International VLBI Service for Geodesy and Astrometry

Working Group 3

VLBI2010

Report of the Subgroup on Data Transport
24 March 2004

Members

Hans Hinteregger
Tetsuro Kondo
Yasuhiro Koyama
Alan Whitney

Abstract

In contrast to the first ~30 years of VLBI development, where highly specialized equipment for VLBI data-acquisition was designed and built at great cost, the last few years are being driven more and more by adapting to rapidly developing technology in the computer and networking industry. This trend is only likely to accelerate, and VLBI must position itself to take maximum advantage of these technologies. Already, the transition from magnetic tapes to magnetic disks has been very rapid, and disks will almost certainly be the mainstay of VLBI data-acquisition for the next few years. However, development of e-VLBI continues to be rapid and will accelerate. Already, international e-VLBI links of more than 500 Mbps have been demonstrated, with speeds in excess of 1-10 Gbps surely achievable in the near future. The advantages of real-time and near-real-time VLBI made possible by e-VLBI are significant, but major potential stumbling blocks are 'last-mile' connectivity for many telescopes and potentially high networking costs. The future of VLBI data-acquisition and transport in the light of current and projected developments will be examined.

Introduction

Since the beginnings of VLBI in the late 1960's, a large segment of VLBI instrumentation development has been directed towards the development of specialized data systems to cater to the extreme demands of both bandwidth and total data quantity that VLBI demands. Until recently, these systems have primarily been focused on magnetic tape as the only practical media. Beginning with modest data rates of <1 Mbps in 1967 recorded on ½" open-reel 2400-ft computer tapes at 800 bpi/track [Ref ?], several generations of specialized magnetic-tape-based data systems have evolved. In the past few years, however, with the rapid advance and broad commercialization of the computer industry, it is increasingly possible take advantage of relatively inexpensive commodity technology as a basis for very high-performance VLBI data systems. An example is the recent development of magnetic-disk-based systems, such as the K5, Mark 5, and PC-EVN systems operating at 1024 Mbps for a tiny fraction of the cost of the tape-based systems they replace.

Electronic transmission of VLBI data, dubbed 'e-VLBI', is not new. As early as the mid-1970's, an experiment using satellite data transmission was successfully executed [Ref ?], but was not economically sustainable. In the late 1970's some VLBI data were transmitted via standard voice telephone lines at 1200 bps and successfully correlated, but this data transport mode was practical only for short tests. In the mid-1990's, dedicated fiber networks spanning limited geographics areas in Japan were successfuly put into operation at data rates of 256 Mbps [Ref ?]. However, only in the last few years, with the advent of a global grid of high-speed fiber connections, has the possibility of transmitting a substantial amount of VLBI data become a practical possibility.

In this paper, we will examine the state-of-the-art in VLBI data systems and try to project to both the demands and the available technology in 2010 and beyond to meet these demands in the most cost effective and efficient manner.

VLBI Data Transport Demands in 2010

For continuing significant improvement in the geodetic VLBI technique, it is clear that several things must happen. Among the likely changes:

- A global dedicated network of 12-25 geodetic-VLBI stations, probably co-located with other space geodetic techniques (SLR, GPS among them), will be established.
- For both economic and scientific reasons, the VLBI antennas at these locations will be small and fast-moving. Antenna diameters in the range of 8-20 meters are a reasonable guess.
- In order to ameliorate the burgeoning RFI problems encountered by some of today's VLBI stations, particularly at S-band, it is likely that a shift to higher observing frequencies will take place, perhaps to 8/15 GHz or maybe higher. Or even systems that are able to observe over extremely wide single bands from, say, 5 to 20 GHz.
- Due to smaller antennas and likely weaker source strengths at higher observing frequencies, observing bandwidths must be widened to re-gain the necessary SNR for precise geodetic observing. A reasonable guess is that aggregate data bandwidths that must be captured for correlation will be at least several GHz wide, resulting in corresponding data rates of at least several Gbps/station.
- The duty cycle of observing will substantially increase at these dedicated stations, perhaps up to 50-75% averaged over a week.
- In order to operate efficiently and economically, these stations will be largely unattended by operators and must be remotely managed.

We will examine the possible options for data transport based on this likely scenario.

Modes of Operation

Though direct transmission of the data via e-VLBI to a correlator is clearly the most desirable from an efficiency point-of-view, there are a range of alternative data-transport strategies that should also be considered. Furthermore, the data-transport strategy may evolve with time as technology permits or economics dictates, so the design of the data-transport system may need to be designed from the beginning with this evolution in mind. The range of possible data-transport options is shown in Table 1, listed in order of least desirable to most desirable.

Situation at station	Data-transport mode
No data link to correlator	Record, ship to correlator
Nearby high-speed POP with link to correlator	Record, transport to POP, transmit to correlator and record
'Slow' network connection to correlator	Use disk FIFO buffer – transmit at whatever speed is available; constrains sustainable observation duty cycle
High-speed link at station, but not to all stations	Transmit to correlator and record
HIgh-speed links to all stations	Transmit and correlate in real-time (no recording)

Table 1: Possible modes of data transport (in order from least to most desirable)

Magnetic Tape

Until very recently, magnetic tape has been the VLBI recording media of choice and could potentially become a contender again, though this seems unlikely. Some advantages and disadvantages of tape are as follows:

Advantages

- Tape is <u>potentially a much lower cost media</u> than disk since all the complicated parts of a tape system reside in the tape drive, which need be purchased only once to service many tapes; however, this is not the current trend as the computer industry spends relatively little on magnetic-tape development compared to magnetic-disk development. Currently (early 2004) we are in a crossover period where the per-GB cost of computer tape media is roughly the same as disk, though the cost of disk media continues to drop faster.

Disadvantages

- Messy mechanical head-to-tape interface: The head-to-tape interface is a notoriously messy problem, particularly with the high bit densities in use today. Extreme measures are often necessary to guarantee proper tracking and close head-to-tape contact.
- Poor reliability: Because the head-to-tape interface is messy and the mechanics are virtually impossible to seal from contamination, tape systems are likely to almost always be more unreliable than magnetic disks, which have a sealed environment. Furthermore, since there is direct head-to-tape contact, headwear will result in a slow but constant change in performance, resulting in eventual failure is head replacements are not timely.
- Expensive Drives: The tape drives are expensive, and there is little guarantee of either forward or backwards compatibility of old media with new generations of tape drives. Old drives, in general, must be discarded. Head replacements are likely to be quite expensive, if they are available at all.
- Difficult to take advantage of new technology: Because compatibility must exist at the mechanical head-to-tape interface, it is difficult and expensive to take advantage of new technology.
- Relatively small market: The relatively small size of the market for magnetic tape technology, compared to disk technology, virtually guarantees slower technological progress and higher prices.
- Limited data rates: Single tape rates of about 1 Gbps seems to be difficult to significantly improve. The messy and imprecise nature of the head-to-tape interface limits areal bit densities, while head-to-tape tape speeds must be constrained to prevent 'flying'; this combination severely constrains significant advances in single-tape data rates. Aggregation of multiple tapes to increase data rate are very awkward and difficult to build, not to mention the complexities of managing the tapes, which must be packaged somehow so as to allow them to be extracted individually to be inserted into tape drives.

On balance, it is clear that the disadvantages of magnetic tape outweigh the advantages, relegating them to a historical museum piece.

Magnetic Disk

The rapid advance of magnetic-disk technology over the past few years has seen magnetic disks literally burst onto the VLBI stage as performance and capacity rapidly increase and price per GB decrease at an equally rapid rate. Still, we should assess the advantages and disadvantages of magnetic disks and try to project where the technology may be in 2010 and beyond. Some of the advantages and disadvantages are:

Advantages

- Fast technology advance: In terms of capacity, the technology advancement rate of disks over the past 10 years has outstripped the Moore's Law advancement rate of semi-conductor technology by about 50%; this shows some signs of slowing, but many analysts attribute this more to market forces rather than technological advances. Nevertheless, advances in capacity are remarkable even over the past 3 years, moving single-disk capacity from ~120MB in 2001 to 1TB today (February 2004). Improvements in single-disk data rates have been much slower, but aggregation of multiple disks into single physical modules to increase aggregate data rates is both practical and economical.
- Rapidly falling prices: Price per GB has been dropping about 50% per year for the past 5 years. Though there is some evidence this rate may be slowing, it is still a much faster decline that any other competing media. Typical disk media costs have now (early 2004) dipped below $1/GB and are expected to drop to ~$0.50/GB over the next year.
- Easy to take advantage of new technology: Because disk drives connect to the system at a standard electrical interface, the technology used in the disk drive is invisible to the user and may advance and change radically with no impact on the design of the data system.
- Very large market: It is obvious that the vast market for disk drives is the driving force in the rapid technology advance and decrease in price. Currently, there are no obvious challengers to the domination of magnetic disks for large mass-market random-access storage systems.

Disadvantages

- Shipping fragility: Compared to magnetic tapes, magnetic disks are more susceptible to shipping damage due to the delicate nature of their internal mechanisms. However, experience has shown that use that shipping susceptibility is somewhat vendor dependent. At least one vendors provides internal mechanisms to lock heads off the magnetic platter in a power-off condition, which appears to be more robust to shipping damage compared vendors who park heads on the non-rotating magnetic platter; with proper packing, however, both types survive with small losses, though the latter tend to have a higher rate of loss than the former. It should be pointed out, however, that some of the disk-based VLBI data systems can continue operate in the face of the loss of one more drives with only slightly degraded performance. And, of course, not all data-transport scenarios that involve local media require shipping the media (see Table 1).

– Limited data rates: Compared to magnetic tape, multiple-disk modules can easily support data rates to 1 Gbps. However, if one is driven to date rates of the order of 10 Gbps or higher, even disks become problematic due to the large number necessary to support these high rates. For example, at least 30 disks of today's commodity variety, operating simultaneously, would be necessary to sustain operation at 10 Gbps; the practicalities and economics of such situation may be marginal at best.

Other Types of Recordable Media

Over the years we have seen much ballyhooed promises for alternative types of media, particularly optical-based media, but in practice these have not become contenders for VLBI data storage. For examples, recordable optical disks are still more expensive that comparable magnetic-disk storage, especially re-writable optical media; some non-re-writable optical media is comparable in cost to magnetic disk cost, but data densities are rather now and the costs become recurring. Additionally, sustainable data rates on optical media a generally lower than for magnetic disks. Optical holographic storage has been long promised but has never become a commercial success.

In terms of recordable media, magnetic disks reign as king and show no hint that they will relinquish their crown over the next decade.

e-VLBI

e-VLBI, particularly of the real-time variety, is clearly the Holy Grail of VLBI data transport; e-VLBI is rapidly coming of age, though it also brings its share of problems as well. Some of the advantages and disadvantages:

Advantages

– Potential for very high data rates: Since increasing data rate translates into increasing sensitivity for geodetic VLBI, the promise of very high data rates for e-VLBI is very enticing. This is particularly attractive and important if observations are done with small antennas. Global networks operating at 10 Gbps/wavelength/fiber are standard commercial practice today, with the possibility of up to ~100 wavelengths on a single fiber for an aggregate of 1Tb/sec/fiber. Over the next few years, this is likely to improve by at least a factor of several.
– Rapid turnaround for results: e-VLBI, particularly real-time or near-real-time e-VLBI, has the potential for reducing the correlation processing delay to near zero and the final analyzed geodetic results in hours. This will meet or exceed the turnaround requirements specified in the IVS Working Group 2 for IVS Product Specification and Observing Programs.
– Lower or eliminate media costs: This is obvious, of course. In Table 1, each of the succeeding scenarios requires less recording of data at a station and hence a reduction of media cost, down to zero in the full 'real-time data transfer' scenario.
– Fully automated operations: The last three scenarios in Table 1 require no shipping of media and therefore allow for completely unattended operation of a station, helping to lower operational costs.
– Monitor station health: With the capability to return data rapidly to the correlator, it is possible to monitor the overall health on a comprehensive basis, as well as assist in the remote diagnosis of problems for assistance in repair.
– Enables distributed correlation: Though not practical today, distributed correlation using arrays of commodity PC's, perhaps geographically dispersed, is a possibility for the 2010 era. With e-VLBI, the data are easily dispersed to many geographically distributed sites. Experiments in distributed 'software' correlation are now taking place in Japan, and though the performance is not now sufficient to accommodate the needs of geodetic VLBI, the continued improvement in commodity PC computing power may change the equation by 2010.

Disadvantages

– Last-mile connectivity: Though high-speed global networks are now a reality, the 'last-mile' connection to stations is a potentially expensive issue that must be faced to move to full real-time capability.

There is no doubt that e-VLBI, particularly real-time or near-real-time, offers many advantages. One-time last-mile connectivity costs and recurring leasing costs are the primary issues to be resolved. Clearly, co-location of several space-geodetic techniques at each station is advantageous in help to reduce these costs.

The cost of fiber data connections can be kept to a minimum by ownership or lease of dark fiber or dark wavelengths on fiber; in these cases, the owner is also responsible for terminal equipment, but is free to upgrade or replace this equipment with no additional costs associated with the fiber itself. In general, the total communications-infrastructure costs are much lower for this approach compared to leasing a configured service from a fiber service provider.

Questions Still to be Answered

There are still a number of questions to be answered to complete the VLBI2010 analysis of the data-transport problem. Among the questions that come to mind?

1. Need more details in costing out the last-mile connections (hfh).
2. What are the technology options for the last-mile terminal equipment? (hfh)
3. A more detailed cost-benefit analysis of recording vs. e-VLBI (and other intermediate options) needs to be undertaken.
4. How can stations such as O'Higgins, which cannot be connected by fiber, be managed (satellite?)? (all)
5. Are current VSI (H,S and E) specifications adequate? Do they need to be modified/expanded? How? (yk, tk, arw)

Summary

Clearly, real-time e-VLBI is the data-transport method of choice for 2010 and beyond. However, the costs of full real-time e-VLBI must be weighed against the benefits; one of the less aggressive options in Table 1 may turn out to be more cost practical and cost effective.

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.