fewer disks than full data buffering and may allow real-time correlation of a subset of observing stations for diagnostic purposes.

Full Real-time transfer – data are transmitted to a correlator or processor in real-time. This mode requires that *all* observing stations must have a real-time link to the processing center.

Gbps Demonstration from e-VLBI from Boston to Washington

In mid-2001, Haystack Observatory received support from DARPA to explore the possibility of Gbps e-VLBI demonstration between antennas at MIT Haystack Observatory in Westford, MA and NASA/GSFC in Greenbelt, MD, a distance of ~700 km (see Figure 7).

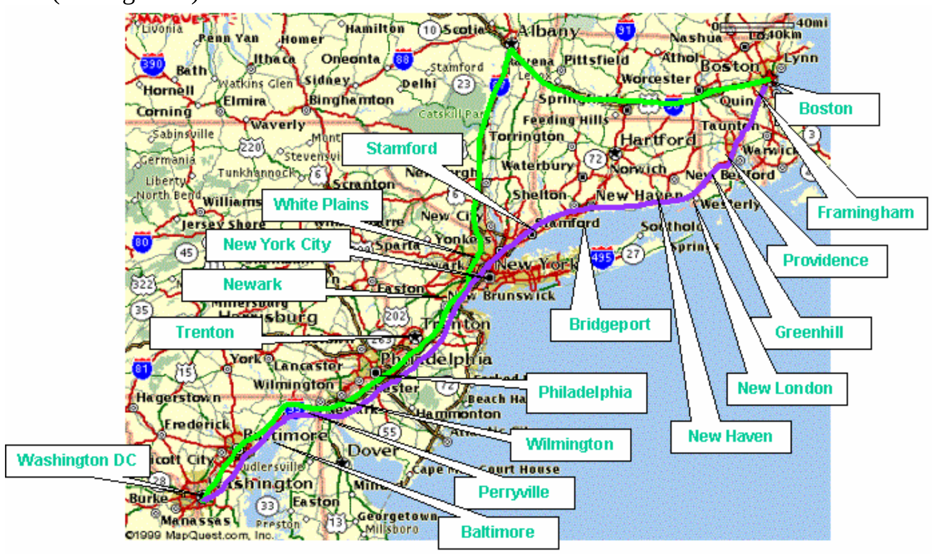

Figure 7: Bossnet route from Boston to Washington, D.C.

The connection from Westford, MA to Washington, D.C. traverses several networks. From Haystack Observatory to MIT Lincoln Laboratory, a local network called GLOWnet was used. The Bossnet link is an dedicated experimental high-data-rate link from MIT Lincoln Laboratory in Lexington, MA to Washington, D.C. The link from Washington, D.C. to NASA/GSFC in Greenbelt, MD, a distance of ~30 km, is primarily over the shared MAX (Mid-Atlantic Crossroads) network.

The goal of the experiment was to establish a reliable ~1 Gbps connection between the antenna at NASA/GSFC and the processing center at Haystack Observatory. The complications in this demonstration experiment lay in the fact that there are many routers and switches of diverse types and models between NASA/GSFC and Haystack Observatory, particularly in the shared segments of the network, and the ensemble had never been used at Gbps rates.

Figure 8 shows the complete path, detailing the many switches and routers in the end-to-end path. An intensive several-month effort had to be undertaken to achieve near-Gbps performance.

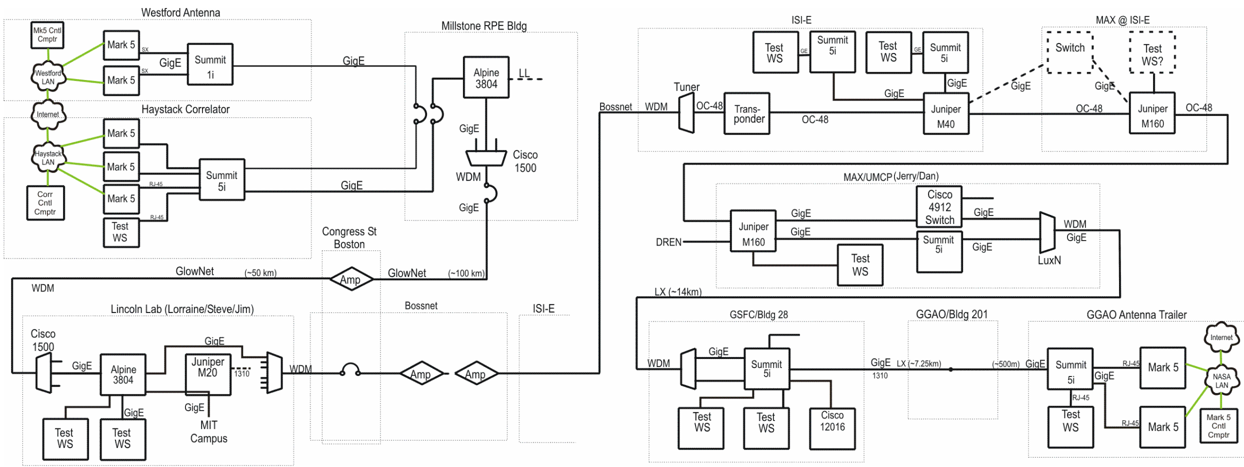

Figure 8: Detailed diagram of Haystack-to-GGAO network

Figure 9 shows measured throughput performance of the Haystack-GSFC link in both directions for a period of ~10 hours. An average sustained data rate of ~960 Mbps was achieved in the direction GSFC to Haystack, while the average sustained rate in the direction Haystack to GSFC was measured at ~900 Mbps. The short periods of lower data rates are due to other traffic on the shared segments of the link.

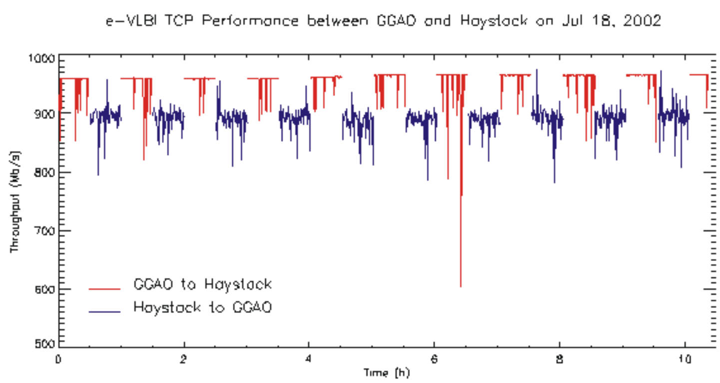

Figure 9: TCP performance measurements over 10-hour period

On 6 October 2002 the first successful high-data-rate e-VLBI experiment was performed using the antenna at Westford (near Haystack Observatory) and the GGAO antenna at NASA/GSFC. Data were simultaneously recorded on disks on two Mark 5P systems at each antenna at 512 Mbps on each system (the maximum rate achievable on a Mark 5P system with a Mark 4 formatter). The GGAO data were then transferred to disks at Haystack Observatory at an average rate of ~788 Mbps; the Westford data were transferred to Haystack over a local GigE link. The data were then correlated using the standard Mark 4 correlator at Haystack Observatory. Fringes were nominal.

On 24 October 2002 another step was taken towards the ultimate goal of full real-time correlator. Data were transferred in real-time from GGAO and recorded at Haystack Observatory on disk. Due to scheduling complications and an imminent and potentially lengthy fiber changeover in the Washington, D.C. area, this experiment had to be done as a piggyback to an ongoing 15-day continuous experiment already in progress at the Westford antenna. A data rate of 288 Mbps was used in the GGAO-to-Haystack link to match the experiment data rate. The data were then correlated with the disk data from Westford. Fringes were nominal in all respects.

Plans are in progress for a full real-time near-Gbps experiment using the Westford and GGAO antennas, but a significant amount of correlator software must be upgraded before this is attempted.

Trans-Pacific e-VLBI between Westford, MA and Kashima, Japan

On 15 October 2002 a successful simple e-VLBI experiment was performed between the Kashima, Japan 34m antenna and the 20m Westford, MA antenna [Ref 12]. X-band and S-band data spread over 14 channels were collected at rate of 256 Mbps on a K5 system at Kashima and a Mark 5 system at Westford. Data were electronically transmitted in both directions and correlated. Fringes were obtained on both the Mark 4 correlator at Haystack Observatory and the PC software correlator at Kashima (see Figure 10). This simple e-VLBI experiment is believed to be the first between the U.S. and Japan.

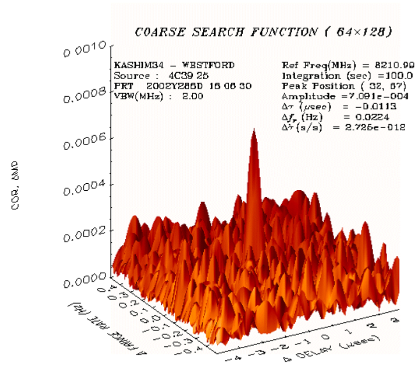

Figure 10: Fringes from first intercontinental e-VLBI experiment (Kashima to Westford)

The network used for this experiment consisted of GLOWnet from Haystack Observatory to MIT Lincoln Laboratory, where GLOWnet was peered to the U.S. Abilene network. In Sunnyvale, CA, Abilene was peered to the GEMnet network, which carried the data to Japan. The slowest link in the network connection is believed to be GEMnet at ~20 Mbps, but actual data transfer rates never exceeded ~2 Mbps for reasons that are still being investigated. In the future, we plan to use the TransPAC network for the trans-Pacific link, which supports two OC-12 connections and should achieve a much higher data rate.

Plans for e-VLBI with Intensive UT1 Observations

VLBI measurements are made almost daily for a period of about 1 hour between Kokee Park, Hawaii and Wettzell, Germany for the purpose of measuring UT1 (earth rotation). Currently these observations are recorded on disk and shipped to the U.S. Naval Observatory in Washington, D.C. for processing. Because these data are important for the *prediction* of UT1, and because the quality of predictions decreases rapidly with longer prediction time, it is important that these observations be processed as quickly as possible. Approximately 100GB of data are recorded in each daily session at each antenna, which is sufficiently small that even a link with a capability of ~100Mbps transfer rate would allow a data transfer to the correlator within a few hours.

Work is currently is progress to make the necessary links to make these connection into reality. Plans are underway to install a high-speed fiber into the correlator at the U.S. Naval Observatory. There is currently no fiber connection available at Wettzell, but the University of Regensburg, about an hour's drive away, has an OC-3 (~150 Mbps) link to which the recorded disks can be transported for data transfer. Plans are also proceeding to connect Kokee Park to the DREN network through facilities at the Pacific Missile Range Facility on the island of Kauaii. The DREN network will carry the data to the Washington, D.C. area where it will enter the MAX (Mid-Atlantic Crossroads) network and thence to USNO. We hope to be able to do first trials with actual observations in early 2003.

Development of Shared-Use Networking Protocols for e-VLBI

e-VLBI has the potential to use a significant amount of the currently unused capacity on existing research network data networks. For example, Figure 11 shows usage statistics for the U.S. high-speed research network Abilene for a typical week. Figure 11a shows that, over a network connection capable of supporting ~600 Mbps, actual average short-term data rates during the week exceed 20 Mbps less than 1% of the time. Figure 11b shows the distribution of 'bulk transfers', defined as transfers >10MB; we see that >90% of these bulk transfers are less that 2 minutes in duration. Clearly, there is a much unused capacity which e-VLBI might utilize. Of course, usage is steadily increasing, but the capacities of networks such as Abilene are also increasing very rapidly. In fact, in order for the 'primary' Abilene user to observe good performance most of the time, it is necessary that a network have extra capacity to gracefully absorb peak loads.

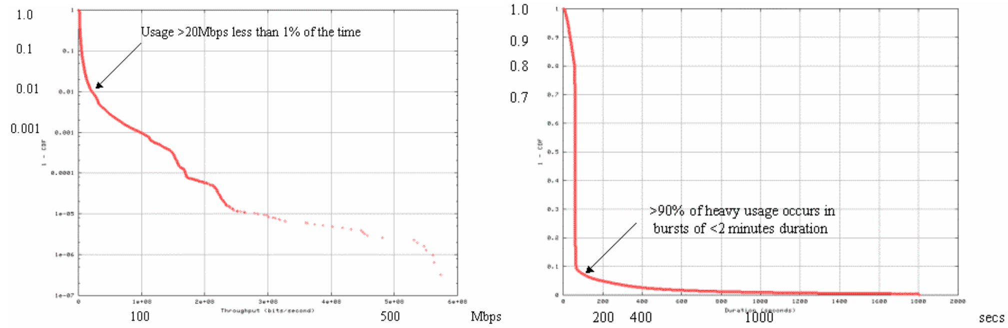

Figure 11a,b: Typical Abilene usage profiles averaged over a typical week in 2002

If e-VLBI is to effectively use a significant fraction of the available bandwidth on shared networks left over from 'primary' users, it must do so in a manner that has little or no impact on 'ordinary' users. In this sense, e-VLBI has the potential to work in the background on the network as a 'secondary' user, scavenging and using available bandwidth, dubbed 'secondary' bandwidth, on an 'as-available' basis, and staying out of the way of 'primary' users.

Haystack Observatory has recently been awarded a grant from the National Science Foundation to develop a new network protocol tailored to the needs of e-VLBI and other similar users. Working in collaboration with the MIT Laboratory for Computer Science, we plan to develop, test and deploy this new protocol for e-VLBI use over the next three years.

Summary

The Mark 5 system takes advantage of low-cost technology developed for the computer industry and applies it to VLBI. Clearly, disk-based recording has many advantages over traditional tape-based recording and is likely to rapidly replace tape-based systems over the next few years. In additional, because the Mark 5 system is based on standard computer technology, it interfaces easily to standard networks to facilitate the further development and use of e-VLBI.

References

Mark 5

1. "Second interim report on COTS-VLBI project" by Alan R. Whitney, 8 Mar 2001, Mark 5 memo 3, available at http://web.haystack.mit.edu/mark5/Mark5.htm

2. "VLBI Standard Hardware Interface Specification – VSI-H", Revision 1.0, 7 August 2000, available at http://dopey.haystack.edu/vsi/index.html.

3. "VLBI Standard Software Interface Specification – VSI-S", under development.

4. FPDP Specification, available at http://www.fpdp.com/.

5. "Mark IIIA/IV/VLBA Tape Formats, Recording Modes and Compatibility - Rev 1.2" by Alan R. Whitney, 28 Sep 2000, Mark 4 memo 230, available at ftp://dopey.haystack.edu/pub/mark4/memos/index.html.

6. The Mark 5 web site at http://web.haystack.mit.edu/mark5/Mark5.htm contains much additional information about the Mark 5 system.

e-VLBI

7. Yen, J. L., K. I. Kellerman, B. Rayher. N. W. Broten, D. N. Fort, S. H. Knowles, W. B. Waltman, and G. W. Swenson, Jr., Real-Time, Very Long Baseline Interferometry Based on the Use of a Communications Satellite, *Science*, **198**, 289-291, 1977.

8. Levine, J. and A. Whitney, Mark III Real-Time Fringe-Detection System, Radio Inteferometry Techniques for Geodesy, *NASA Conference Publication 2115*, 1979.

9. Yoshino, T. et al, Keystone Project and VLBI real-time data processing, *Proceeding of 4th APT Workshop, Sydney, Australia*, 1995.

10. Levy, G. S. and 31 coauthors, VLBI Using a Telescope in Earth Orbit. I. The Observations, *Astrophys. J.*, **336**, 1098-1104, 1989.

11. Hirabayashi, H. and 52 coauthors, Overview and Initial Results of the Very Long Baseline Interferometry Space Observatory Programme, *Science*, **281**, 1825-1829, 1998.

12. K. Yasuhiro, T. Kondo, J. Nakajima, M. Kimura, H. Uose and S. Iwmura, International e-VLBI experiments, *IVS CRL-TDC News No. 21*, 23-26, Nov 2002.