Alan R. Whitney: Mark 5 Disc-Based Gbps VLBI Data System, IVS 2002 General Meeting Proceedings, p.132–136
http://ivscc.gsfc.nasa.gov/publications/gm2002/whitney1

Mark 5 Disc-Based Gbps VLBI Data System

Alan R. Whitney

MIT Haystack Observatory

e-mail: awhitney@haystack.mit.edu

Abstract

The Mark 5 system is being developed as the first high-data-rate VLBI data systm based on magnetic-disc technolgy. Incorporating primarily low-cost PC-based components, the Mark 5 system will support data rates up to 1024 Mbps recording to an array of up to 16 inexpensive removable IDE discs. An initial demonstration system was used in March 2001 to record data at 576 Mbps, with correlation on the Mark 4 correlator at Haystack Observatory.

With the continuing fall in disc prices, IDE discs are already becoming competitive with the cost of Mark 4/VLBA/K4 tape, with the expectation that prices will continue to fall to a level below ~$1/GB with capacities of hundreds of GB per disc.

The Mark 5A system, which is a direct replacement for a Mark4 or VLBA tape transport, will deploy ~12 prototype units in mid-2002 to stations and correlators around the world, at a cost of ~$20K/unit. A fully VSI-compliant Mark 5 system, dubbed Mark 5B, will be available in late 2003, along with the necessary adapter interface required to utilize the system with existing Mark 4 and VLBA data-acquisition and correlator systems.

The Mark 5 development effort at Haystack Observatory is supported by BKG, EVN/JIVE, KVN, MPI, NASA, NRAO and USNO.

1. Introduction

The Mark 5 system is being developed as the first high-data-rate VLBI data system based on magnetic-disc technology. Incorporating primarily low-cost PC-based components, the Mark 5 system will support data rates up to 1024 Mbps, recording to an array of up to 16 inexpensive removable IDE discs. An initial demonstration system has been used to record data at 576 Mbps, with correlation on a Mark 4 correlator at Haystack Observatory. A program is now in place for the development of an operational Mark 5 system.

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 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 disc technology. The first Gbps Mark 5 systems will be put into operations in 2002.

2. Why Discs?

Though both magnetic-disc technology and magnetic-tape technology have made great strides over the past few years, the pace of magnetic-disc development has been so great that it is very likely that disc storage will become cheaper than magnetic tape storage by ~2004. Figure 1 shows the comparison of disc and tape prices since 1980. Current consumer IDE disc costs are ~$2/GB and falling; current Mark4/VLBA tape prices are ~$2/GB and remaining steady. By ~2005-2006, industry projections suggest the price of discs will fall to ~$0.5/GB. Similarly, current single-disc capacities are ~160 GB and rising; by ~2004-2005, single-disc capacities are expected to rise to 500-1000 GB! A Mark 5 system with sixteen 700 GB disc drives will record 1024 Mbps continuously for 24-hours unattended.

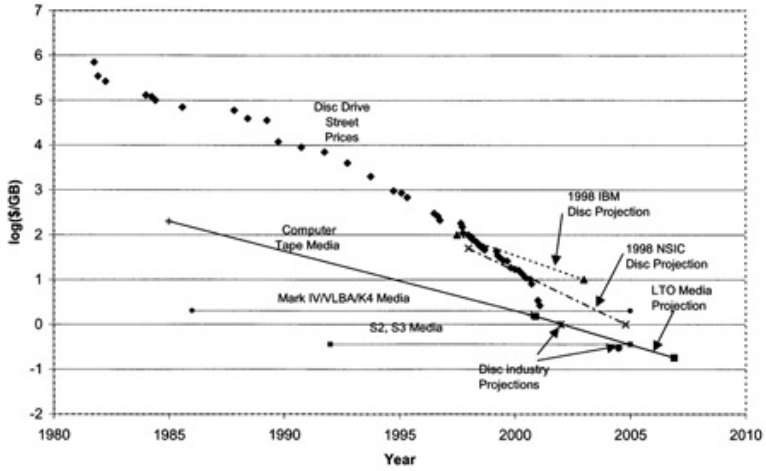

Figure 1. Disc/Tape Price Comparisons

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

- Readily available inexpensive consumer product
- continually improving in price/performance with standard electrical interface
- Self contained; do not have to buy expensive tape drives, 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!

3. Mark 5 Development Program

Based on the success of the 512 Mbps Mark 5 demonstration unit in early 2001 (developed and demonstrated in 3 months time, shown in Figure 2), Haystack Observatory is now undertaking the development of 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. Mark 5A: Records 8, 16, 32 or 64 tracks from a Mark4/VLBA formatter, up to 1024 Mbps, and plays back in the same Mark4/VLBA format. As such, the Mark 5A is a direct replacement for a Mark4/VLBA tape drive. We anticipate deployment of ∼20 systems mid-2002.
2. Mark 5B:VSI-compliant system, up to 1024 Mbps; no external formatter necessary. Will be backwards compatible to existing Mark4/VLBA correlator systems. We anticipate deployment in 2003.

A Mark 5A system may be upgraded to a Mark 5B system simply by replacing PCI boards in the host PC.

The cost of either the 1 Gbps Mark 5A or Mark 5B recording or playback system (without discs) is expected to be <∼$15K with a do-it-yourself kit. These costs are more than an order-of-magnitude below current costs of available tape-based Gbps systems. With currently available disc drives of ∼160 GB each, a Mark 5 system will record 1024 Mbps of user data for ∼5.5 hours using 16 discs. Normally, the 16 discs will be divided into two 8-disc sets. The Mark 5 will automatically and seamlessly ping-pong between the two 8-disc sets. Disc drives are mounted in carriers made for multiple insertion/removal cycles. When modern disc drives are powered down, they are quite robust to external handling forces and can be shipped easily in padded containers. Including the carriers, the shipping weight per disc is <∼0.9kg, so that the shipping weight of 8 discs plus shipping container is ∼9kg.

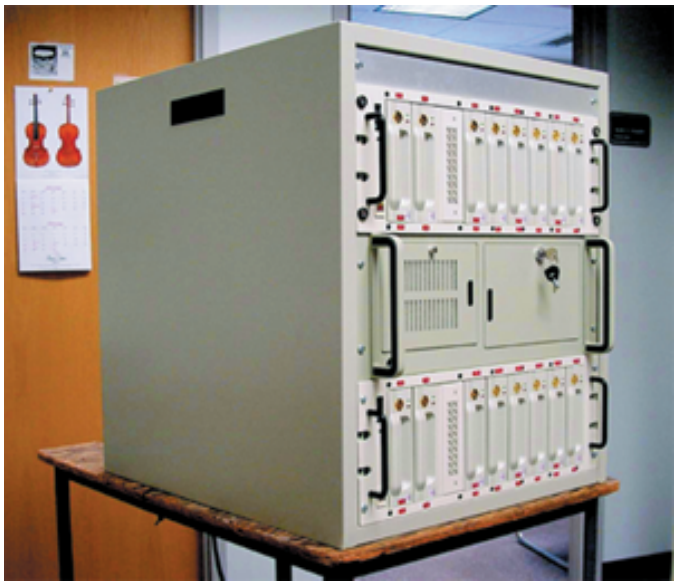

Figure 2. Photograph of demonstration Mark 5 System

Figure 3 shows a simplified block diagram of the Mark 5 system. The heart of the system is a streaming disc interface which supports 16 disc drives at data rates in excess of 1024 Mbps. For recording, external data from a Mark 4 or VLBA formatter is transformed to 32 parallel bit streams to be accepted by the disc interface for recording. On playback, the 32 bit streams played back from the discs are re-transformed to re-create the formatter input. The VSI version of the Mark 5, called Mark 5B, will operate in a similar manner.

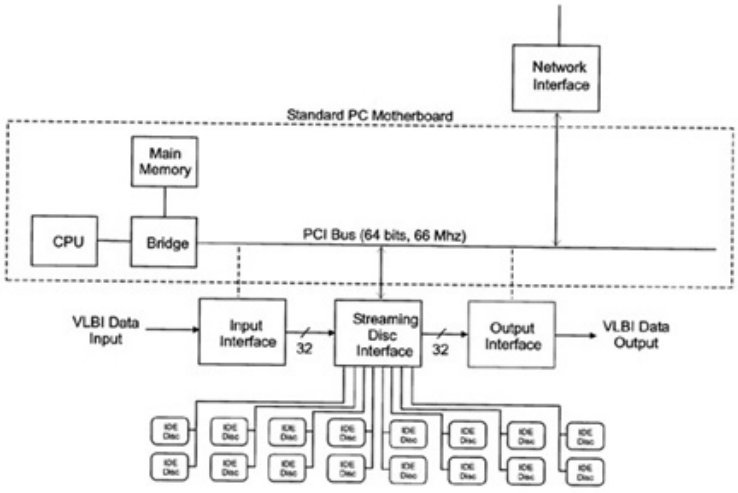

Figure 3. Simplified block diagram of Mark 5 system

Figure 4 shows the anticipated packaging of the operational Mark 5 system, which will be housed in a single 7"-high chassis.

4. 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 a 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.

5. e-VLBI Support

The Mark 5 system allows easy connection of a VLBI data system to a high-speed network connection. Because the Mark 5 system is based on a standard PC platform, any standard network connection is supported.

Depending on the availability of high-speed network connections, this can be accomplished in at least two ways:

1. Direct Station to Correlator: If network connections allow, data may be transferred in real-time at up to 1 Gbps from Station to Correlator, either for immediate real-time correlation or buffering to disc at the Correlator.
2. Station Disc to Correlator Disc: If network connections are not sufficient to allow real-time transmission of data to the Correlator for processing, data may be recorded locally to disc at the Station, then transferred to disc at the Correlator at leisure for later correlation.

Depending on the available network facilities, either entire experiments or small portions of experiments may be transmitted electronically. The latter may be particularly useful for verifying fringes in advance of important experiments.

Haystack Observatory is being supported by DARPA to demonstrate Gbps e-VLBI data transmission between Haystack Observatory and NASA/GSFC (~700 km) using the Mark 5 system. Data will be collected at the Westford antenna at Haystack Observatory and the GGAO antenna at NASA/GSFC and transmitted in real-time to the Mark 4 correlator at Haystack Observatory.

6. Summary

The Mark 5 system promises to move VLBI data systems to dramatic new levels of high-performance and low-cost by leveraging the enormous investments of the computer industry in high-speed data technology. Within only a very short time, the possibilities to economically expand VLBI observing programs by large factors appear to be within reach.

Figure 4. Mark 5 prototype system