Mark 5 Disc-Based Gbps VLBI Data System

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

Abstract

The Mark 5 system is being developed at MIT Haystack Observatory 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. The system is being developed with support from BKG, EVN, KVN, MPI, NASA, NRAO and USNO; prototype Mark 5 units are now in routine use at several stations and correlators, with ~20 units expected to be in place by summer 2002.

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. It is expected that disc-based VLBI systems will completely 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 disc technology.

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 no less than spectacular, far exceeding even disc-industry projections. Figure 1 displays a comparison of disc and tape prices over the past several years, showing that disc prices (on a $/GB basis) continue downward in a still-accelerating trend. Current (spring 2002) consumer IDE disc costs are ~$1.5US/GB and falling rapidly; current Mark4/VLBA tape prices are ~$2US/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 ~120 GB and rising; by ~2005-2006, single-disc capacities are expected to rise to 500-1000 GB! A single Mark 5 system with sixteen 700 GB disc drives will record continuously 1024 Mbps for 24-hours unattended.

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; don't 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!

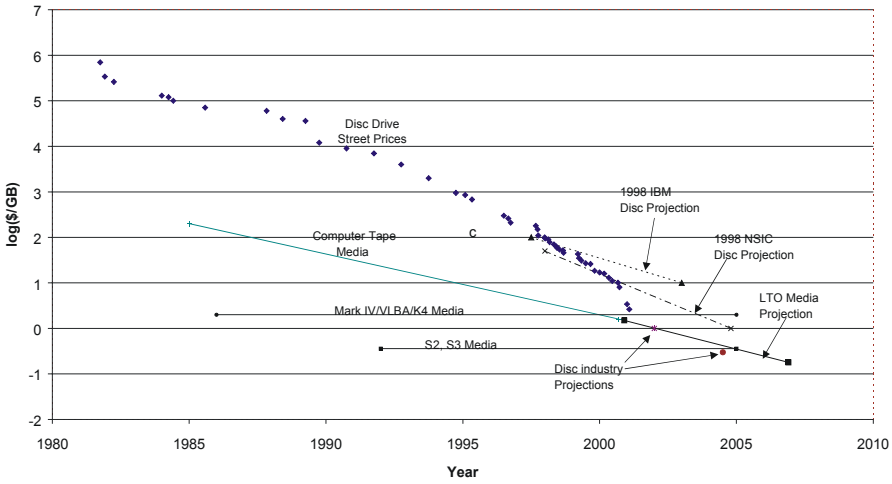

Figure 1: Disc and tape prices vs. time

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!), 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. Mark 5A: 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 will record 8, 16, 32 or 64 tracks from a Mark4/VLBA formatter, and will play 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 routine use at several antennas and correlators, with ~20 Mark 5A prototype systems to be deployed late summer 2002.
2. Mark 5B: The Mark 5B is VSI-compliant 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 2003.

Figure 2 shows a photographs of the prototype Mark 5A system. A Mark 5A system may be upgraded to a Mark 5B system simply by replacing a PCI board in the host PC.

Figure 2: Prototype Mark 5A VLBI data system

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 ~120 GB each, a Mark 5 system will record 1024 Mbps of user data for ~4.3 hours using 16 discs. The number of discs used may be 1 to 16 depending on the data-rate and data-capacity requirements of the user. Eight discs are sufficient to support a data-rate of 1024 Mbps.

Disc drives are mounted in standard 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 <~1.0 kg; the boxed shipping weight of 8 discs is ~9kg.

Due to the intricacies of handling many individual disc drives, serious consideration is being given to designing an '8-pack' disc carrier which holds 8 discs as an individual removable unit. Such a carrier would significantly simplify disc handling in routine heavy-usage situations for which the Mark 5 system is designed.

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.

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.

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.