Mark 5 Overview

Alan Whitney
MIT Haystack Observatory

Mk1

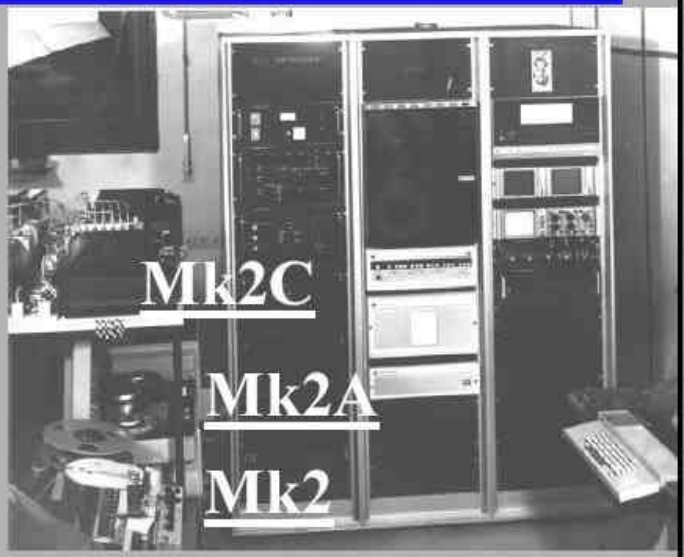

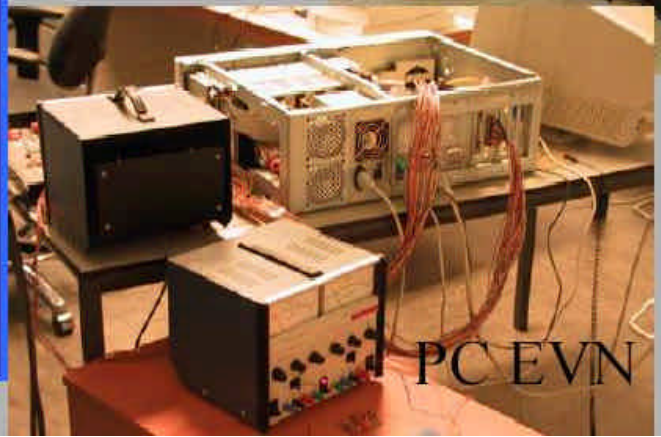


K3


K4/VSOP


PCVLBI

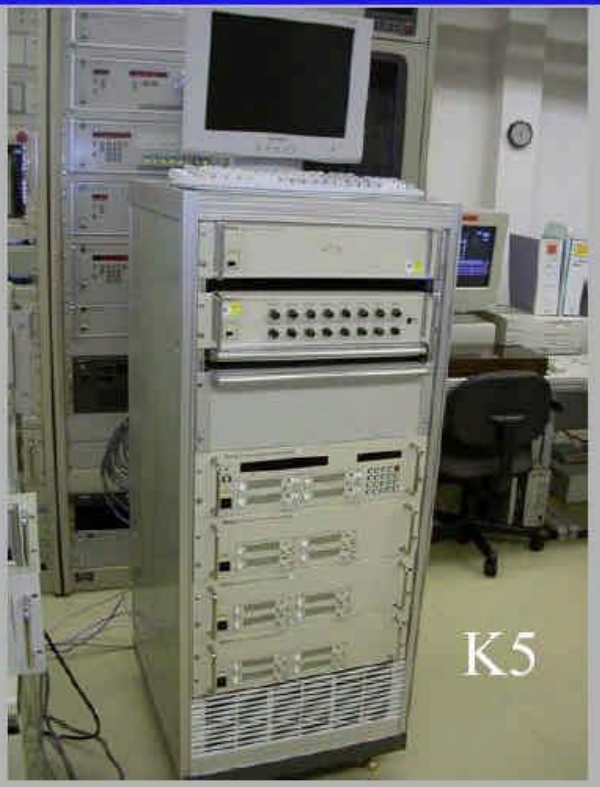
K5


K1


GBR2000


VDR2000

Goals of Mark 5 VLBI Data System

- Low cost
- Based primarily on unmodified COTS components
- Modular, easily upgradeable
- Robust operation, low maintenance cost
- Easy transportability
- Conformance to VLBI Standard Interface specification (VSI)
- Compatibility with existing VLBI systems
- Flexibility to support e-VLBI
- Minimum of 1 Gbps data rate
- 24-hour unattended operation at 1 Gbps

Tape vs. Disc Price Comparison

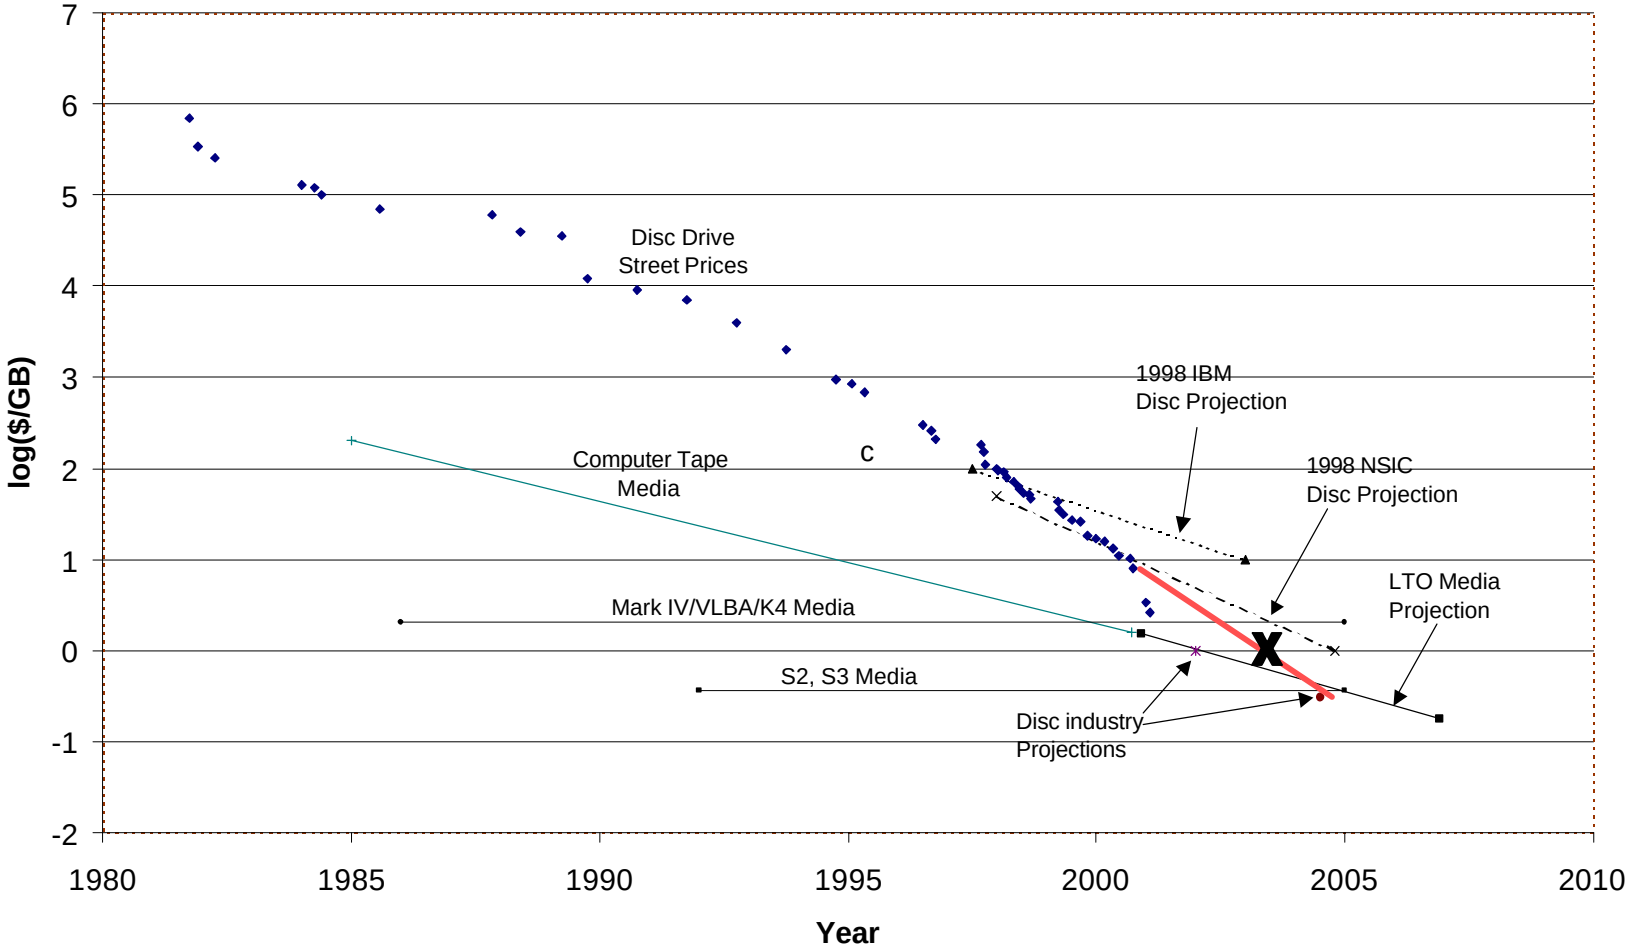

Mark 5 VLBI Demonstration System – March 2001

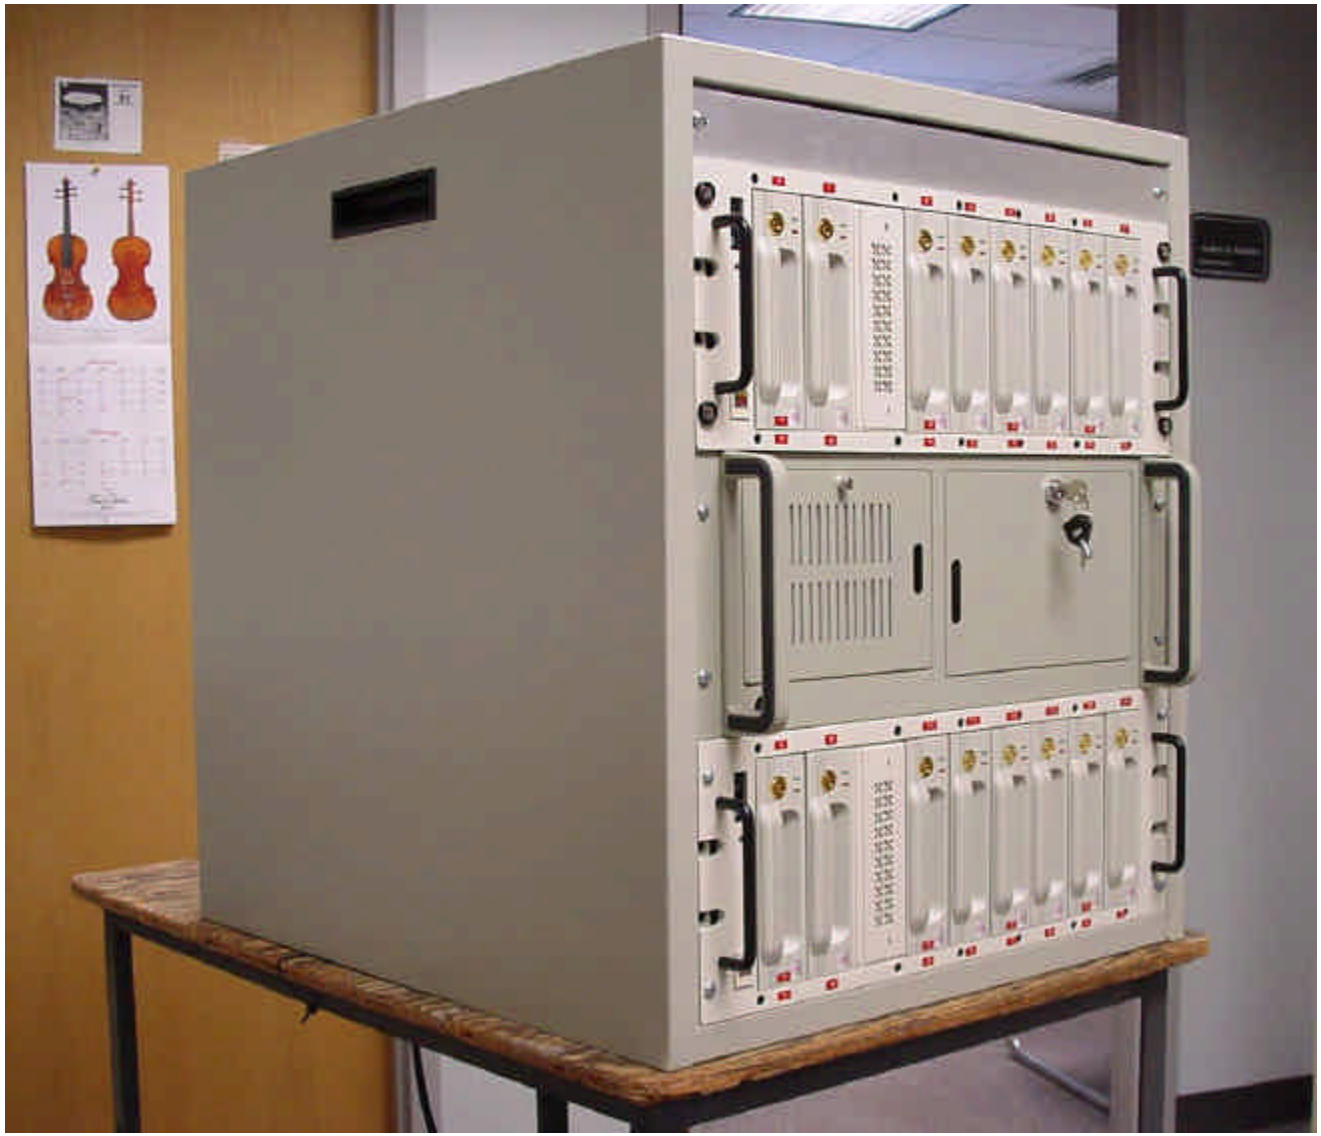

3 months start to finish!

Mark 5P Data System – mid 2002

- Direct replacement for 32-track Mark4 or VLBA tape drive
- Records 32 tracks from Mark4 formatter (1024 Mbps max) or VLBA formatter (512 Mbps max)
- Playback at any rate up to 1024 Mbps
- Uses 1 to 16 disks depending on data rate & capacity needs

BUT

- Large numbers of individual disks are nearly unmanageable!

Mark 5A ‘8-pack’ chassis - August 2002

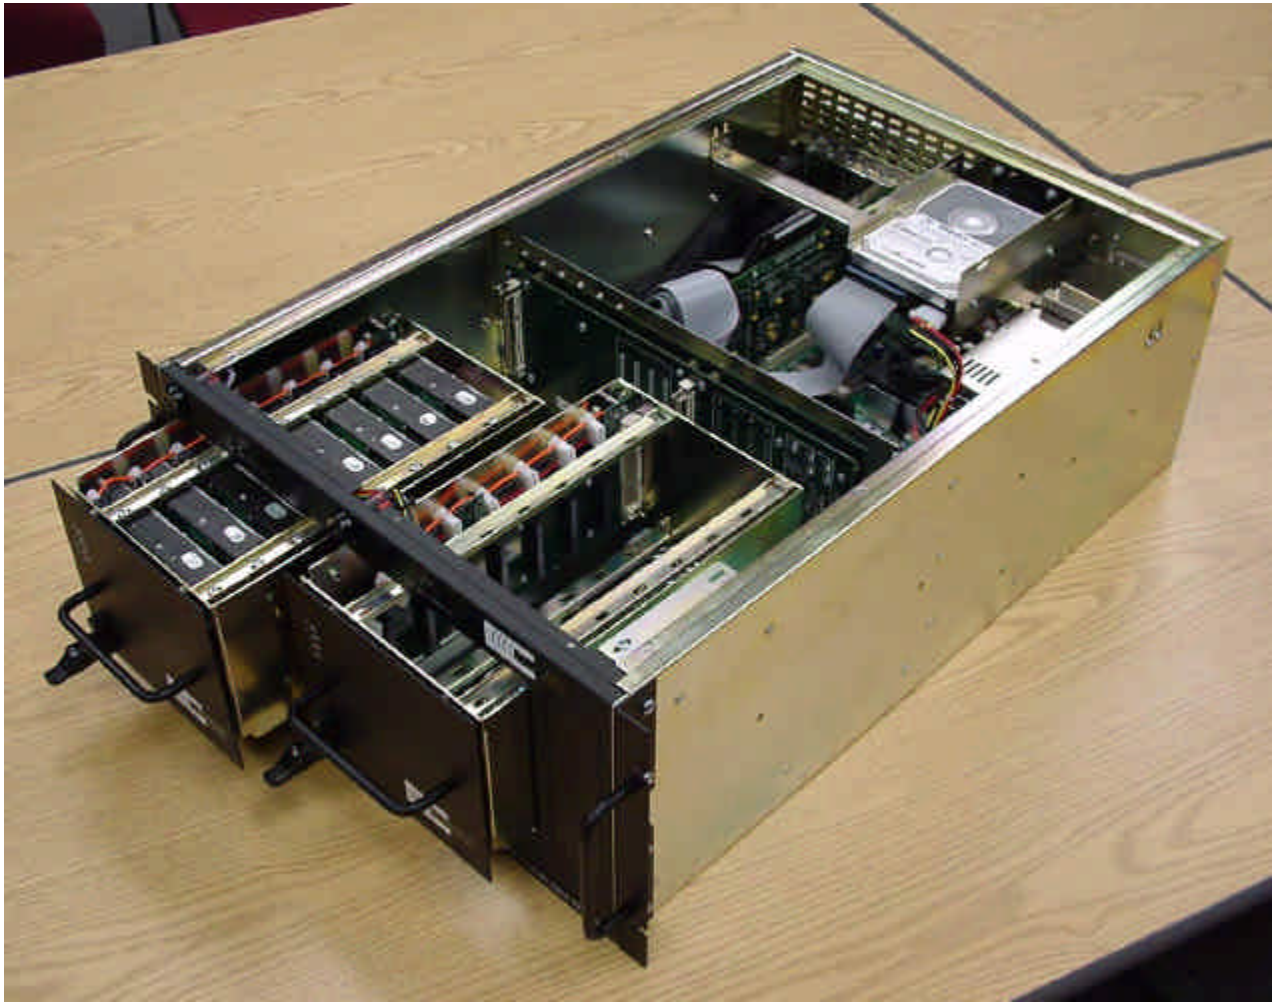

Mark 5 ‘8-pack’ disk module

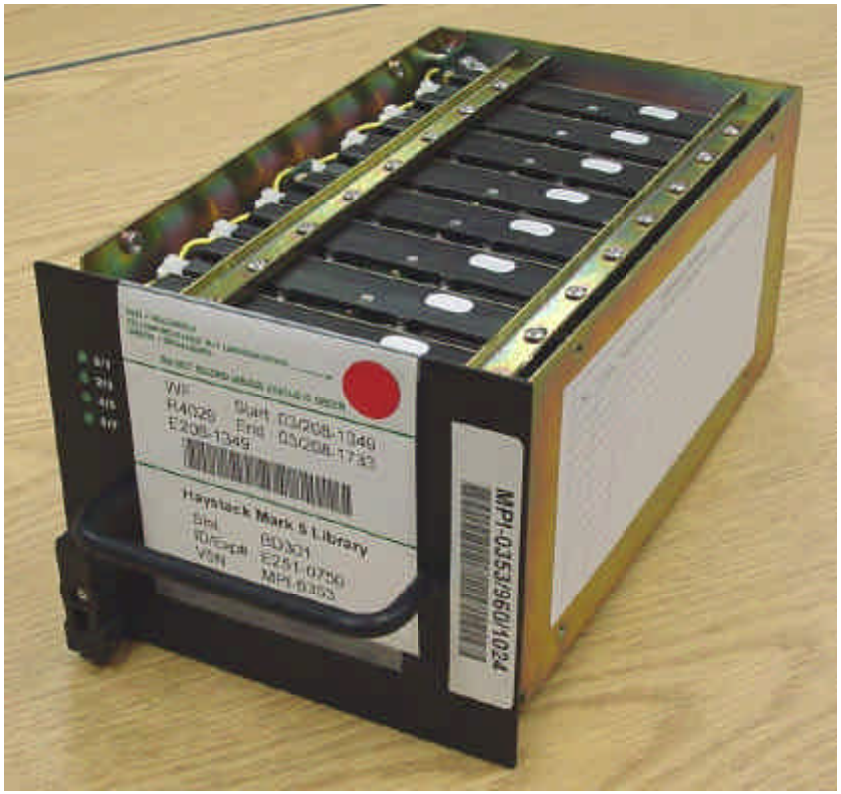

- Simplifies disk-handling logistics by keeping all disks together as a single unit
- Each module can be managed just like a Mark4/VLBA tape, using nearly identical management and tracking software

Mark 5A Data System

- Direct plug-compatible replacement for 64-track Mark4 or VLBA tape drives
- Record/Playback at rates up to 1024 Mbps
- Two independent ‘8-pack’ disk modules per system can be used in ‘ping-pong’ fashion
- 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; re-inserted on playback
- Arbitrary mixing of modes (#tracks, data rate, bits/sample) is allowed, always using 100% of installed disk capacity

Mark 5A I/O Panel

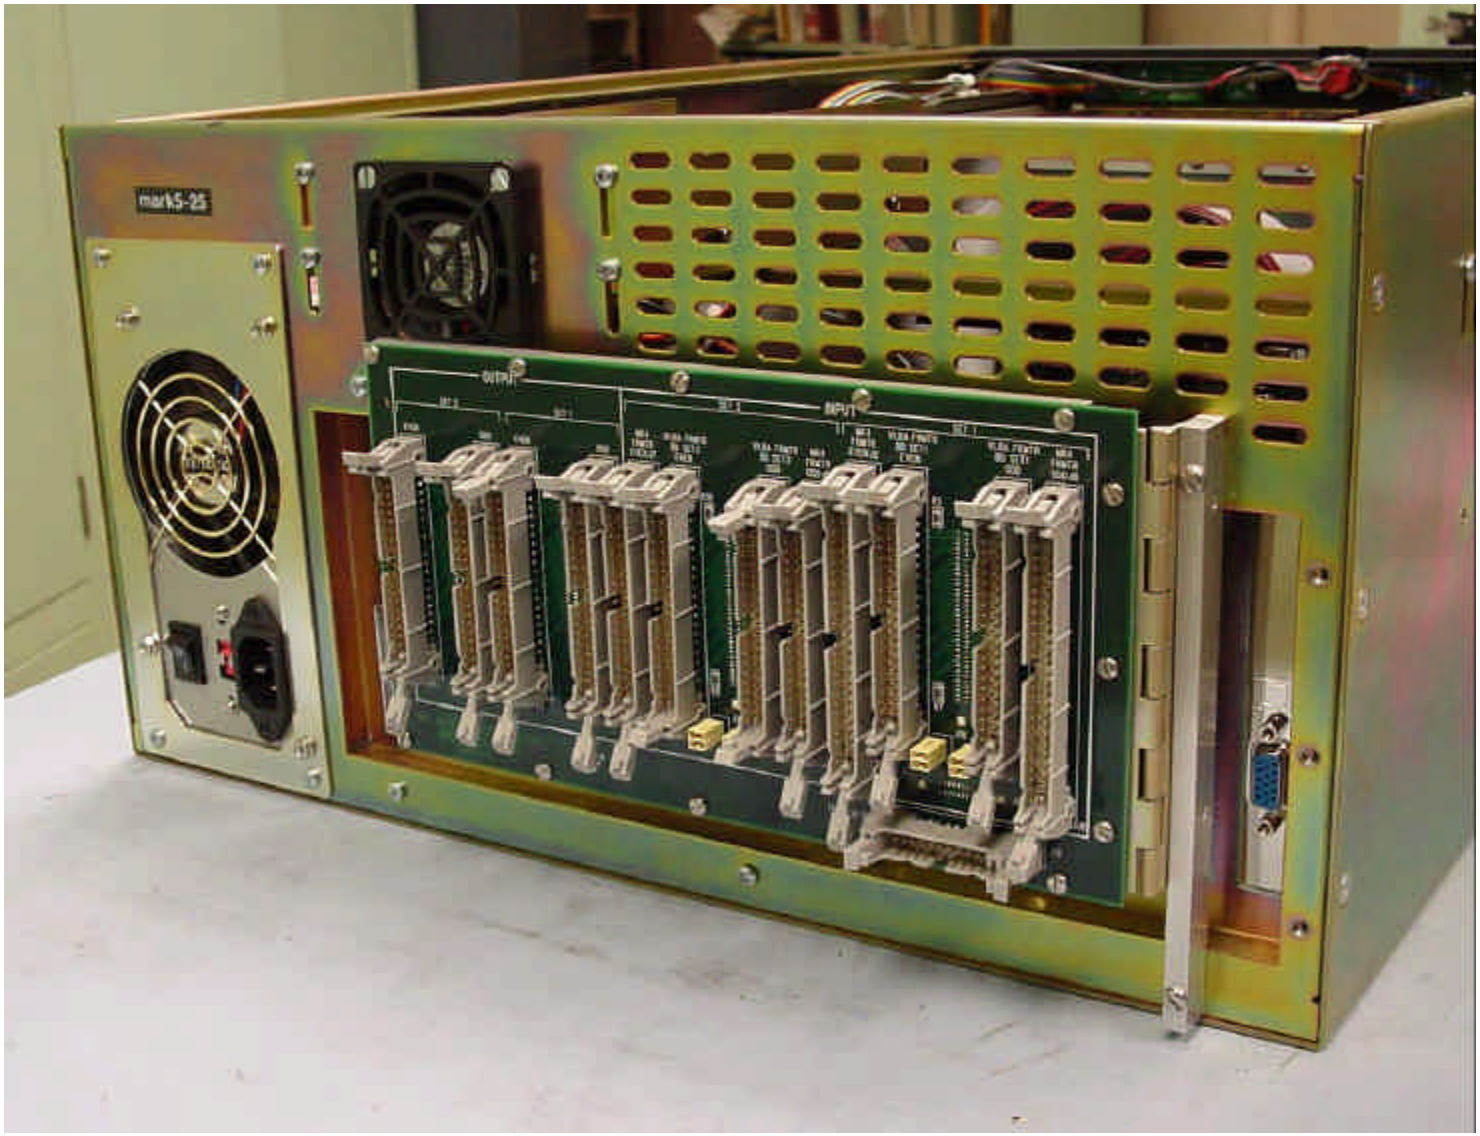

Current Mark 5 Status

- ~40 Mark 5 systems deployed
- ~250 Mark 5A ‘8-pack’ disk modules now in use
- Daily Intensive UT1 observations Wettzell-Hawaii have been exclusively Mark 5 for 1 year with almost no problems
- Westford recorded 15-day CONT02 experiment entirely on Mark 5
- Hobart now entirely Mark 5A
- Several astronomy experiments have now successfully used Mark 5A at rates to 1 Gbps, including a recent very successful mm-wavelength experiment
- A number of e-VLBI experiments have been conducted with Mark 5A
- Current Mark 5A price ~$16.3K from Conduant Corp

Simplified Mark 5A Block Diagram

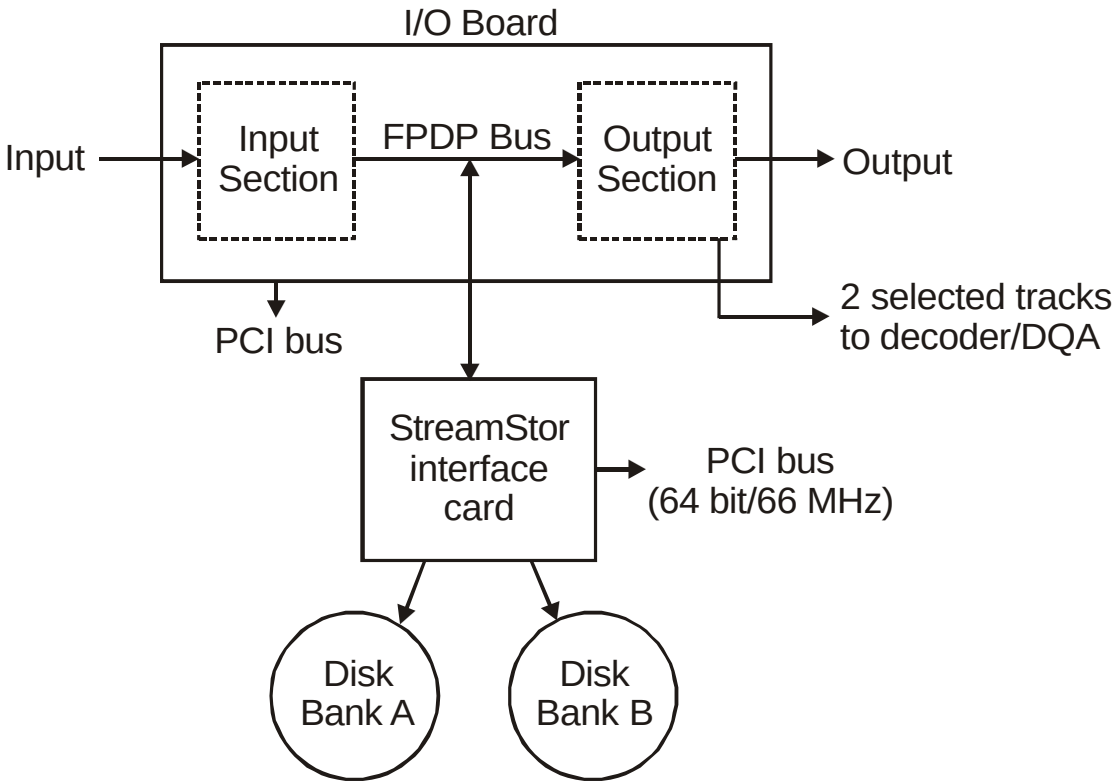

Note:

- Disk Banks A and B operate independently (but not at the same time)
- During recording, the disk interface simply ‘listens’ on the FPDP bus and records
- During recording, the Output Section is active, so the Input is actively duplicated at the Output
- When playing back, the disk data are put back onto the FPDP in the same form they were recorded.

Mark 5 Connectivity

- Mark 5 supports a triangle of connectivity for recording and e-VLBI

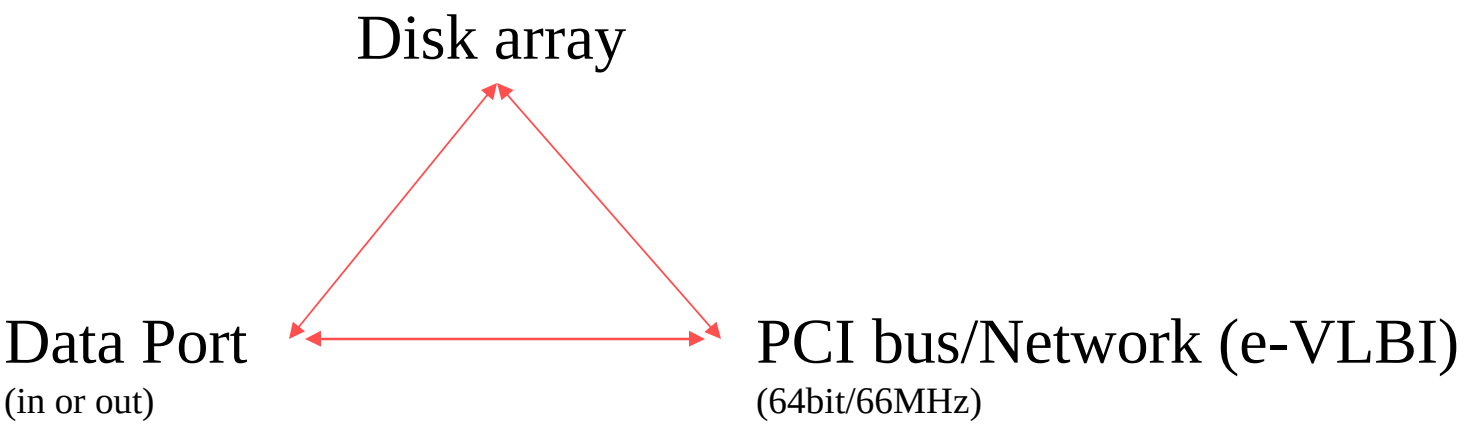

Mark 5 can support several possible e-VLBI modes:

- e-VLBI data buffer (first to Disc Array, then to Network); vice versa
- Direct e-VLBI (Data Port directly to Network); vice versa

Mark 5A Basic Operating Modes

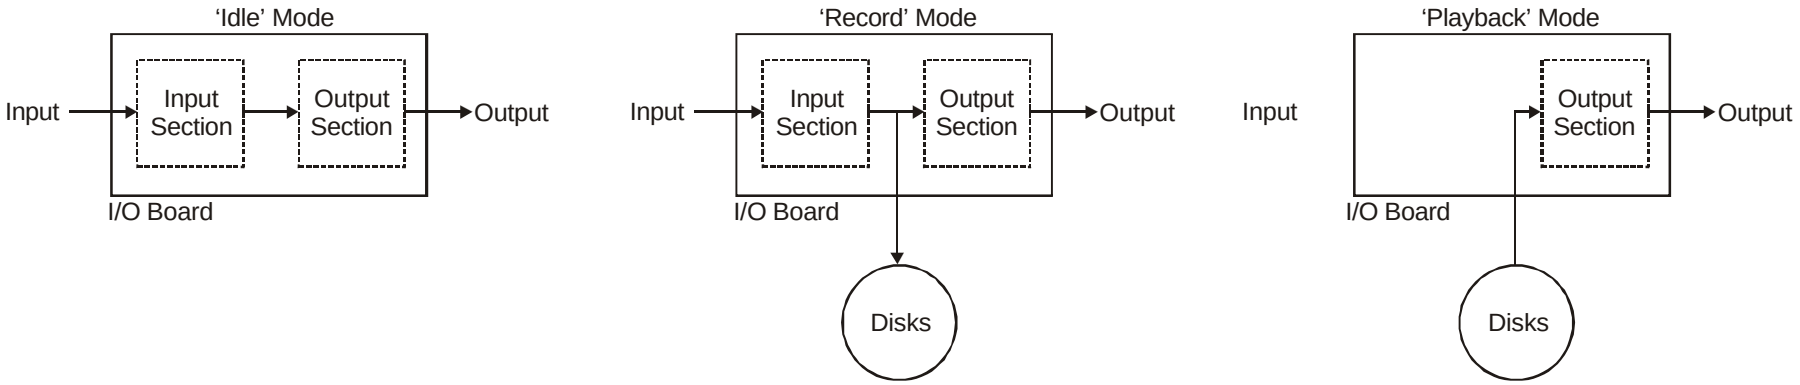

Idle Mode

- Input is actively duplicated at the Output

Record Mode

- Same as Idle Mode, except disks record what passes across FPDP bus

Playback Mode

- Disk data are re-played onto FPDP bus into Output Section

Mark 5A Basic Data Modes

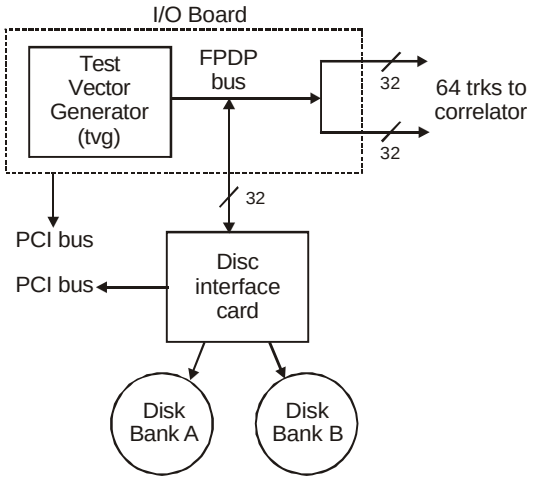

'tvg' Data Mode

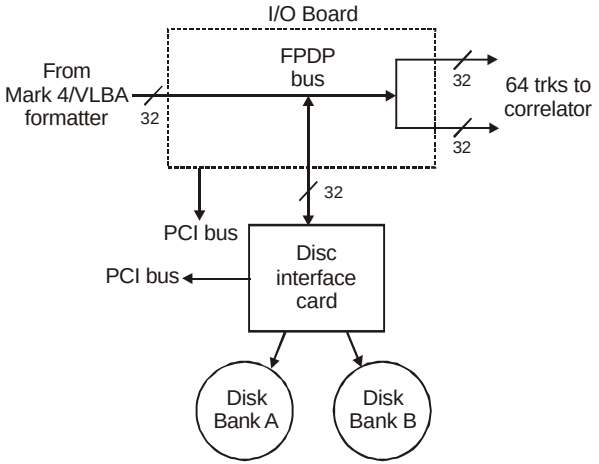

'st' Data Mode

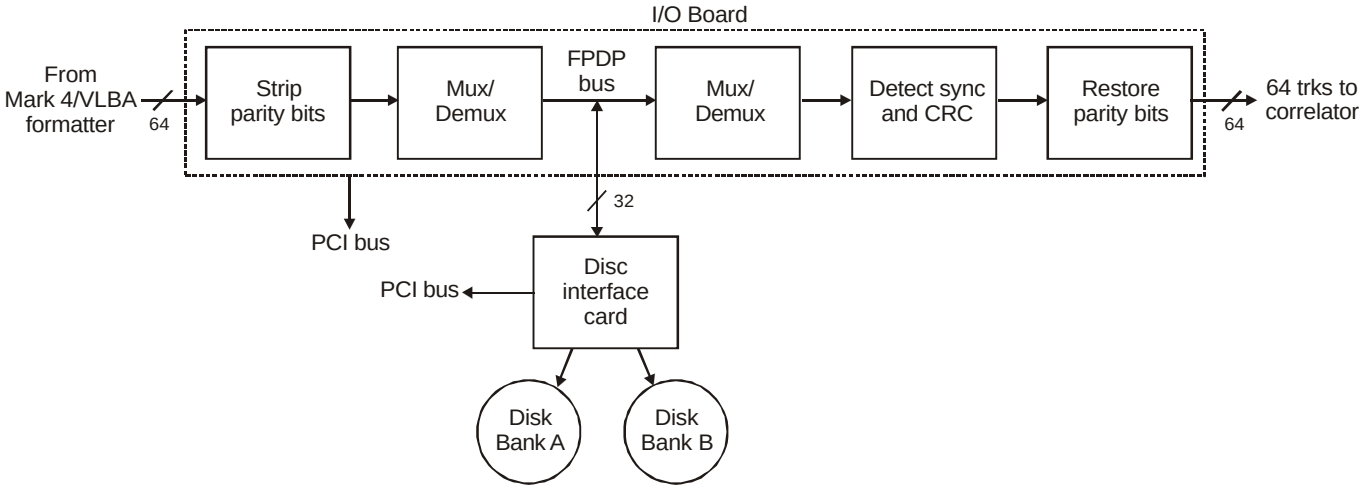

'mark4' and 'vlba' Data Modes

Mark 5A Data Modes

Data Mode	Active input trk#'s	Comments
'mode=st;'	2-33 all (Set 1)	Parity not stripped ('Mark 5P' mode)
'mode=tvg;'	Data taken from internal tvg	No data massaging; Uses on-board clk gen
'mode=mark4:8;' 'mode=vlba:8;'	2-17 even (Set 1)	Parity-stripped on record; restored on playback
'mode=mark4:16;' 'mode=vlba:16;'	2-33 even (Set 1)	Parity-stripped on record; restored on playback
'mode=mark4:32;' 'mode=vlba:32;'	2-33 all (Set 1)	Parity-stripped on record; restored on playback
'mode=mark4:64;' 'mode=vlba:64;'	2-33 (Sets 1 and 2)	Parity-stripped on record; restored on playback

Note:

- When recording in 'mark4' mode, Mark 4 rack-ID must be set to even value

Reliability Features

- Recording:
  - All channels are always distributed over all recorded disks; no need for barrel-roll or data-modulation
  - Loss or substandard performance of a disk on record is compensated for automatically by transferring load to other disks
- Playback
  - Missing data blocks due to slow or bad disks are replaced by ‘even parity’ data that is invalidated by correlator
  - Individual disk-performance statistics are kept to sort out marginal or failing disks
    - Transaction-time statistics on each individual disk
    - # of replaced blocks on each individual disk

Mark 5 Data-recovery Example

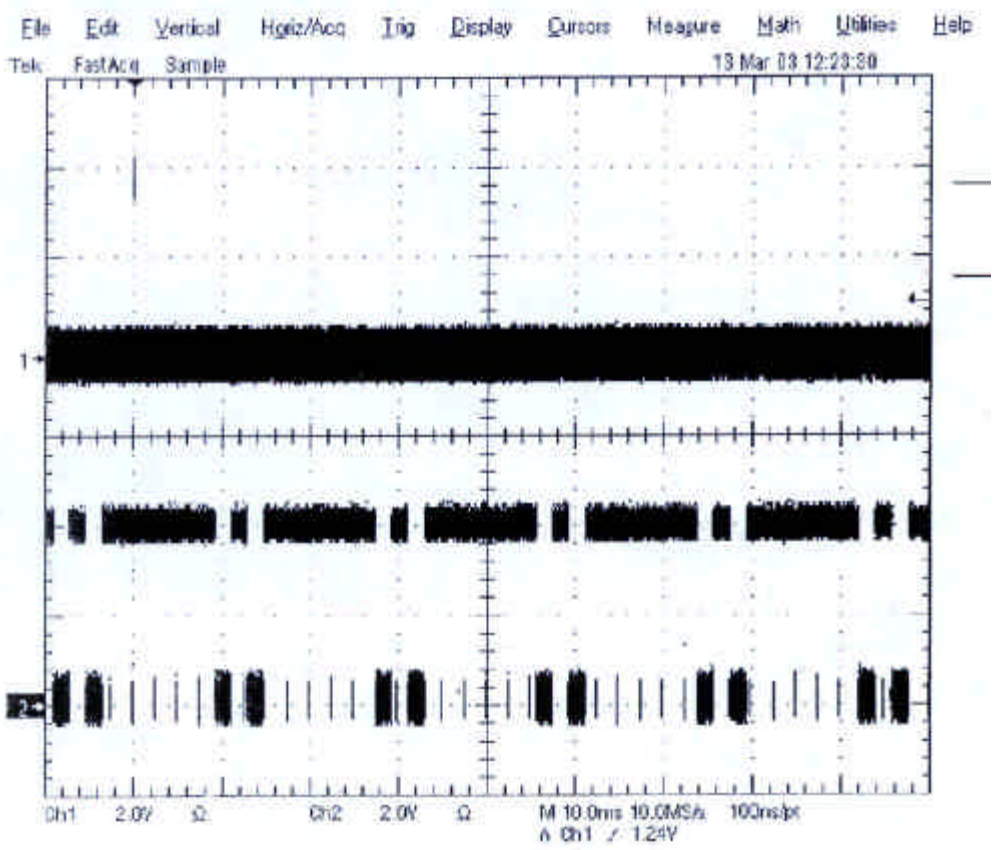

CONT02 experiment, Oct 2002

Westford antenna

2 of 10 disks missing

~77% data recovered

Requires careful correlator setup

Disk-Media Status

- Hard disk price vs capacity/performance will continue to drop rapidly
  - Now ~$1.00/GB, expected to drop to ~$0.50/GB by ~2005
    (Mark 4/VLBA tape is ~$2.00/GB)
- 250 GB disks now available – ~35 hours @ 256 Mbps unattended
- (comparable to ~7 VLBA/Mark 4 tapes)
- 320 GB disks expected soon – ~22 hours @ 512 Mbps unattended
  (comparable to ~9 VLBA/Mark 4 tapes)
- 700 GB disks expected ~2005 – ~24 hours @ 1 Gbps unattended
  (comparable to ~19 VLBA/Mark 4 tapes)

IBM/Hitachi Disks currently recommended

- We have recently learned a significant difference between WD and IBM/Hitachi disks
  - Powered-down 3.5” WD disks leave head in contact with surface (at either innermost or outermost diameter)
  - Powered-down 3.5” IBM/Hitachi disks move heads off surface onto a ‘ramp’, offering better protection to the heads and the surface.
- Apparently, 2.5” disk drives used in notebook computers use a similar ‘ramp’ parking scheme for ruggedization
- Head-parking schemes used by other manufacturers is not known
- 180GB disks are being phased out in favor of new 250GB disks

Disk-module handling and management

- Disk-modules must be carefully assembled and tested carefully according to instructions on Mark 5 website
- Disk-modules are managed the same as tape, except:
  - Erasure is done at station before recording
  - Switching from one module to next on recording is now usually done dynamically and automatically by Field System
- Disk modules should be erased and conditioned at station before use
- Labeling conventions should be carefully followed; unerased disks will be shipped stations with yellow dots (released, but unerased)
- Module tracking is done by entering into ‘track’, same as tape

Disk-module VSN's

- Each module is assigned a permanent 8-char VSN just like tape
- VSN format: Owner-ID followed by '-' followed by serial number
  - Example: ABC-0132 ('-' identifies as disk module)
  - Posted on rear-panel bar-code label for use with 'track' program
  - 2-6 char owner-ID must be registered with Jon Romney at NRAO (jromney@nrao.edu); can also be coordinated through Nancy Vandenberg
- Extended-VSN format: 'VSN/GB/max-data-rate'
  - Example ABC-0132/960/1024 (960GB, 1024Mbps)
  - Posted on front-panel bar-code label, which 'track' will eventually use
  - Extended-VSN is written into protected place on all disks unaffected by erasure or conditioning; extended-VSN is read by field system and recorded in log

Disk Conditioning

- Why conditioning?
  - Disks are not fully tested at factory
  - Bad sectors encountered during playback can cause loss of data
  - Conditioning checks all sectors and re-assigns bad sectors to spares
  - Sectors can go bad during shipping due to junk floating around inside disk drive; disks will tend to stabilize after a few conditionings
- Condition is done using standalone program ‘SSErase’

  ‘SSErase –m0 –c1 &’
  - Can take several hours depending on disk sizes and number of problems found (example: ~5 hrs for 8 * 200GB)
  - Examine disk statistics printed at end of conditioning to see if any disks stand out as being different from others
  - Please note successful conditioning on module side label

Other Considerations

- Cabling: High-quality cabling between the formatter and the Mark 5A is a requirement; replace cables if necessary with new high-quality cables of minimum length
- Disk modules should not be handled on hard surfaces
  - Very easy to subject disk to quite high shocks
- For shipping: Shipping covers should be installed to protect disks from debris and, perhaps, prying eyes
- Use special shipping boxes

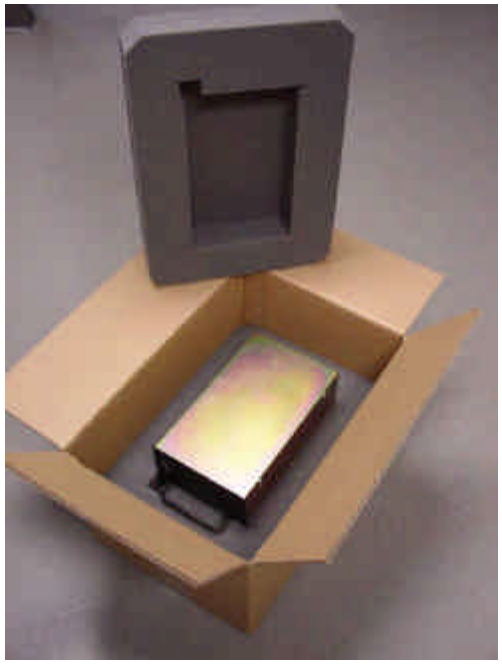

Mark 5 Software Status

- Nearly all planned features are now supported with the exception of automatic bank switching
- All software/firmware is available on-line at Mark 5 web site
- Mark 5A systems are now supported on all Mark 4 correlators
  - USNO
  - MPI
  - Haystack
  - JIVE

Documentation

- Extensive on-line documentation at www.haystack.edu/mark5
  - Mark 5A Users Manual
  - Mark 5A Test Procedures
  - Assembly and Test of Modules
  - Disk-module management and handling
  - Mark 5A command set
  - Mark 5A FAQ and tutorials
  - Software updates
  - Mark 5 memo series
  - Mark 5 newsletter series

Serial Signaling in Computers

Hard Drives

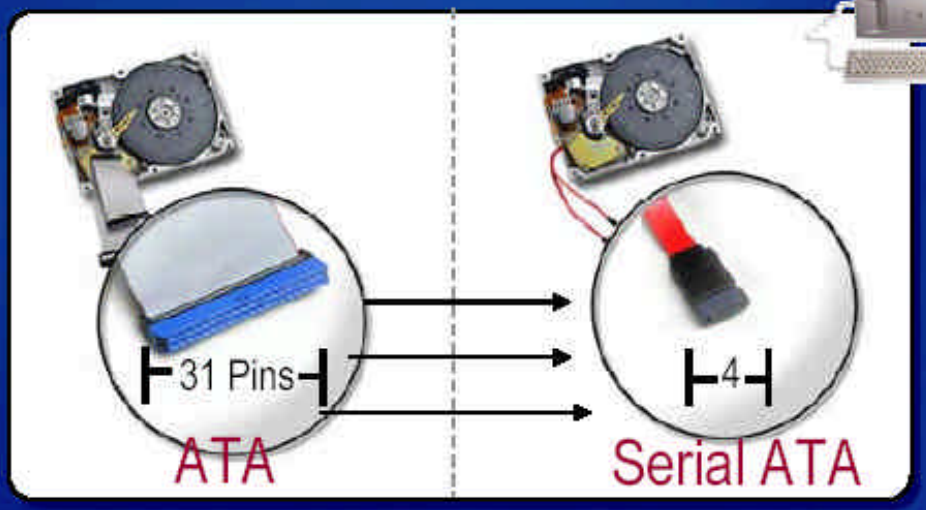

Motherboards

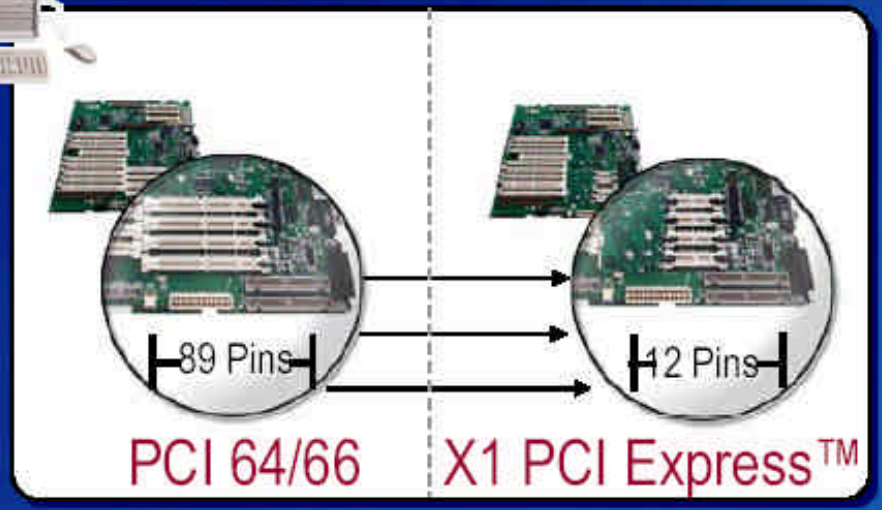

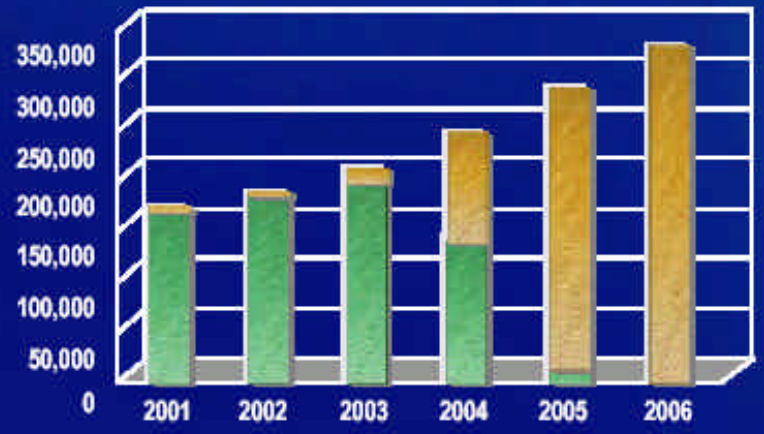

Serial	0.0%	0.3%	7.3%	42.8%	95.6%	100.0%
Parallel	00.0%	99.7%	92.7%	57.2%	4.4%	0.0%

XILINX®

Plans for Serial-ATA Support

- PATA and SATA disk modules will be interchangeable in Mark 5 chassis
- How this will be done:
  - New SATA disk module with connector on R (right) side of module (connector on PATA module is on L)
  - Ejector lever moved to R (on L on PATA module)
  - New chassis backplane
  - New sheet metal piece to accept module with ejector lever on R or L
  - No other changes
- Expected upgrade price ~$1000
- SATA module price expected to remain about same
- Hope to complete by end 2003/early 2004

VLBI Standard Interface (VSI)

- Joint effort by astronomy (GVWG) and geodetic (IVS) communities
- **The purpose of VSI is to define a standard interface to and from a VLBI Data Transmission System (DTS) that allows heterogeneous DTS's to be interfaced to both data-acquisition and correlator systems with a minimum of effort.**
- Focuses on those functions independent of DTS technology
- VSI-H – complete
  - Data and control interfaces for recording and playback, including connectors and pinouts
  - Electrical and timing specs
- VSI-S – complete
  - Communications model
  - Application Protocols
  - Command/Response Syntax
  - Suggested base command set
- VSI-E – new, in progress!
  - Will define e-VLBI standards
- VSI-H and VSI-S specs available at http://www.haystack.edu/vsi

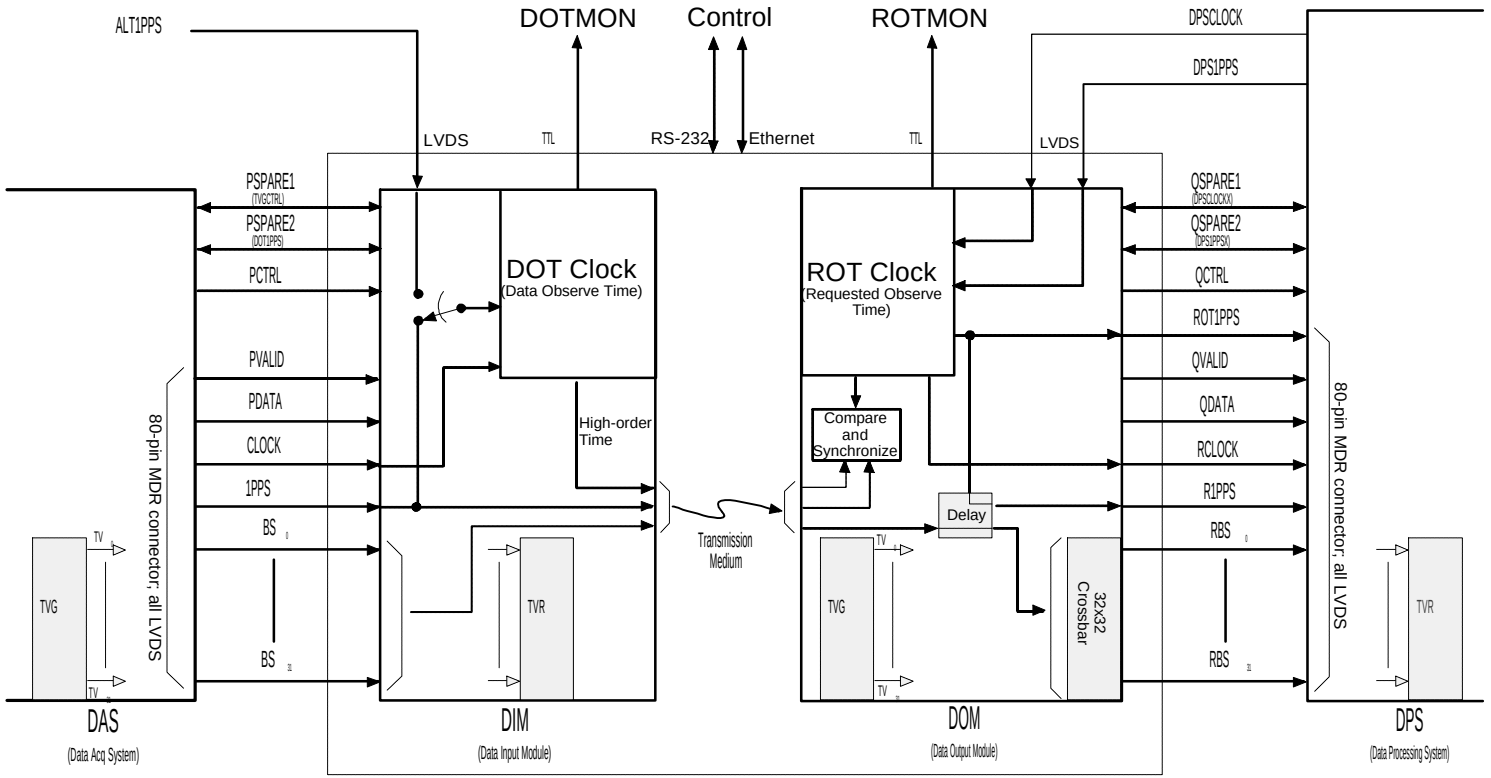

Notes:

1. Shaded items are for illustrative purposes only.
2. PVALID is optionally transmitted from DIM to DOM.
3. PDATA is optionally transmitted from DIM to DOM.
4. Data delay in DOM is required only for storage-based systems.
5. See text for discussion of use of optional use of P/QSPARE1/2 signals.
6. If DIM/DOM in single box, ALT1PPS/DPSCLOCK/DPS1PPS share single MDR-14 connector.
7. This diagram does not show all functions and options -- see VSI-H specification for details.

VSI6.DRW
ARW 21 Jun 2000

Figure 1: VSI-H Functional Block Diagram

VSI Committee

• Wayne Cannon	York University	Canada
• Brent Carlson	DRAO	Canada
• Dick Ferris	ATNF	Australia
• Dave Graham	MPI	Germany
• Ed Himwich	NASA	U.S.
• Richard Hughes-Jones	U. of Manchester	England
• Nori Kawaguchi	NAO	Japan
• Tetsuro Kondo	CRL	Japan
• Ari Mujunen	Metsahovi	Finland
• Sergei Pogrebenko	JIVE	Netherlands
• Sergey Likhachev	ASC	Russia
• Jon Romney	NRAO	U.S.
• Ralph Spencer	Jodrell	England
• Harro Verkouter	JIVE	Netherlands
• Alan Whitney (chair)	Haystack	U.S.

Mark 5B Data System

- Full VSI (VLBI Standard Interface) capability
- Up to 1024 Mbps
- Requires new Mark 5B I/O card (currently being designed)
- Eliminates need for Mark 4 or VLBA formatter
- Same chassis as Mark 5A
- Will include Station Unit capabilities for connection to Mark 4 correlators
- Will also provide adapter for Mark 4/VLBA DASs to provide VSI-compatible output
- Expect Mark 5B to be ready early-mid 2004

Mark 5B Compatibility Matrix

			Playback					
			System →	Mark 5A		Mark 5B		
Record	System ↓	Data input ↓	Play format →	Mark 4	VLBA	VSI	Mark 4	VLBA
	Mark 5A	Mark 4		ü				
	Mark 5A	VLBA			ü			
	Mark 5B	VSI			ü	ü		

In summary:

- Mark 5B can play only Mark 5B recordings (VSI format in/out)
- Mark 5A can play
  - All Mark 5A recordings
  - Mark 5B recordings will playback in VLBA-track-format (requires upgrade to current Mark 5A)

Conclusions and predictions

- Disk-based systems are rapidly displacing tape systems
  - High-performance, high-reliability, low-cost are the main drivers
  - Expect most of the world (except VLBA!) to be disk-based by end of 2004
- VSI-compatible systems are now appearing and are likely to dominate within a few years
  - Japan has been the leader in developing and deploying VSI-compatible systems
- E-VLBI is now in infancy and is developing rapidly
  - We will hear more about e-VLBI from David Lapsley this afternoon