

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
HAYSTACK OBSERVATORY
WESTFORD, MASSACHUSETTS 01886

Telephone: 978-692-4764
Fax: 781-981-0590

24 January 2001

TO: Distribution
FROM: Alan R. Whitney
SUBJECT: Interim report on COTS-VLBI Phase I study

Introduction

Over the past several months we have been studying options to create a new COTS-based VLBI recording and playback system that can serve the VLBI requirements for the future. The study (Phase I) is being sponsored jointly by NASA, NRAO and JIVE. This report is intended to provide an interim update on the progress of the Phase I study now underway. A major result to date is a recommendation to adopt magnetic discs for the next generation of VLBI data systems.

Goals

As you will recall from the COTS-VLBI white paper [1], the goals for the COTS-VLBI system are as follows:

- Minimum of 1 Gbps data rate
- Economically upgradable/expandable to ~8 Gbps over the coming decade
- Design based primarily on unmodified off-the-shelf subsystems and components
- Modular, easily upgradeable as better/cheaper technology becomes available
- Robust operation, low maintenance cost
- Easy transportability
- Conformance to VLBI Standard Interface (VSI) specification
- Flexibility to support electronic transfer ('e-VLBI') and computer processing of recorded data
- Easy adaptability to HDR applications other than VLBI
- Minimum of 24-hour unattended operation

Areas of Investigation

At the outset of the Phase I study, a series of questions, summarized as follows, were posed with respect to the COTS-VLBI system:

1. Which recording technology is most suitable for COTS-VLBI? The potential contenders are:
 - a. Magnetic tape
 - b. Magnetic disc
 - c. Optical disc
2. What is the best overall implementation strategy? The goals are:
 - a. To maximize the leverage of COTS-based components and minimize custom design
 - b. To create a useful implementation in the shortest amount of time, with a full VSI-compatible system following.
 - c. To be easily and inexpensively upgradable as new COTS components become available to maximize the useful life of the system
3. What is a realistic development schedule? Two particular goals are in mind:
 - a. A useful, but perhaps not complete, COTS-VLBI implementation as soon as possible.
 - b. A fully VSI-compatible implementation

We will discuss the current state of the Phase I study in each of these areas.

Recording Technology

Summary of Conclusions

The following conclusions have been reached to date regarding choice of recording technology:

1. Tape technology is progressing as projected in the COTS-VLBI white paper and is usable for COTS-VLBI.
2. Exceeding industry predictions made even as recently as two years ago, magnetic-disc technology development is outpacing magnetic-tape technology development by a factor sufficiently large that a cost-crossover point for VLBI is likely to occur in the next few years, with rapid continuing advances continuing beyond that date.
3. Magnetic-disc recording is the technology of choice for COTS-VLBI. *We recommend that Haystack's efforts be focused on this technology.*

Discussion

The original COTS-VLBI concept, as outlined in the COTS-VLBI white paper, showed projections for increasing magnetic-tape data density and data rate based on both past history and industry roadmaps. As best as can be judged, these projections are coming into reality pretty much according to schedule.

What was not foreseen was the *rapid acceleration* of progress in magnetic discs over the past two years. Figure 1 shows a cost comparison, as a function of time, of the cost of magnetic-disc drives and magnetic-tape media (*not* including tape drives). The black diamonds show the progression of actual street prices of disc drives in \$/GB; the points for the last few years are for consumer-grade 'dimstore-variety' disc drives. Also shown are two 1998 disc-industry projections, which have been far surpassed, as well as the cost of computer tape media and the projected cost of the LTO tape media from our original COTS proposal. For reference, the costs of Mark IV/VLBA/K4 and S2/S3 media are also approximately plotted. As you can see from

Figure 1, the magnetic disc industry is projecting a cost of ~\$.30/GB by 2004 that will be incorporated into consumer-grade IDE-type drives with a single-drive capacity of ~1000 GB (1 Terabyte) for ~\$300. These costs are considerably lower than Mark IV/VLBA/K4 and even lower than projected S3 costs (more on that later). The current Mark 4/VLBA tape has a capacity of ~500 GB at a cost of ~\$1000, or ~\$2/GB. As early as late 2001, the cost of magnetic discs is expected to fall to ~\$1/GB for drives with a capacity of ~120 GB, so that even *by the end of 2001, disc storage prices will be roughly half of current Mark IV/VLBA media costs.*

It is important in this discussion to note that historically the cheapest disc drives are consumer-grade devices with IDE (aka ATA) interfaces, as opposed to more corporate-grade discs typically equipped with SCSI interfaces and somewhat higher performance. For the COTS-VLBI application there is no reason not to use the cheapest reasonable-quality consumer grade IDE disc drives; the necessary aggregate VLBI data rates can be achieved simply by operating as many discs as necessary in parallel. Unlike magnetic tapes, which require expensive drives to handle them, disc drives are fully self-contained devices so there is virtually no penalty for using many of them simultaneously.

There are a number of other obvious benefits of magnetic discs over magnetic tape:

- Rapid random access to any data
- Essentially instant synchronization on playback into a correlator
- Self contained; don't have to buy expensive tape drives
- Can always buy and use the latest and cheapest 'dimestore' IDE disc drives, while still using older ones
- Easily upgradeable host platform: the COTS host computer and interface cards can be easily and inexpensively upgraded as newer models become available

Optical discs have not developed to the capabilities projected a few years ago and remain uncompetitive with magnetic discs for the VLBI application. This could possibly change in the future. However, should optical discs overtake magnetic disc in the future, it is more than likely they would be supplied with IDE interfaces which are fully interchangeable with magnetic discs so that optical discs could be substituted transparently in a COTS VLBI data system.

Given these facts, *we recommend that the COTS-VLBI project at Haystack focus attention on magnetic-disc storage.*

HP LTO Magnetic Tape Drive Tests

In September of 2000 Haystack Observatory obtained a pre-production prototype of a Hewlett-Packard LTO tape subsystem of the type that was discussed in the COTS-VLBI white paper [1]. After some initial difficulties, this drive was successfully connected to a PC-based Linux system using a high-performance SCSI interface card (Adaptec 29160). Initial tests showed the drive performed roughly to the expected level. And independent industry reviews confirmed that it met the ~15 Mbyte/sec data rate claimed by its manufacturers. However, it has become apparent that further effort pursuing magnetic tape is unjustified, and attention has henceforth focused on magnetic discs. Further tests on magnetic tape systems are not planned.

Implementation Strategy

Summary of Conclusions

The following conclusions have been reached to date regarding design strategy:

1. The COTS-VLBI system will be based on a standard PCI bus commonly available in desktop PC systems.
2. A 'core' COTS-VLBI recording/playback system will first be created, likely utilizing fully commercial components, to implement a device which can simply record and playback a set of externally-clocked parallel bit streams. Such a system will be almost directly and immediately useful as a replacement for Mark 4 and VLBA tape systems, at data rates to at least 512 Mbps.
3. VSI-specific recording and playback interfaces will be designed to interface to the 'core' COTS-VLBI system to fully implement the VSI specification to 1 Gbps.
4. The operating system of choice is Linux, though early implementation may be under Windows.

Off-the-shelf System Elements

Platform

A survey of possible computer platforms and busses to be used for the COTS-VLBI system was conducted. The candidates examined were:

1. Standard PC platform with PCI bus (or busses)
2. CompactPCI bus in a card-cage format with plug-in processor card
3. VME bus in a card-cage format with plug-in processor card

From a cost-effectiveness point-of-view, a standard PC platform with a standard PCI bus is clearly the best choice. There are several clear advantages to this choice:

1. The cost of these platforms is extremely cost competitive, resulting in a very low price.
2. The widest and fastest PCI busses are almost always available first on these platforms.
3. A single modern PCI bus should easily be able to support a 1 Gbps data rate.
4. The selection of available COTS interfaces and subsystems available for this platform far surpasses that of any other.

Though CompactPCI and VME bus architectures are still common and readily available, there has been a massive shift to the standard PCI bus in recent years, and COTS-VLBI is well positioned to take full advantage of this shift.

COTS Interface Cards

Due to the large mass market for the standard PC platform and its ubiquitous PCI bus, the support available in the form of third-party hardware and software is massive. For the 'core' COTS-VLBI system, there are three important hardware interfaces that will be acquired from one of numerous third-party sources:

1. Parallel-input interface card: This card will receive a set of externally-clocked parallel bit streams to be transmitted to the recording media. Candidate interfaces have been identified from General Standards Corporation [2], Gage Corporation [3], Adlink, Inc. [4] and United Electronic Industries [5]. Other possibilities may also appear in a broader search.

2. Parallel-output interface card: This card will receive playback data from the recording media and transmit a set of externally-clocked parallel bit streams to the outside world. Candidate interfaces are available from the same set of manufacturers as above.
3. Interface cards to magnetic discs: Traditionally, very large disc systems have been implemented using SCSI disc interfaces, but the marketplace has recognized the inequity of pricing of SCSI versus IDE discs and now offers a variety of effective and inexpensive solutions to using large arrays of IDE discs. Interfaces from such companies as Boulder Instruments [6], Promise Technology [7], Raidzone [8] and 3Ware [9] are now being examined for suitability for the COTS-VLBI project. Some of these offer RAID type systems for fail-safe protection against individual disc failures, but such a system is necessarily somewhat more costly and it is not clear whether the benefits to VLBI are sufficient to justify these extra costs.

The 24-hour Challenge

Disc industry expectations are that in 2004 a terabyte (TB) disc will be available for ~\$300. The recording of a continuous data rate of 1 Gbps for 24 hours will consume ~11 TB of storage, so that an array of only ~12 such discs for a total cost of \$3600 will comfortably operate unattended at 1 Gbps for 24 hours, with a shipping weight <~18 kg. This compares to today's cost of >\$20K for VLBA/Mark IV media of the same storage capacity with a shipping weight of ~200 kg.

Media Removability and Transportability

In any disc-based system, removability and transportability of drives is a significant concern. Unlike for magnetic-tape media, there are no known 'disc changers', so a disc-based COTS-VLBI system must be designed for easy and reliable removal and insertion of drives on a routine basis. Several possible solutions are being investigated, from rugged single-disc carriers to 7-disc cartridges, all of which are designed for multiple insertion/removal cycles. We expect more potential solutions to this problem to become available over the next year. When modern disc drives are powered down, they are quite robust to external handling forces, and can easily be shipped in padded containers.

'Core' COTS-VLBI System

The 'core' COTS-VLBI system implementation will be non-VLBI specific, but nevertheless will be almost immediately useful for VLBI observations. Illustrated in Figure 1, it will have the following basic specifications:

Input: 32 parallel ECL bit-streams externally clocked at up to ~20 MHz.

Output: 32 parallel ECL bit-streams externally-clocked at up to ~20 MHz, coarsely synchronizable to an externally-defined epoch.

The 'core' COTS-VLBI recording system will simply record a set of parallel bit streams, and the COTS-VLBI playback system will reconstruct them. No time or synchronization information will be added to these bit streams. However, *this system will be almost immediately useful for Mark 4 and VLBA systems*. For recording, the output of a Mark 4 formatter or VLBA formatter can be directly attached to the input interface, which will then record fully formatted VLBI data in almost exactly the same way as a standard Mark 4/VLBA recorder, recording at the usual 4.5, 9.0, or 18 Mbits/sec/bit-stream. On playback, the output can be directly attached to the correlator

input in the same way as current Mark 4/VLBA playback drives; synchronization will be controlled in a similar manner as is currently done with Mark 4/VLBA playback drives.

For VLBA systems, the 'core' COTS-VLBI system will support the equivalent of 32 tracks at 9 Mbps/track (including parity) for a useful aggregate data rate of 256 Mbps. For the Mark 4 system, the 'core' COTS-VLBI system will support the equivalent of 32 tracks at 18 Mbps/track (including parity) for a useful aggregate data rate of 512 Mbps.

The effort in creating the 'core' COTS-VLBI system is almost entirely software, with only minor hardware-design requirements in physical interfacing to formatters and correlators. There will be some required software updates in VLBI data-acquisition and correlator systems, but these should be relatively minor.

Phase II COTS-VLBI System

The Phase II COTS-VLBI System will add full VSI-H compatibility on top of the 'core' system without affecting the 'core' capabilities, as shown in Figure 2. This will be accomplished with a pair of custom-designed VLBI-specific interface cards which attach to the 'core' COTS-VLBI system:

VSI-H formatter: This card will accept sampler and 1-pps outputs, according to the VSI-H specification, add the necessary timing and synchronization information, and transmit a set of 32 parallel bit streams to the 'core' COTS-VLBI system for recording. Though this card may be designed to reside on a PCI bus for control purposes, it will not transmit data through the PCI bus.

VSI-H deformatter: This card will accept the output from the 'core' COTS-VLBI playback system, strip all time and formatting information, and re-construct the data according to the VSI-H specification. We are also investigating the required additions to implement all of the functions of a Mark 4 Station Unit in the deformatter card as well; if this is possible, the playback output of the COTS-VLBI playback system will be fully corrected for the station-based delay model and may be input directly into a Mark 4 correlator, bypassing the existing Station Units. This may be particularly useful for expansion of any of the four existing Mark 4 correlator systems to accommodate more stations, since all of the correlator systems have been built with fewer than the number of playback systems that can be accommodated by the correlator proper.

S3 as an Alternative?

As can be seen from Figure 1, the media cost for the S3 system, now under development by Crestech of Canada, is projected to be ~\$.40/GB. This is quite competitive with the ~\$.30/GB expected of disc drives in ~2004. However, consideration of the S3 system should take into account the capital cost of an S3 system which includes a number of fairly expensive custom-modified digital-video helical-scan tape drives plus an automated tape changer. Furthermore, whether the S3 will be able to take advantage of the expected continuing developments, and lower prices, in computer-related consumer-grade products, is an open question.

Development Schedule and Funding

We project the Phase I COTS-VLBI study will be completed by September 2001. At that time we expect to have identified all the relevant technologies and created a detailed design strategy for implementation of both a 'core' COTS-VLBI system and a full 1-Gbps VSI-compatible COTS-VLBI system along the lines we have described above. Further work will depend on the availability of additional funding. A specific proposal for Phase II (as defined in our original proposal) will be prepared and submitted in summer 2001 to allow sponsors to program the necessary funds in FY02 (beginning October 2001).

Summary

Based on our studies to date of the COTS-VLBI data system, we recommend that development focus on magnetic discs as the media of choice. We feel the advantages of magnetic discs, both in terms of performance and projected costs, are very compelling. Given funding continuation, a useful 512 Mbps 'core' COTS-VLBI system should be operational by early 2002, with a full 1-Gbps VSI-compatible system operational by ~2003.

References:

- [1] "Concept for an Affordable High-Data-Rate VLBI Recording and Playback System", MIT Haystack Observatory, 30 September 1999.
- [2] <http://www.generalstandards.com/>
- [3] <http://www.gage-applied.com/>
- [4] <http://www.adlink.com.tw/products/nudaq/7300a.htm>
- [5] <http://www.ueidaq.com>
- [6] <http://www.boulderinstruments.com>
- [7] <http://www.promise.com>
- [8] <http://www.raidzone.com/>
- [9] <http://www.3ware.com/>

Distribution:

- J. Bosworth, NASA
- T. Clark, NASA
- J. Romney, NRAO
- R. Schilizzi, JIVE
- J. Webber, NRAO
- W. Wildes, NASA




Figure 1: Disc/Tape Price Comparisons

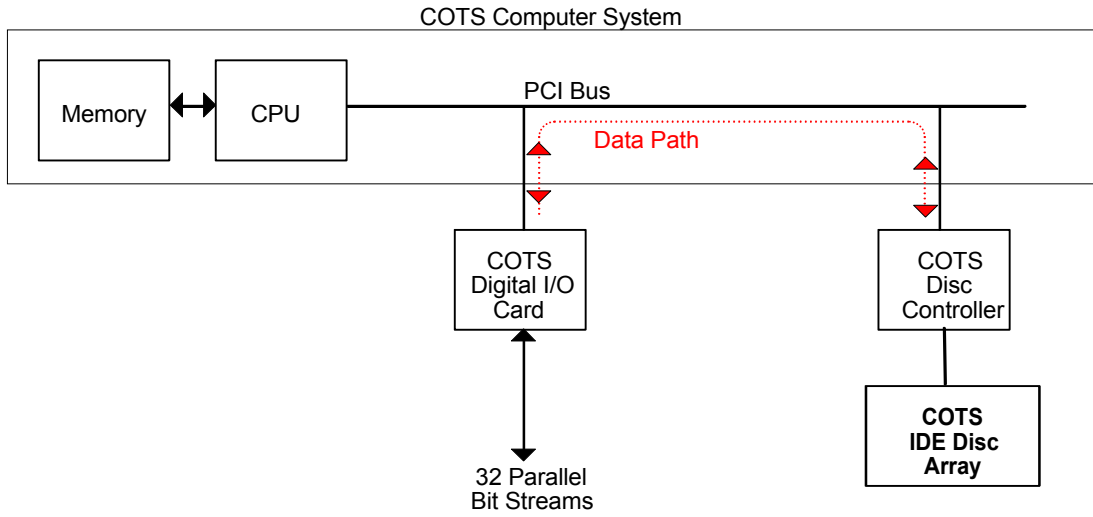


Figure 2: Block diagram of 'core' COTS-VLBI record/playback system

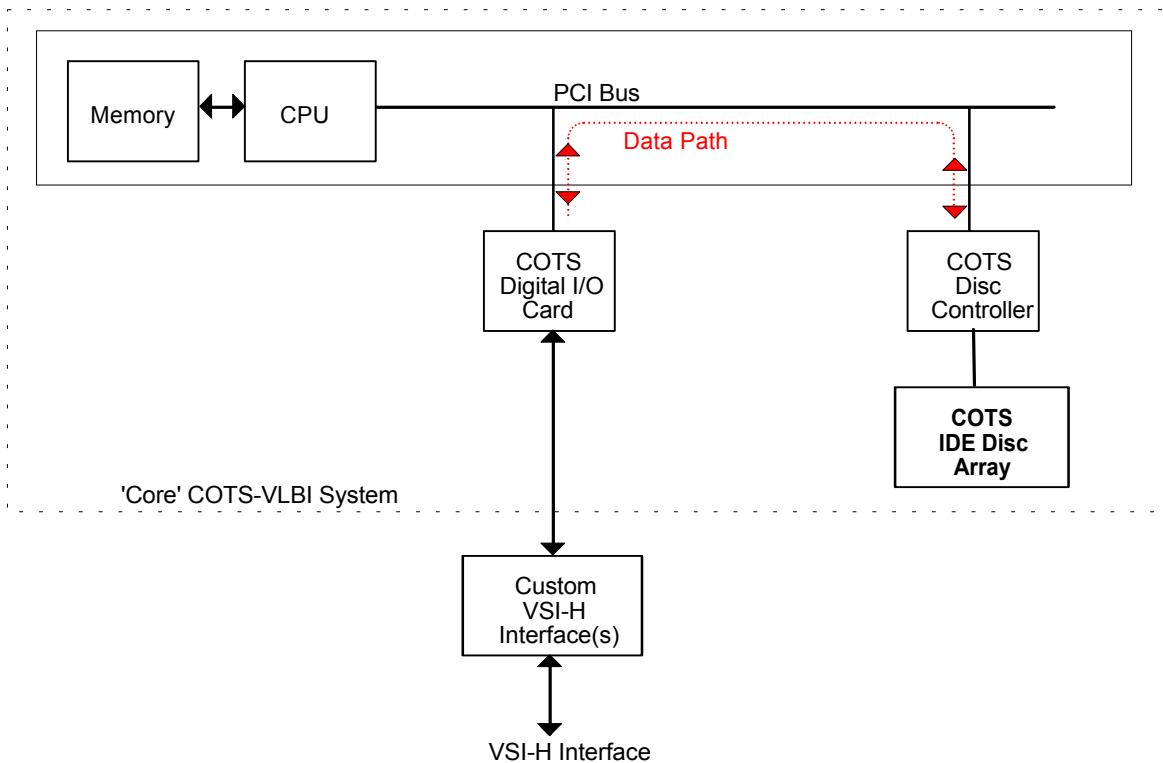


Figure 3: Block diagram of VSI-compatible COTS-VLBI record/playback system