Future Directions in High Resolution Astronomy: The 10th Anniversary of the VLBA
ASP Conference Series, Vol. 340, 2005
J. D. Romney and M. J. Reid (eds.)

Disk-Based VLBI Recording

Alan Whitney

MIT Haystack Observatory, Route 40, Westford, MA 01886

Abstract. Ever since the first VLBI experiments in 1967, magnetic tape has been the mainstay medium for recording VLBI data. However, in late 2000 it became evident that the price of computer disks was dropping far faster than that of computer magnetic tape, and would likely drop below the price of tape within a few years. Additionally, disks have the advantages of relatively high data-rates, random access and easy availability. Used in disk arrays, data rates of 1 Gbps are readily achieved. These realizations have led to the development of several disk-based systems for VLBI data, including the Mark 5 system in the U.S., the K5 system in Japan and the PC-EVN system in Europe. We will examine the current state of disk-based VLBI systems, their expected future, and their close relationship to the developing technique of e-VLBI.

1. Introduction

VLBI recording has traditionally been the domain of magnetic tape. In the beginnings of VLBI in 1967, data were recorded on standard $^1/_2$-inch open-reel tape at 720 kbps for a total of 3 minutes on a single tape! In the early seventies, a succession of Mark II VLBI systems utilized new helical-scan technologies based on TV-recording transports, raising the data rate to 4 Mbps, first on 2 inch-wide open-reel tape and progressing to VHS cassettes. In 1979, the Mark III system raised the bar to 224 Mbps on 1 inch-wide open-reel tape on 14 inch-diameter reels based on tape transports originally designed for analog instrumentation recording, but adapted to the needs of VLBI by adding movable multi-track heads and re-designing all signal electronics. In Japan in the early 90s, the K4 system, operating at 256 Mbps and based on the helical-scan technology developed for digital video, made its appearance. Also in the 90s, the S2 system from Canada, operating at 128 Mbps and based on multiple-parallel video cassette transports, gathered much interest due to its relatively low cost. And in the late 90's, two Gbps systems appeared. The Mark 4, based on the continued evolution of the Mark III and Mark IIIA series of machines, pushed with difficulty to reach its design goal of 1024 Mbps, while in Japan a beautiful, but very expensive, 1-Gbps helical-scan system developed by Toshiba appeared, based on an HDTV video recorder transport. Even as recently as three years ago, there was little serious thought that magnetic tape would be displaced anytime soon as the VLBI media of choice. Optical-recording technology looked promising at several junctures, but has never lived up to its promise. Magnetic-disk technology simply did not seem competitive, primarily due to high cost and relatively low data rates for individual disks. Even now, as magnetic disks are becoming the current media of choice, global high-speed networking is beginning to show

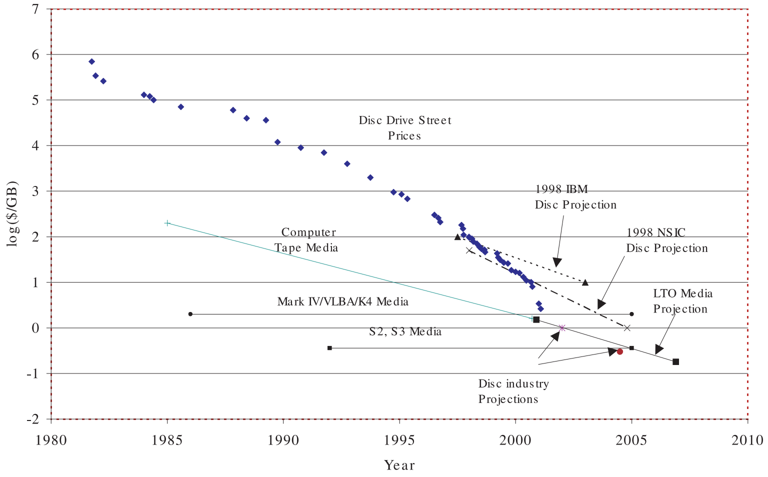

Figure 1. Disk and tape prices vs. time.

promise of superseding disks, though exactly when and at what cost is not yet clear.

2. Why Disks?

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, in terms of cost, capacity and data-rate capability. Figure 1 shows 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 (fall 2003) consumer IDE/ATA disk costs are ~$1/GB and falling; current Mark 4/VLBA tape prices are ~$2/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 ~250 GB and rising; by ~2005-2006, single-disk capacities are expected to rise to ~500-1000 GB! A single disk-based VLBI data system with sixteen 700 GB disk drives would record continuously at 1024 Mbps for 24-hours unattended!

Over the same period of time, sustainable read/write data rates to individual disks have risen dramatically, today averaging in excess of 25-30 MBps over the entire disk recording area and continuing to rise. These increases have been accompanied by continual evolution in reliability, including sophisticated error-correction codes and transparent sparing of bad magnetic sectors to spare pools of good sectors.

For VLBI, disks have several obvious advantages: 1) Readily available inexpensive consumer product with a standard electrical interface, 2) Self contained drive mechanism, so host system can be inexpensive 3) Technology improvements independent of electrical interface, 4)Rapid random access to any data, 5) Essentially instant synchronization on playback to correlator (no media-wasting early starts needed) 6) No headstacks to wear out or replace - ever!

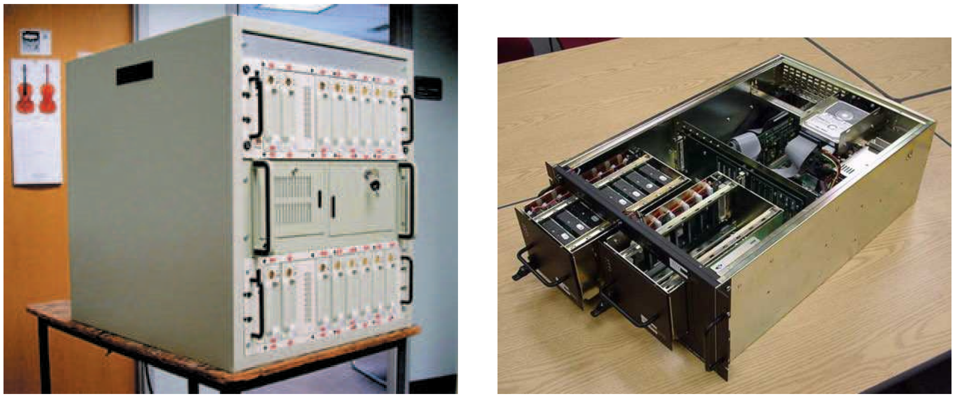

Figure 2. (*left*) Mark 5 demonstration system. (*Right*) Mark 5A VLBI data system.

3. The Mark 5 VLBI System

In early 2001, with a development time of only 3 months, a 512 Mbps disk-based VLBI system (Whitney et al. 2001) was constructed and demonstrated (Figure 2). VLBI data were recorded and processed successfully on the Mark 4 correlator at Haystack Observatory. The rapid speed of development was primarily due to the fact that commercial off-the-shelf (COTS) technology was used entirely throughout the system. Based on the success of the demonstration unit, Haystack Observatory embarked on a program to develop and an operational 1 Gbps disk-based system, dubbed ‘Mark 5’, with support from BKG, KVN, MPI, NASA, JIVE, NRAO and USNO.

The Mark 5 system is being developed in two stages:

1. **Mark 5A:** Direct plug-compatible replacement for a Mark 4 or VLBA magnetic-tape transport; now (fall 2003) operational with ∼40 Mark 5A systems deployed to stations and correlators worldwide.
2. **Mark 5B:** The Mark 5B is a VSI-compliant (Whitney 2000) system with capability up to 1024 Mbps; no external formatter is necessary; expected to be deployed in mid-2004.

Figure 2 shows a photograph of the Mark 5A system. Primary features of the Mark 5A system are: 1) Two removable ‘8-pack’ modules of standard IDE/ATA disks, 2) Primarily based on PC COTS technology 3) Linux OS, 4) Each disk pack can record/playback at 1024 Mbps and can be used in ‘ping-pong’ fashion for nearly continuous recording, 5) Support for e-VLBI data transfer, 6) Commercially from Conduant Corp for ∼\$16K.

4. Strawman Plan for Mark 5 Systems for VLBA

Though funds are not currently available, the VLBA has chosen the Mark 5B as its target system for eventual upgrade. Upgrading the VLBA to Mark 5B technology will require the purchase of ∼36 Mark 5 systems at a cost of ∼\$600K (without media). Outfitting the VLBA with disk-media over a 3-year period would allow a gradual increase from sustained 256 Mbps capability to a sustained 1 Gbps capability and take advantage of the continuing decrease in the cost of disk media. Table 1 shows a strawman plan for purchasing disk media to implement this plan, based on the assumption that 300 station-days of media are necessary for smooth VLBA operation. Note: 300 station-days at 1 Gbps = 3300 TB.

Year	Sustainable Rate (Mbps)	Disk purchase (TB)	Disk cost ($/GB)	Other costs ($/GB)	∼ cost
1	256	825	1	.15	$1.0M
2	512	825	.75	.15	$0.8M
3	1024	1750	.50	.10	$1.0M
Total		3300			$2.8M

Table 1. Strawman plan for purchasing disk media for VLBA

Total cost to implement this plan, including Mark 5 systems and disks, is ∼$3.4M. Some additional relatively modest costs will be incurred for upgrading support software as well as upgrading the correlator interface to VSI. In any case, for a 3-year total cost of <∼$4M, which is ∼5% of the original cost of the VLBA, the average data rate will be raised by a factor of 8, with an average corresponding sensitivity increase of almost a factor of 3!

According to Jon Romney, the current VLBA correlator, with relatively minor upgrades, will be able to keep up with 512 Mbps playback for processing up to 20 stations and 1024 Mbps playback for processing up to 10 stations.

5. Other Disk-Based Systems

Two other disk-based VLBI data systems are known to exist in addition to Mark 5:

5.1. K5

The K5 system (Kondo et. al. 2002), was developed by CRL in Japan and is deployed to several stations in Japan. It is based on 4 Pentium-class PC‘s operating in parallel, each capable of recording four channels of data at 64 Mbps/channel for an aggregate rate of 1024 Mbps. The K5 was primarily developed for use as a large data buffer for e-VLBI and does not have removable disks. A companion software correlator has also been developed at CRL, which has been used in several international e-VLBI experiments.

5.2. PC-EVN

The PC-EVN VLBI data system (Ritakari & Mujunen 2002) was developed at Metsahovi Radio Observatory and has been used in several test VLBI and e-VLBI experiments. It is based primarily on commodity PC components and implements the VSI specification. Data are stored in ordinary Linux disk files. Several PC-based systems operating in parallel are normally required to support data rates above 256 Mbps, similar to the K5 system.

6. e-VLBI: The Next Step

Because the new disk systems are all based on standard PC platforms, they easily support standard network connections necessary to support the electronic transmission of VLBI data, called ‘e-VLBI’. e-VLBI is coming of age more rapidly than anyone predicted only a few years ago. High-speed national and international networks, some now operating at speeds of 10 Gbps or more, coupled with low-cost off-the-shelf networking equipment, make for a potentially very useful tool for VLBI.

Japan has led the way in e-VLBI since the mid-1990’s, with the Keystone project in ∼1995 (Yoshino et al. 1995) linking four antennas in real-time at 256 Mbps and, more recently, dedicated Gbps networks. Until recently, however, it has not been possible to take significant advantage of the rapidly developing high-speed national and international research and development networks. As more telescopes become connected to this global network, e-VLBI will move from a research curiosity to routine global operations.

7. Advantages of e-VLBI

The potential advantages for scientific productivity and technical operations of e-VLBI over traditional VLBI are: 1) Higher sensitivity, 2) Fast turnaround of results, 3) Fully-automated observations, 4) Quick diagnostics and tests.

Note also that even if network connections do not allow full real-time transmission of VLBI data, disk-based systems such as the Mark 5A can act as large temporary buffers while data are transmitted over a slow network.

8. e-VLBI Demonstrations

In October 2002, a ∼1 Gbps e-VLBI experiment was conducted as a feasibility demonstration by MIT Haystack Observatory (Whitney et al. 2002). With support from DARPA and NASA, data were collected on Mark 5A systems at the Westford antenna in Westford, MA and the GGAO antenna at NASA/GSFC in Greenbelt, MD; the Westford-antenna data were transmitted over a local network to Haystack Observatory while the GGAO data was transmitted to Haystack over a combination of dedicated (Bossnet) and shared networks at a average speed of ∼790 Mbps. The data were correlated on the Mark 4 VLBI correlator at Haystack Observatory.

In several experiments from October 2002 through June 2003, antennas at Kashima (Japan) antenna and Westford (MA) were used to conduct e-VLBI experiments. Data at Kashima were collected on the K5 system while data at Westford were collected on the Mark 5A system. Data were exchanged in both directions and correlated at both Kashima and Haystack Observatory, using software conversion of data formats as necessary for correlation. Experiments are continuing, with recent demonstrations of Haystack-Japan connectivity speed of ∼400 Mbps, though the Kashima antenna is currently connected at only ∼100 Mbps.

In Europe, several e-VLBI demonstration experiments have been successfully conducted at data transmission rates to ∼400 Mbps, some using the PC-EVN system for data collection.

Efforts are now underway to connect Kokee, Hawaii and Wettzell, Germany to high-speed networks to facilitate e-VLBI as the routine data-transfer method for daily Intensive UT1 measurements. These experiments are relative low in data volume (∼100GB/station/experiment) and can easily utilize networks with transmission capacities of 20-40 Mbps to send data to a correlator at USNO or Haystack Observatory. By reducing turn-around time from several days to several hours, the value of the results of these observations is significantly increased.

9. New e-VLBI Protocols

e-VLBI has the potential to use a significant amount of the currently unused capacity on existing research networks. High-speed backbone networks, such as Abilene in the U.S., carry an average load that is typically a small fraction of their carrying capacity. However, if e-VLBI is to effectively use a significant fraction of this available bandwidth, 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 1secondary' user, scavenging and using available 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.

10. 'Last-mile' e-VLBI Costs

Unfortunately, most of the world's VLBI telescopes are not well connected to national and international high-speed networks. Furthermore, since many telescopes are purposefully in remote locations, the problem of the 'last-mile' connection to the majority of the world's radio telescopes is a serious obstacle to widespread e-VLBI connectivity.

The cost of outdoor armored fiber of the type that would be required for most 'last-mile' installations is ∼\$0.14/meter/fiber-pair in a multi-pair cable, which translates to ∼\$3500/km for a 48-fiber cable, and is not expected to decrease significantly over time. Installations costs range widely and are also likely to be relatively constant over time:

- Do-it-yourself lay-on-ground installation: ∼\$0.5-1.0K/km
- Plow into ground: ∼\$9-15K/km
- Hang on existing poles: ∼\$10-30K/km
- Urban bury: much greater than \$30K/km

In some cases, microwave or open-air optical links may be less expensive alternatives to fiber links.

The cost of end-terminal equipment has been dropping rapidly. Within just the last two years, for example, the cost of Gigabit Ethernet switches has dropped from ∼\$15K to ∼\$1.2K, and GigE transceivers have dropped from ∼\$750 to ∼\$180. Furthermore, CWDM (coarse-wavelength division multiplexing) transceivers, which allow up to ∼8-16 wavelengths to be multiplexed onto a single fiber with working distances of 50-100km distances, are now in the range of \$400-800/wavelength.

11. Summary

The rapid development of personal-computer technology has enabled the first major shift in VLBI data systems in more than 30 years. Disk-based VLBI data systems based on PC technology are now rapidly displacing magnetic-tape-based systems due to lower entry cost, lower media cost, and significantly improved performance. Furthermore, the natural extension of these systems to support e-VLBI opens yet another avenue for higher bandwidth and more sensitivity, along with near-real-time results, in the foreseeable future.

References

Whitney, A. et al, 'Second interim report on COTS-VLBI project', 8 Mar. 2001, available at `http://web.haystack.mit.edu/mark5/Mark5.htm`.

"VLBI Standard Hardware Interface Specification", Aug. 2000, available at `http://dopey.haystack.edu/vsi/index.html`.

Kondo, T. et al, 'Internet VLBI system developed at CRL', 2002, available at `http://web.haystack.mit.edu/e-vlbi/kondo.pdf`.

Ritakari, J. and A. Mujunen, 'A VSI-H Compatible Recording System for VLBI and e-VLBI', (2002) available at `http://kurp.hut.fi/vlbi/instr/IVS2002.pdf`.

Yoshino, T. et al, 'Keystone Project and VLBI real-time data processing', Proceedings of 4th APT Workshop, Sydney, Australia, 1995. See also `http://ksp.crl.go.jp/index.html`.

Whitney, A. et al, 'The Gbps e-VLBI Demonstration Project', 2002, available at `ftp://web.haystack.edu/pub/e-vlbi/demo_report.pdf`.