

Mark 5 Disc-Based Gbps VLBI Data System

A. R. Whitney

MIT Haystack Observatory, Rt. 40, Westford, MA 01886 USA

Abstract. 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/playback to a modular array of 8 inexpensive removable ATA/IDE discs.

ATA/IDE discs have already fallen to <\$2.00/GB, below the cost of Mark 4/VLBA tape, with the expectation that prices will continue to fall and capacities continuing to expand to several hundred-GB per disc.

Besides recording and playing from disc, the Mark 5 system will be fully e-VLBI compatible, utilizing standard Gigabit Ethernet connections. For real-time e-VLBI usage, data may be either directly transmitted or received; for quasi-real-time usage, e-VLBI data may be buffered through the disc array.

A development effort is now underway at Haystack Observatory, with support from BKG, EVN, KVN, MPI, NASA, NRAO and USNO to fully develop the Mark 5 system. The development plan is highly attentive to compatibility requirements with existing Mark 4 and VLBA data-acquisition and correlator systems, with prototype deployment of ~20 systems expected in 2002. A fully VSI-compliant Mark 5 system will be introduced in 2003.

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 8 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.

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 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, show-

ing that disc prices (on a \$/GB basis) continue downward in a still-accelerating trend. Current (spring 2002) consumer IDE disc costs are ~\$2US/GB and falling; 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 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!), 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 three stages:

1. Mark 5P: This MarkP records 32 tracks from a Mark4/VLBA formatter. The maximum data rate is

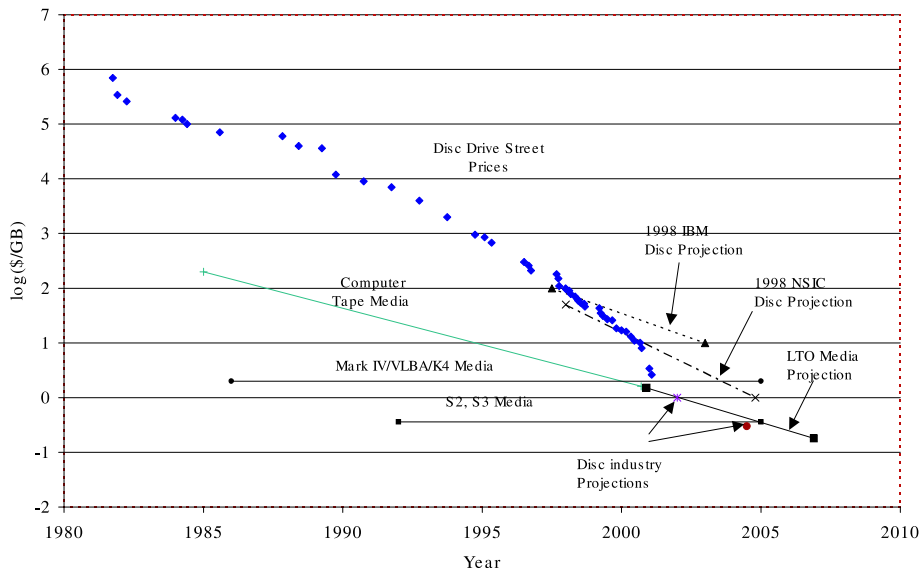


Fig. 1. Disc/Tape Price Comparisons

512 Mbps (16 Mbps/tk) with a Mark 4 formatter or 256 Mbps (8 Mbps/tk) with a VLBA formatter. The Mark 5P system is in routine use at several locations, with ~ 15 Mark 5A prototype systems to be deployed by late summer 2002.

2. 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 correlator. 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. Approximately 20 Mark 5A systems will be deployed in fall 2002.
3. 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.

The upgrade path from Mark 5P to Mark 5A to Mark 5B is, in each case, just the simple replacement of a single PCI card in the Mark 5 chassis.

A new 5U chassis is being developed for Mark 5 that holds two removable modular '8-pack' disc carriers (see Figure 2). Recording or playback is to/from a single 8-pack at a time, always spreading data across 8 disks. The system is being designed to 'ping-pong' seamlessly between the two 8-pack modules for uninterrupted data recording and playback. Idle 8-packs may be 'hot-swapped' at leisure for easy media management. Prototypes of the new chassis are expected to be delivered in early August 2002.

The '8-pack' disk module will be managed logistically in exactly the same way as a Mark 4 or VLBA tape, with each 8-pack assigned an 8-character VSN. A bar-coded

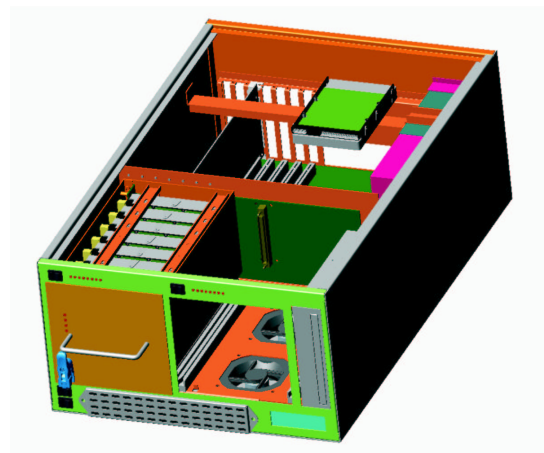


Fig. 2. Mark 5 prototype system.

permanent VSN label will be attached, just like tape, and the VSN will be electronically readable by the Mark 5 unit from a 'permanent' data field on the disks. This will allow the easy migration of media management from tape to disc, with almost 100 % re-use of existing tape management software transferable to disks.

The initial design of the chassis and 8-pack module will accommodate standard parallel-bus IDE/ATA disks. However, the unit will be upgraded to support both parallel-bus and serial-ATA disks when the expected arrival of serial-ATA disks becomes reality.

From a software point-of-view, the Mark 5P and Mark 5A appear to the correlator in much the same way as

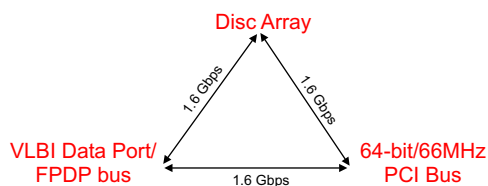


Fig. 3. Mark 5 'Triangle of Connectivity'.

tape, except that a footage count is replaced by a corresponding byte count. This means that existing correlator software can be easily adapted to the Mark 5P/A system. The scans on a Mark 5B system will be separated into standard named Linux files that appear identical to the files on a CD-R.

The cost of either the 1 Gbps Mark 5A or Mark 5B recording or playback system (without discs) is expected to be $< \sim \$15K$ with a 'do-it-yourself' kit. These costs are more than an order-of-magnitude below current costs of available tape-based Gbps systems.

4. Brief Theory of Operation

The Mark 5 is based on a standard PC platform using a combination of COTS and custom-designed interface cards.

The heart of the system is a 'StreamStor disc interface card from Conduant Corp that is specially designed for high-speed real-time data-collection and playback. The StreamStor card supports three physical interfaces in a 'triangle of connectivity' as shown in Figure 3:

1. Data Port/FPDP: This port is present as a 32-bit card-top bus which supports the industry-standard 'Front-Panel Data Port' interface specification. This is a two-way port through which high-speed real-time data may be either input or output. All 32-bits of the FPDP bus are always active.
2. Disc array: This port supports up to 2 arrays of 8 standard IDE discs for reading or writing, arranged in 8 master/slave pairs.
3. PCI bus: This is the standard connection to the host PC platform, however the StreamStor card supports a 64-bit/66MHz bus, though it is backwards compatible with standard 32-bit/33MHz buses.

The triangle of connectivity shows that data may be moved in either direction between any two of the three ports. The StreamStor card supports a maximum sustained data transfer rate of up to 1600 Mbps between any two ports, though only one connection path may be exercised at a time and the maximum data rate for VLBI usage is anticipated ~ 1024 Mbps. The path exercised for

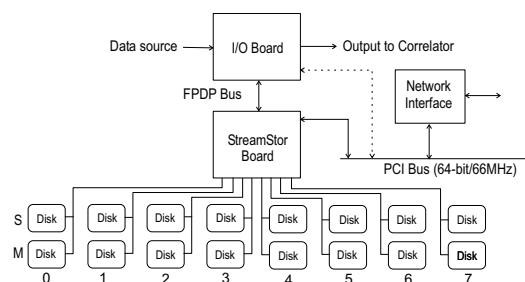


Fig. 4. Simplified block diagram of Mark 5 system.

traditional VLBI observations is between the FPDP bus and the disc array; note that in this mode, the VLBI data never touch the PCI bus, so the speed of the PC platform is largely irrelevant. Of course, the path between the disc array and the PCI bus allows the PC to read and verify VLBI data written to the discs via the FPDP port. The direct connection of FPDP bus to PCI bus will be used in upcoming e-VLBI experiments where data are transferred directly to a high-speed network and are not recorded locally. An on-board 512 MB buffer provides the necessary elasticity between the three connection nodes to support full real-time operation.

In order to record or playback VLBI data, a translation must be made between 'normal' VLBI data interfaces (formatters, correlators, VSI) and the 32-bit FPDP bus. This is done by an 'I/O board' connected to the StreamStor and host computer as shown in Figure 4. The I/O Board accepts data from the output of a Mark 4 or VLBA formatter and translates it to FPDP-bus format. Conversely, the I/O Board accepts data from the FPDP bus (from either the input part of the I/O Board or from the StreamStor) and translates it back into a form usable by existing Mark 4/VLBA correlators. For the Mark 5P system, 32 track outputs from a formatter are translated directly to corresponding 32-bits of the FPDP bus; this means that disc usage will be 100 % only if the data can be 'fanned-out' to 32 tracks by the formatter, which is normally the case. The Mark 5A, on the other hand, will allow 100 % disc utilization when recording 8, 16, 32 or 64 tracks from the formatter, and will also strip parity bits to improve data capacity by $\sim 12\%$.

The format of the data recorded to disc is such that barrel-rolling and/or data modulation are not necessary. The StreamStor card collects data from all 'tracks' in 64 kB chunks and writes these chunks sequentially in a round-robin fashion to the disc array. Therefore barrel-rolling is not necessary; if barrel-rolling is applied, it may complicate any post-recording analysis by the host PC. 'Data modulation' is an artifact of tape recording and is not needed by discs.

Data are recorded to discs in a special format optimized for high-speed real-time performance. In addition,

the format allows load shifting away from slow or failed discs to maintain error-free recording under less than optimal conditions. However, the data format on disc is entirely transparent to the user.

5. 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.

6. 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.