# Mark 5 Status Update

Alan R. Whitney

MIT Haystack Observatory

**Abstract**

The Mark 5A 1Gbps disk-based VLBI data system has been rapidly adopted with ~100 units now deployed around the world, along with ~1000 Mark 5 disk modules totaling ~1.5PB of storage capacity. As a result, the majority of VLBI observations are now recorded on disk systems. The Mark 5B system, which will be compatible with the VSI-H specification, is being developed at Haystack Observatory. The Mark 5B includes 'Station Unit' capabilities to allow it to replace the existing Mark 4 Station Units on the Mark 4 VLBI correlator. The Mark 5B system is expected to be available for purchase near the end of 2005.

**Introduction**

The Mark 5 VLBI disk-based data system is being developed in two stages:

1. Mark 5A: A 1Gbps disk-based record/playback system that is a plug-compatible replacement for a Mark 4 or VLBA tape transport

2. Mark 5B: A 1Gbps disk based record/playback system that supports the VLBI Standard Interface Specification (VSI). Mark 5A and Mark 5B use the same disk modules, but record different data formats onto disk.

The Mark 5A system was placed into operation in 2003; ~100 Mark 5A units are now deployed worldwide and have largely displaced the Mark 4 and VLBA tape transports. Mark 5B is expected to be ready for deployment in late 2005. Complete information on the Mark 5 system is available online [Ref 1].

**Mark 5A Data System**

The primary characteristics of the Mark 5A data system (Figure 1) are:

- Direct plug-compatible replacement for 64-track Mark4 or VLBA tape drives
- Record/Playback at rates up to 1024 Mbps
- Two independent removable '8-pack' disk modules per system can be used in 'ping-pong' fashion for near-continuous recording
- Will record 8, 16, 32 or 64 tracks from Mark4 formatter (1024 Mbps max) or VLBA formatter (512 Mbps max)
- Parity bits are stripped before recording and re-inserted on playback
- Arbitrary mixing of modes (#tracks, data rate, bits/sample) is allowed, always using 100% of installed disk capacity

The Mark 5A is based on a standard PC platform with two specialized interface cards:

1. A high-speed 'StreamStor' disk interface card from Conduant Corporation which supports recording and playback at 1Gbps to/from a bank of 8 commodity IDE disks

2. A specialized VLBI interface card designed at Haystack Observatory which emulates the data interfaces of a Mark 4 or VLBA tape transport.

The Mark 5A system is now used routinely for both geodesy and astronomy VLBI experiments with data rates up to 1024 Mbps.

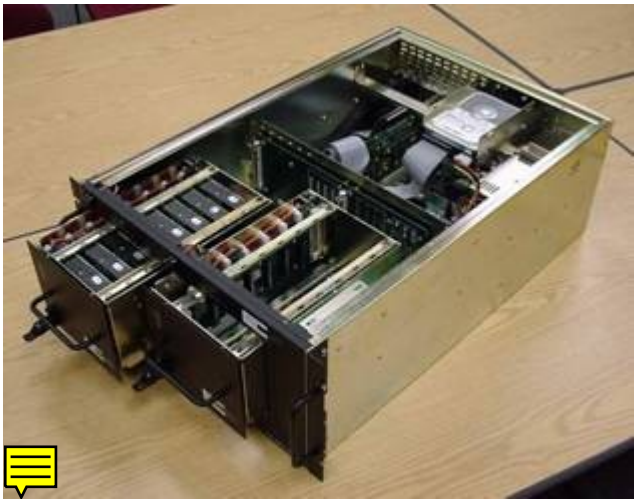

Figure 1. Mark 5A data system

**Disk-Media Status**

The original impetus for the development of the Mark 5 system, starting in 2001, was the rapidly falling price of disk storage and the expectation that disk storage would be less expensive than tape storage within a few years. Now, just four years later, these expectations have already been met and exceeded (see Figure 2).

- Hard disk price vs capacity/performance continues to drop

  - Now approaching ~$0.50/GB and will continue to drop
  (Mark 4/VLBA tape is ~$2.00/GB)

- 250 GB disks are now common –
8-disk module of 250GB disks is comparable to ~7 VLBA/Mark 4 tapes

- 400 GB disks are available –
8-disk module of 400GB disks is comparable to ~11 VLBA/Mark 4 tapes

- 700 GB disks expected by end 2005 –
8-disk module of 700GB disks is comparable to ~19 VLBA/Mark 4 tapes;
*2 modules will support ~24 hours @ 1 Gbps unattended!*

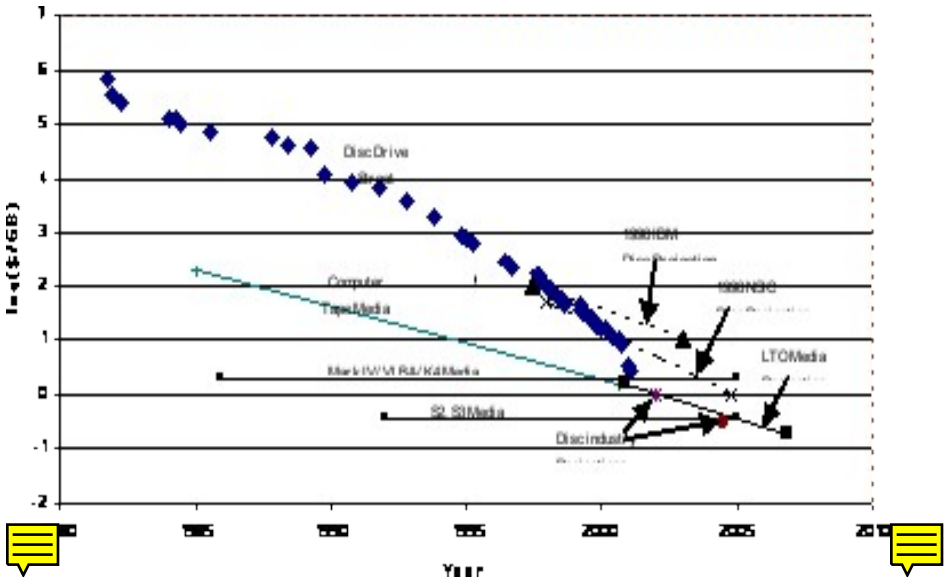


Figure 2: Comparison of disk and tape prices vs. time

Though disks have, by and large, been much more reliable than tape, disks have not been without problems:

- Based on statistics collected at Haystack, average disk drive failure rate is ~0.5% per year
- Failure rate of Hitachi 250GB has been higher than expected and Conduant recommends discontinuing their use
- Conduant has qualified drives from Maxtor, WD and Seagate and is shipping the following disks in assembled Mark 5 disk modules:
    - 400GB – Seagate
    - 250GB – Western Digital RAID Edition (high-reliability)
    - 250GB – Maxtor Maxline III (high-reliability)
- Disk reliability at high altitudes is under investigation. Most disks are specified to operate only to ~3000m altitude

**Mark 5A Software/Firmware Upgrades**

The Mark 5A continues to be upgraded for both more capability and reliability. Some of the recent upgrades include:

- Improved capabilities to deal with bad or slow disks on both write and read
- Ability to recover from inadvertent operational errors using ‘recover’ command
    - recover data from interrupted recording
    - recover from accidental use of ‘sstest’ or ‘WRSpeed’ test programs
    - recover data from accidentally erased module

- Support of disk FIFO mode up to 512Mbps

- Implementation of 'disk state mask' (being used by NRAO)

- Improved data-buffer handling for support of faster network transfers

Considerable effort has been expended to allow the Mark 5A to continue to operate under conditions of slow or bad disks. In particular:

- Writing: Shift the writing load away from the bad disk(s) and continue writing on the good disks. In most circumstances, this will still allow the full data rate to be captured, but at a reduced disk-module capacity (since the storage of at least one disk is lost)

- Reading: Substitute a 'fill pattern' for data that is not received in a timely manner from slow or bad disks. This special fill pattern is recognized by the Mark 5A playback electronics and causes wrong parity to be inserted, which in turn invalidates the data to the correlator. Since data from all recorded channels are spread evenly over the disks, the loss on one disk drive typical means the loss of about 12.5% data from each channel, resulting in an SNR loss of ~6%, which is usually tolerable.

**Plans for Serial-ATA (SATA) Disk Support**

IDE-interface disk drives (also know as Parallel-ATA or 'PATA') are being rapidly replaced in the marketplace by Serial-ATA ('SATA') disk drives. Currently, the Mark 5 supports only disk modules with PATA drives, but support for SATA disks is planned:

- PATA and SATA disk modules will be interchangeable in Mark 5 systems

- Biggest challenge was finding module SATA connectors with sufficient durability for many insertion/removal cycles

- Requires upgrade kit for current Mark 5 systems - new chassis backplane with separate connectors for PATA and SATA (same for both Mark 5A and Mark 5B)

- New SATA disk modules have been designed

- SATA disk module should be able to support 2048 Mbps in the future with Mark 5B, and next-generation StreamStor board, and next-generation disk drives

- SATA module price will be about same as current module

- Expect SATA modules to be available late 2005; upgrade kit will be available from Conduant

**Mark 5B Data System**

The Mark 5B system is being designed to adhere to the VSI-H specification [Ref 2]. It is particularly attractive for new stations as it directly accepts sampler data without the need for an expensive external formatter. The Mark 5B system is being designed with the following characteristics:

- Full VSI compatibility

- Same chassis as Mark 5A; uses same disk modules; requires Mark 5B I/O card

- 1024 Mbps record/playback

- Plans for expansion to 2048 Mbps with next-generation StreamStor card using 64 MHz data clock

- Full e-VLBI capability

- Eliminates need for external formatters

- With a 14-BBC Mark4 or VLBA4 system, up to 1792 Mbps can be recorded with two parallel Mark 5B systems

- Requires sampler adapter for Mark 4/VLBA DASs to provide VSI-compatible output [but will not be needed when digital BBC's become available in the future]

- Station Unit capabilities for connection to Mark 4 correlators is being designed into Mark 5B

- Built-in phase-cal extraction and state counting

- Front-panel status display – 8 tri-color LEDs

- DIM and DOM capabilities are separate FPGA downloads

- FPGA is in-place programmable via software

**Mark 5A/Mark 5B Compatibility**

The Mark 5A records data in a parity-stripped version of the Mark 4 or VLBA tape format and is thus highly specialized. The Mark 5B, being governed by VSI specifications, is freed from the legacy constraints of magnetic tape and allows the adoption of a more logical data format on disk. Hence, the disk-data format of Mark 5A and Mark 5B are different. However, the Mark 5B data format has been designed in such a fashion that the Mark 5A can be upgraded to read Mark 5B data and produce VLBA-tape format data. This will allow correlators with only Mark 5A systems to still read and correlate Mark 5B data.

The upgrade of the Mark 5A system to read Mark 5B disks is a significant engineering task and can be avoided if correlators gradually upgrade their playback units to Mark 5B in step with the observing stations. It now appears likely that this may happen and the engineering expenditure to upgrade the Mark 5A system can be avoided.

**Mark 5B Station Unit Capabilities**

As indicated earlier, the Mark 5B playback system will support a mode of operation that emulates the Mark 4 Station Unit and allows the current Mark 4 Station Units on the Mark 4 correlators to be retired. Figure 3 shows a simplified block diagram of the Mark 5B Station Unit capability. A station-delay model is applied to the data during playback to the correlator; the delay model, as well as model phase and delay-rate are also inserted into the data stream and sent to the correlator.

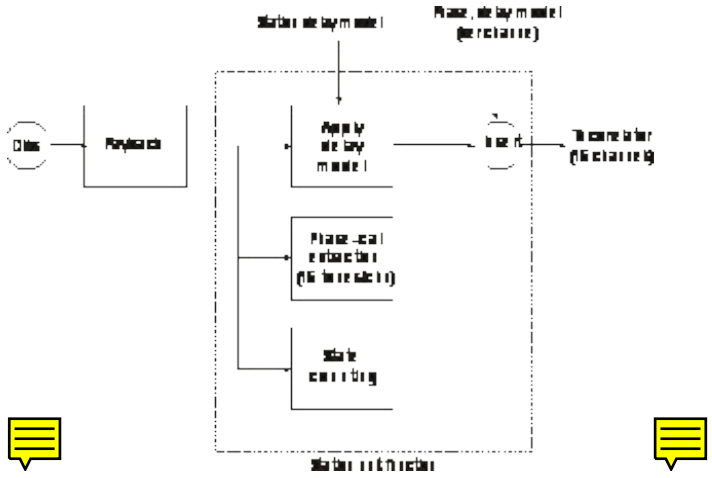

Figure 3. Simplified block diagram of SU function of Mark 5B

A built-in phase-calibration tone extractor will analyze the amplitude and phase of up to 16 individual tones in each of up to 16 channels of data, and will also measure sample-state statistics to aid in proper correlation-amplitude normalization.

**Mark 5B Status**

As of this writing (May 2005), 3 prototype Mark 5B systems are under test. Checkout of the system is a lengthy procedure, particularly for the complicated Station Unit functionality. Though it is difficult to predict, we expect that checkout and interfacing to the Mark 4 correlator will require several months of effort.

Of course, considerable software development is also being undertaken to support Mark 5B, particularly with respect to support of the Mark 5B on the Mark 4 correlator.

Upon completion of the Mark 5B checkout, the design will be transferred to Conduant Corp for replication and sale. We expect the Mark 5B will be somewhat more expensive than the Mark 5A, but not substantially so.

Existing Mark 5A systems will be upgradeable to Mark 5B with the procurement of a Mark 5B I/O interface board to replace the Mark 5A I/O interface board. Of course, new support software will also be required, which will be available from the Haystack web site.

**Mark 5B Interface Boards**

A number of specialized interface boards are required to support the Mark 5B in a world largely populated by the remnants of legacy Mark 4 and VLBA data systems. In particular, four interface boards will be made available to ease the integration of the Mark 5B into the existing VLBI world:

- VSI-4 sampler adapter board
    - Will be mounted inside Mark 4 formatter to provide VSI interface to Mark 5B
    - Uses existing Mark 4 samplers and 1pps generator

- Supports 2 VSI output connecters at 1024Mbps each (though aggregate data rate is restricted to 1792 Mbps by availability of only 14 BBC's)
- Prototype boards ready for checkout

- VSI-C sampler adapter
  - Interfaces VLBA samplers to VSI
  - Designed at Metsahovi; purchased from Metsahovi
- Correlator Interface Board (CIB)
  - Interface between Mark 5B and Mark 4 correlator
  - Currently in fabrication
- Upgraded Serial Link boards for Mark 4 correlator
  - Designed at MPI
  - Prototypes have been tested; replication underway at MPI

**Summary**

The Mark 5A and Mark 5B disk-based systems represent a new direction in VLBI data technology and have opened new capabilities for higher performance and lower cost than previous tape systems. Because these systems are based on technology that is rapidly developing due to the demands of the global marketplace, we can expect that performance and capabilities for VLBI will dramatically improve in the coming years.

**References**

[1] Mark 5 information available at http://web.haystack.edu/mark5/Mark5.htm.

[2] VSI information available at "VLBI Standard Interface Specification – VSI-H," August 2000, Revision 1.0, available at http://web.haystack.edu/vsi/index.html.