Mark 5C VLBI Data System

Alan Whitney
Chet Ruszczyk
Kevin Dudevoir
MIT Haystack Observatory

Walter Brisken
Jon Romney
National Radio Astronomy Observatory

17 September 2007
e-VLBI Workshop
MPI
Bonn, Germany

Mark 5 Data Acquisition System

(Mark 5A/B/B+/C all look the same)

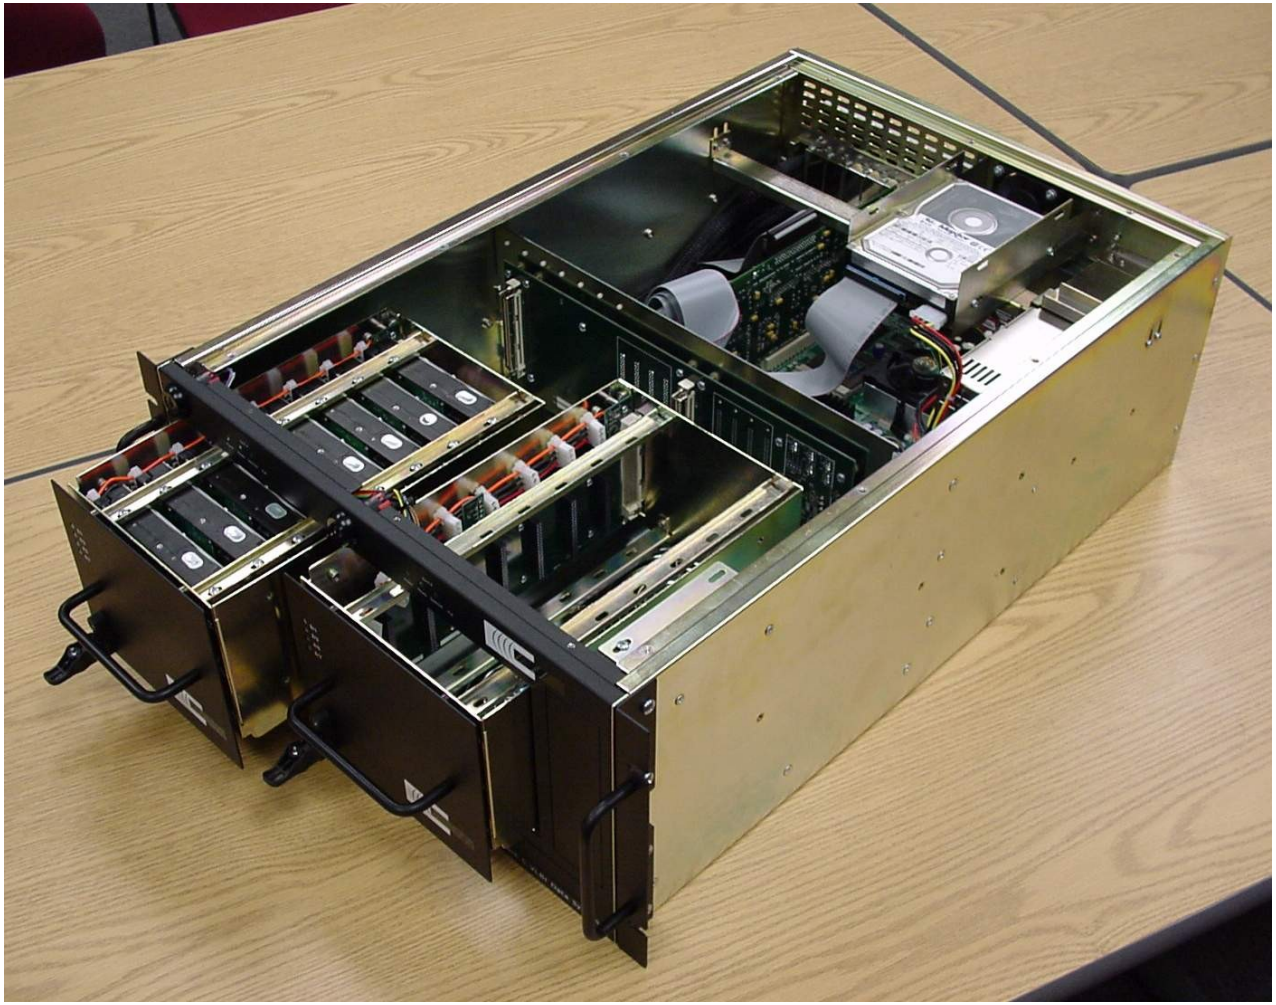

Mark 5A introduced in 2002, Mark 5B in 2005, Mark 5B+ in 2006

Mark 5A Status

- Direct plug-compatible replacement for 64-track Mark4 or VLBA tape drives
- 1024 Mbps maximum record/playback data rate
- ~150 Mark 5A units deployed to stations and correlators
- ~1300 Mark 5A '8-pack' disk modules deployed (>2 PB of storage!); growing rapidly!
- 1 Gbps experiments are now routine for both geodesy and astronomy
- Correlator efficiency has improved by factor of ~2 over tape!

Development support 2001-2004 from Mark 5 development consortium – NASA, NRAO, USNO, MPI, BKG, EVN, KVN, JPL

Mark 5B Status

- Compatible with VSI-H specification
- 1024 Mbps record/playback
- Requires VSI-H data source (dBBC, DBE1, ADSn000, VSI4, etc)
- Same chassis and disk modules as Mark 5A
- Station Unit capabilities for connection to Mark 4 correlators is built-in
- Westford and Badary stations are now full-time Mark 5B;
  Parkes will soon join
- ~30 Mk5A-to-Mk5B upgrade kits have been distributed
- Haystack, MPI and USNO correlators support Mark 5B playback
- JIVE supports playback of Mark 5B recordings on Mark 5A+;
  working on Mark 5B support

Mark 5B+ Status

- Compatible with VSI-H specification
- Same as Mark 5B, except 2048 Mbps max record; 1024 Mbps playback
- Requires upgraded StreamStor disk interface card (“Amazon”)
- Playback on Mark 5B or 5B+ (2048 Mbps plays back at 1024 Mbps max)
- Playback on Mark 5A+ not supported for 2048 Mbps recordings
- Available as standalone or upgrade to Mark 5B from Conduant Corp
- Mark 5B+ is operational and has been used in a number of successful experiments, including a large 3-station mm-VLBI experiment in April 2007 recording at 4 Gbps per station (two Mark 5B+s per station)

Mark 5C (under development)

- Input data on 10 Gigabit Ethernet (OSI Layer 2 or higher)
- 4096 Mbps max data rate to two standard Mark 5 disk modules
- Record through SS 10GigE interface
  - Requires new 10GigE daughter board for Amazon (under development); no separate "I/O board", unlike Mk5A/B/B+
- Playback will be through host computer as standard Linux data file*
  - Natural for software correlators and e-VLBI
  - Goal to replay data at full rate
- Prototype expected by ~early/mid 2008

*Conduant may provide playback path through Amazon card as well, but not mandated

Mark 5 Roadmap for Recording

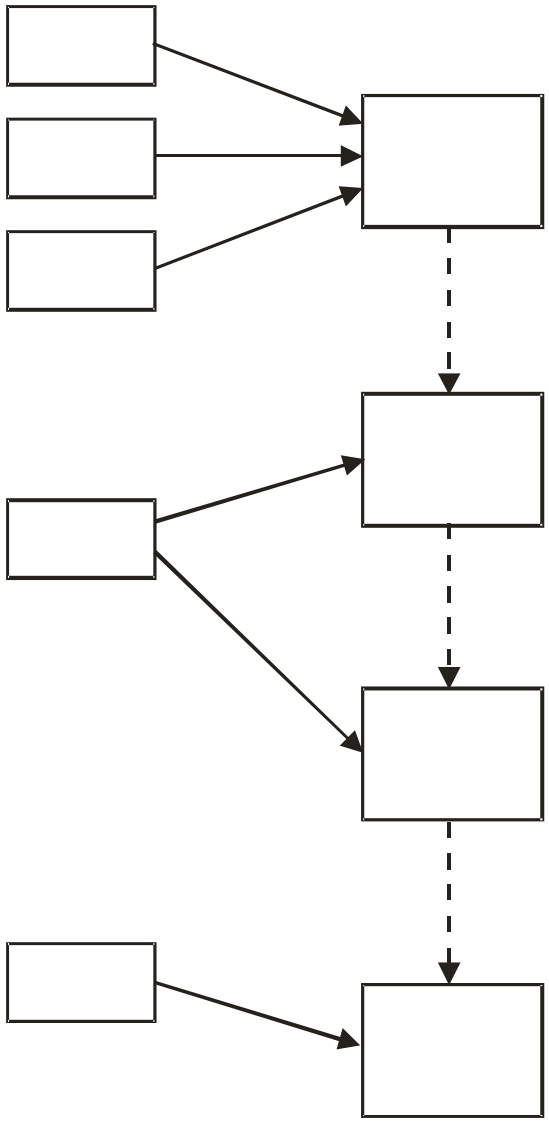

Mark 5 Roadmap for Playback

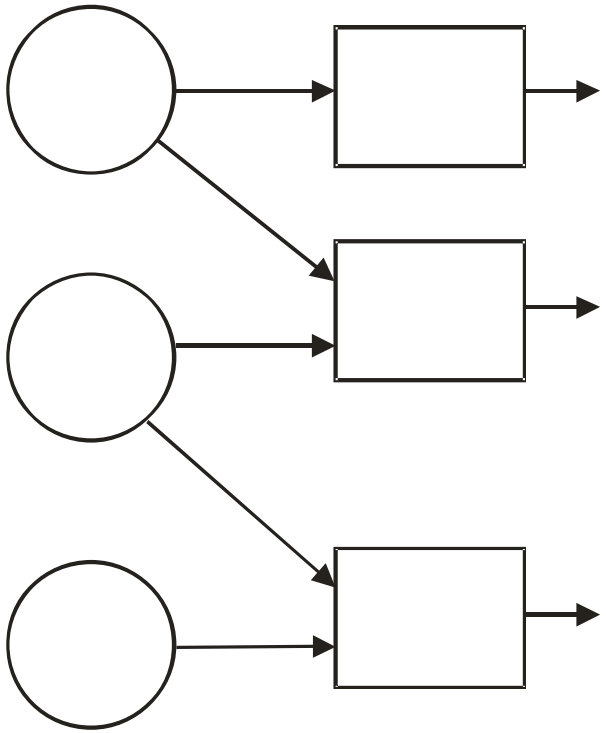

Generalized 10GigE Data Distribution Concept

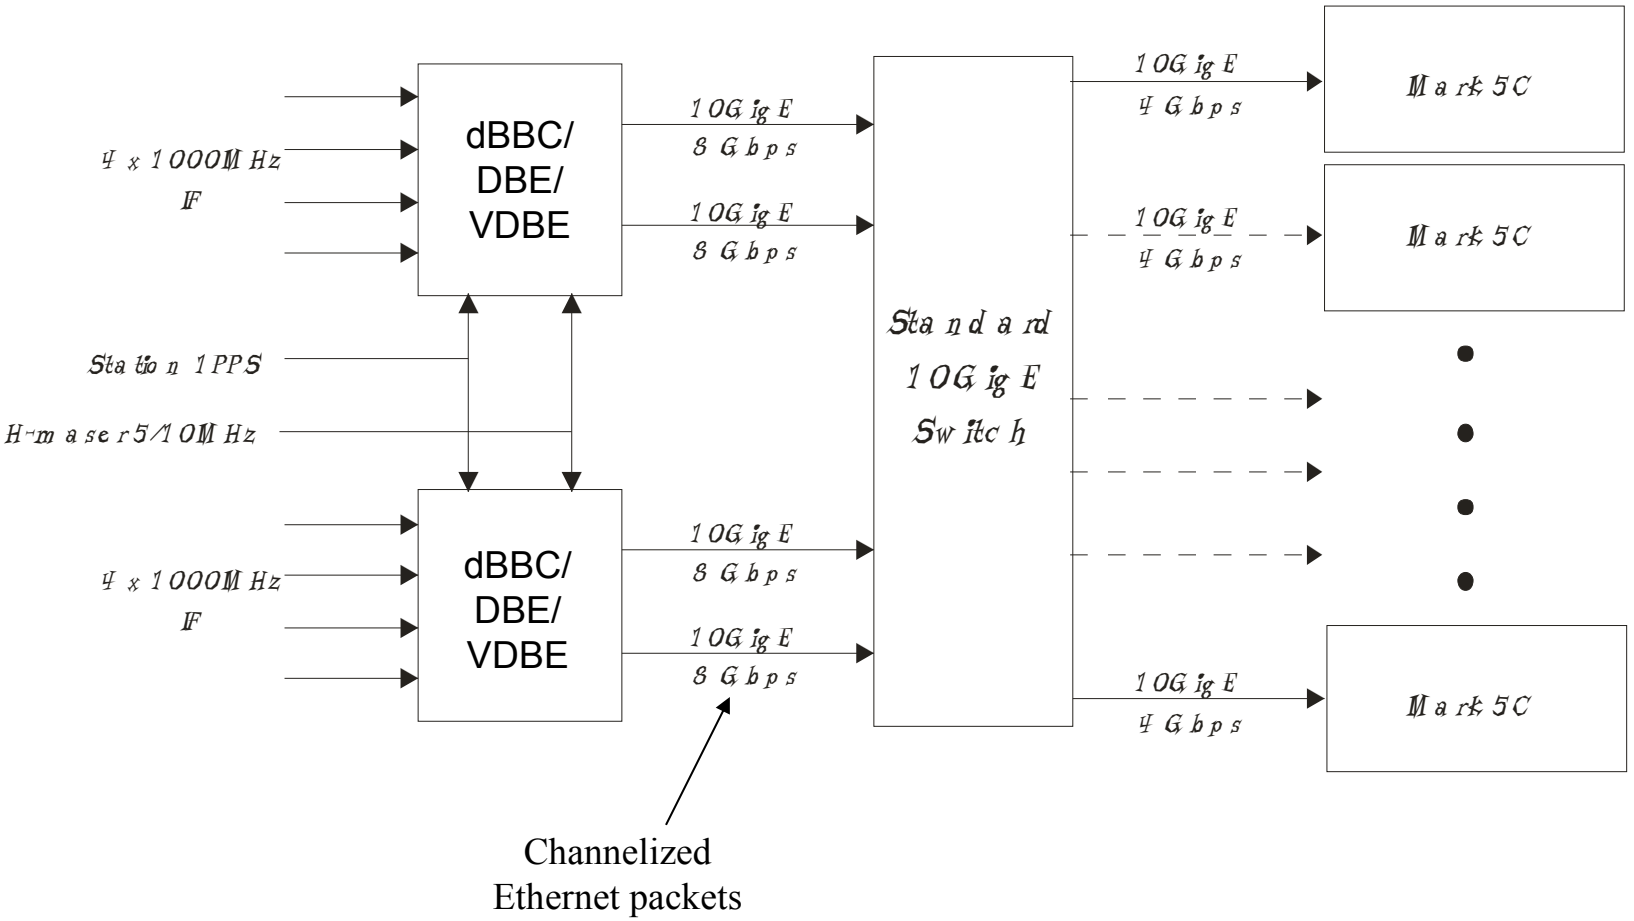

Mark 5C design basics

- Mark 5C is basically a ‘dumb’ asynchronous Ethernet packet recorder that simply records the payload of each Ethernet packet which it receives
- Will work with a ‘transmit only’ data source (i.e. no handshaking required
- User options will exist to make Mark 5C a bit ‘smarter’:
  - Ignore higher-level OSI content and other data not to be recorded
  - Monitor a user-generated Packet Serial Number to
    - Detect missing or bad packets and take certain actions
    - Prevent recording of certain marked packets
- Data source is responsible for creating Ethernet packet, including time-tagging and data-formatting
- A draft ‘VSI-G’ packet specification has been developed (but Mk5C itself doesn’t really care):
  - One freq channel per Ethernet packet
  - Supports an arbitrary # of channels
    (i.e. not constrained to 2**n channels, unlike Mk5A/B/B+)

Ethernet Packet Structure

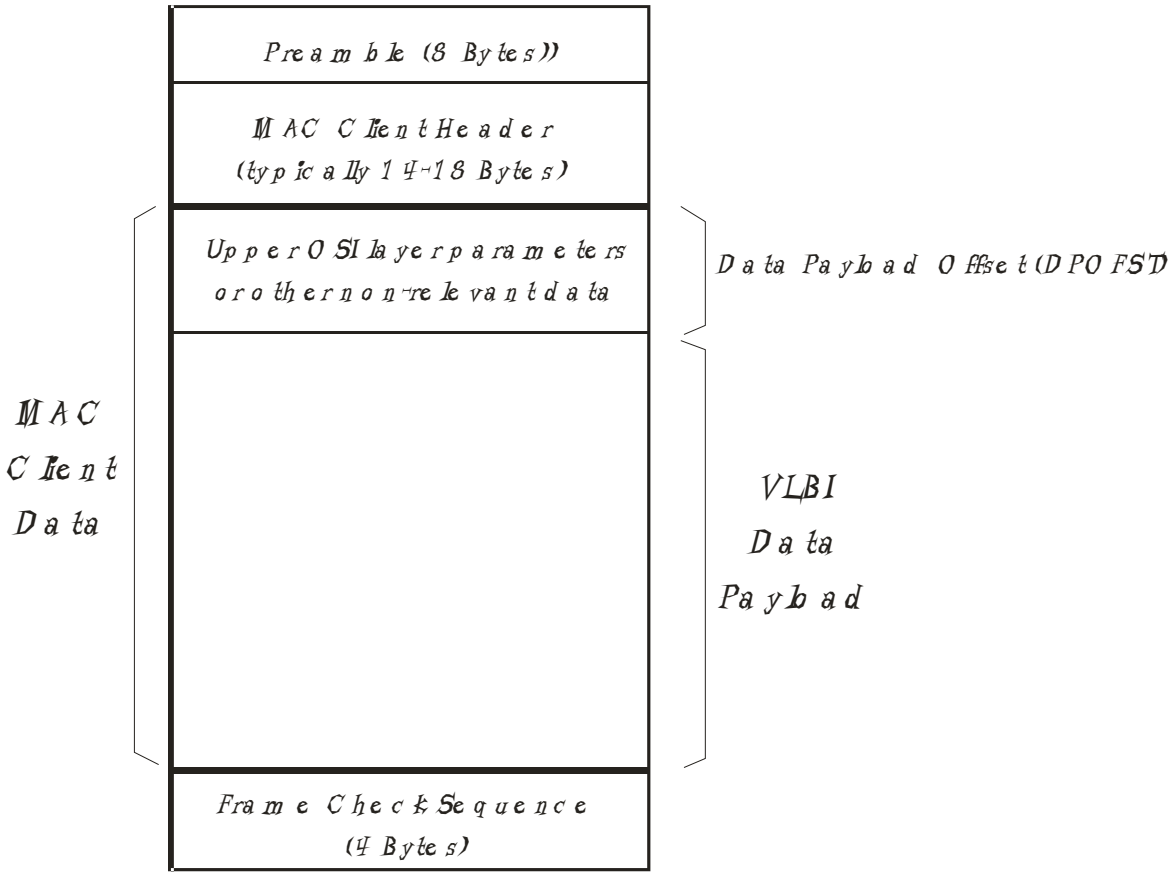

Rule: All packets must have same length within a single scan.

Mark 5C VLBI Data Payload

- Composed of (optional) Data Frame Header followed by Data Frame
- Data Frame Header may be of arbitrary length (up to ~9000 bytes), but must be multiple of 8 bytes
- All data in Data Frame are recorded

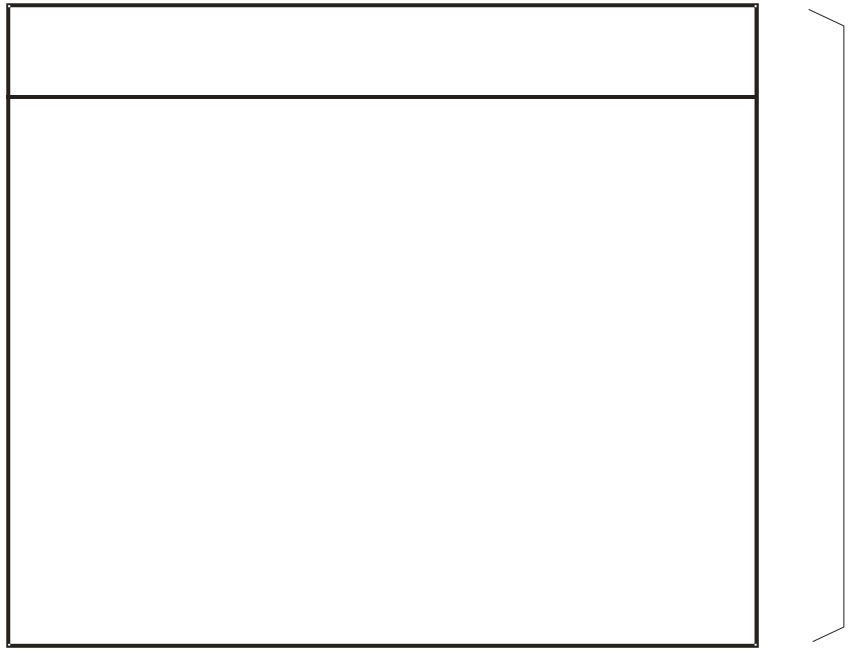

Packet Sequence Number

- VLBI Data Payload may contain a 32-bit 'Packet Sequence Number' (PSN) generated by data source
- May be either in Data Frame Header or Data Frame
- If PSN is to be monitored by Mark 5C, packet PSN must be generated in strict integer order
- Mark 5C may be instructed to monitor PSN to:
  1. Monitor for missing or bad packets and replace recorded Data Frame with user-defined fill pattern (mostly for Mk5B compatibility mode)
  2. Re-order packets arriving out of order (within limits)
  3. Mark a Data Frame as invalid to prevent it from being recorded (only for standard Mark 5C mode)
- PSN may be embedded either in Data Frame Header or Data Frame itself (will be recorded if in Data Frame)

Mark 5C Data Frame Header format

	Bit 31	23	Bit 0
Word 0	Sync word (0xdec0de5c)		
Word 1	ID	I	Data Frame # within second
Word 2	Integer seconds since 00UT 1 Jan 1990		
Word 3	Unassigned		

- Must be integral number of Data Frames per second
- First sample in Data Frame number zero must correspond to second tick
- Mark 5C data may be discontinuous in time, but must be monotonic in time; useful for pulsars, burst-mode mm-VLBI, etc
- 00UT 1 Jan 1990 is standard reference time for RTP data streams used by VSI-E
- 'Unassigned' may be used for Packet Serial Number
- 'I' indicates data in Data Frame is 'invalid'
- Format is almost same as Mark 5B disk frame header format

Mark 5C Data-Frame data format

- Each Data Frame contains samples from one freq channel
- Multi-bit sampling is supported
- Bits within each sample are adjacent in Data Frame word; LSB in lower bit number

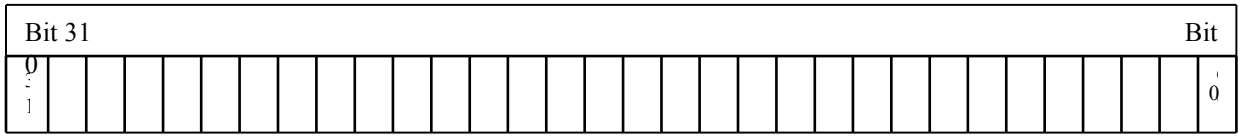

1-bit/sample data word format

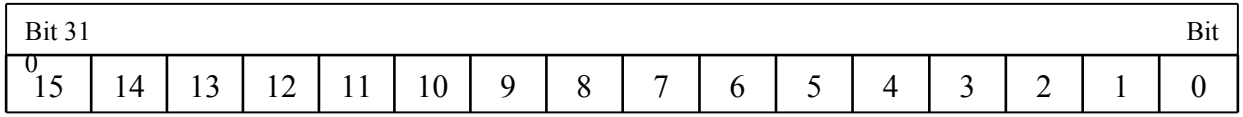

2 bits/sample data word format

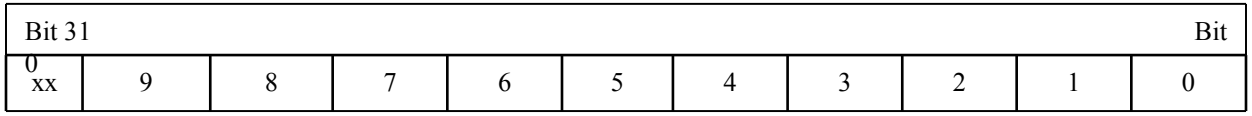

3 bits/sample data word format. ‘xx’ indicates unused bits

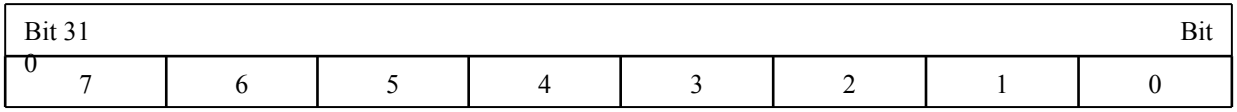

4 bits/sample data word format

Mark 5B emulation mode

- Goal is to create disk in format which can be played back on Mark 5B DOM (but will not play back on Mark 5A+)
- Mark 5B disk frame is 10016 bytes, too long for single Ethernet packet, so data source must create Mark 5B disk frame as two adjacent packets, each 5008 bytes long
- Sample coding same as Mark 5B (VLBA coding)

	Bit 31	28 27	15	Bit 0
Word 0	Sync word			
Word 1	Years from 2000	User-specified data	T	Frame# within second (starting at 0)
Word 2	VLBA BCD Time Code Word 1 (‘JJJSSSSS’)			
Word 3	VLBA BCD Time Code Word 2 (‘.SSSS’ plus 16-bit CRCC)			

Mark 5B emulation Disk Frame Header format
(shaded parameters are created by Mark 5B, but not needed for Mk5B emulation mode)

Mark 5C Sample coding

Fixed-point two's-complement of desired number of bits

2-bit/sample example

Voltage level	Sample code
> +threshold	01
0 to +threshold	00
0 to −threshold	11
< −threshold	10

Mark 5C setup parameters

- List of valid MAC source addresses (to filter out spurious packets)
- Length of MAC Client Data; option to ignore packets of wrong length
- Byte offset to beginning VLBI Data Payload
- PSN monitoring mode
  - 0 - off
  - 1 - write replacement fill pattern for missing or bad packets; primarily for Mark 5B emulation mode
  - 2 - do not write packets tagged as ‘Invalid’
- Byte offset to PSN (if PSN monitoring is enabled)

10 GigE data sources

- Two design approaches are currently being pursued:
  1. ‘dBBC’ (EVN)
     - VSI-E to 10GigE converter being developed
  2. DBE2/VDBE (Haystack/NRAO/Berkeley/South Africa collaboration)
     - Based on next-generation iBOB board (‘iBOB2’); ~$1K each
     - Haystack developing PFB app; two 1-GHz IFs per iBOB2 - 8 Gbps/ iBOB2; two iBOB2s packaged in single chassis – 16 Gbps aggregate
     - NRAO developing ‘VDBE’ app; similar to dBBC;
       plan to outfit VLBA with VDBE and Mark 5C for 4Gbps operation by sometime summer 2008.

iBob1/DBE1

Sampler boards

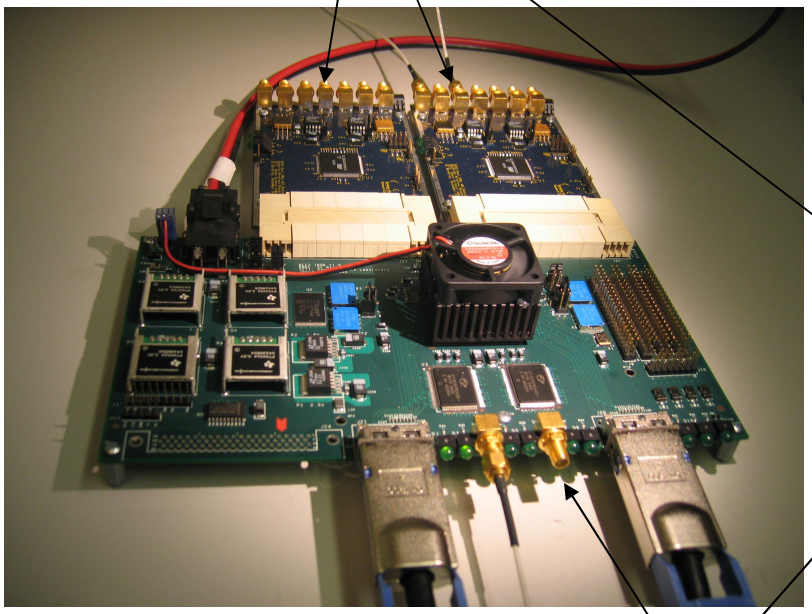

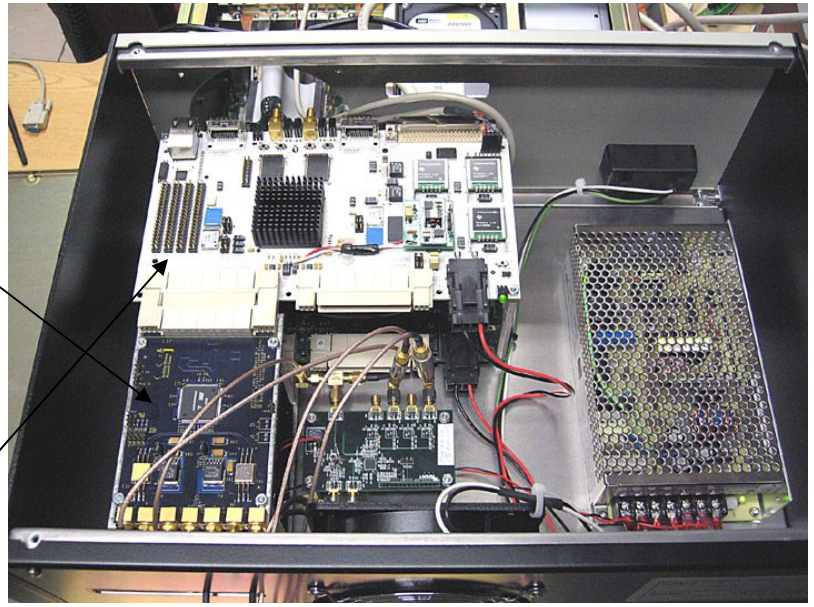

iBOB1

(two iBOB1 bds per chassis)

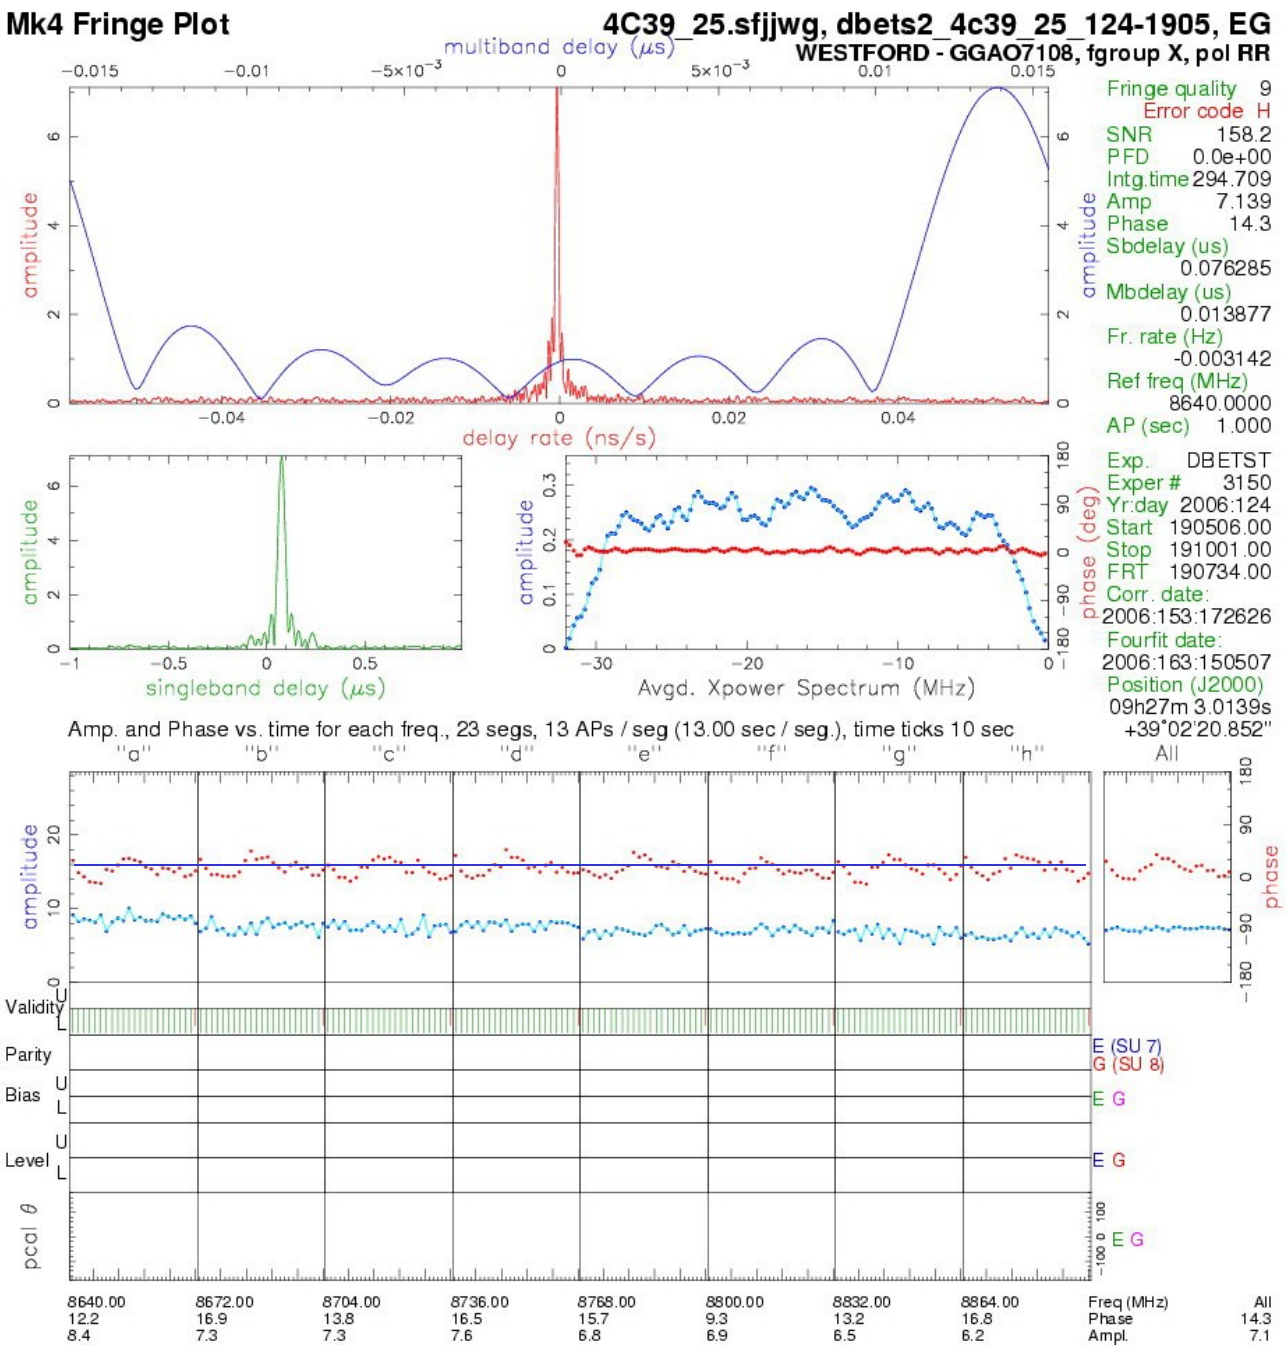

DBE1 vs DBE1 with no phase adjustments whatever

RMS phase across frequency channels:

RMS: 2.5 deg

Theor: 1.0 deg

Comparable to best we've seen even with 'manual' adjustments to embedded phase-cal

Mark 5 Upgrade Costs

Existing \ Target	**Mk5A**	**Mk5B** (requires VSI-H data source)	**Mk5B+**	**Mk5C** (not yet available; rough estimates)
0	Unavailable	$20.8K	~$22.3K	$20-25K
Mk5A	-	~$3.5K (Mk5B I/O)	~$13K (Amazon plus Mk5B I/O)	Est. $12-14K (Amazon plus 10GigE DB)
Mk5B	-	-	~$9.6K (Amazon)	Est. $12-14K (Amazon plus 10GigE DB)
Mk5B+	-	-	-	Est. $2-4K (10GigE DB)

Note: Costs do not include cost of creating data source

Disk-Media Status

- Hard disk price vs capacity/performance continues to drop
  - Now below ~$0.30/GB and continues to drop
    (Mark 4/VLBA tape is ~$2.00/GB)
- 500 GB disks are commonly used –
  Two 8-packs of 500GB disks comparable to ~16 VLBA/Mark 4 tapes
- 750 GB disks –
  Two 8-packs of 750GB disks comparable to ~24 VLBA/Mark 4 tapes;
  ~26 hours @ 1 Gbps unattended!
- 1 TB disks – two 8-packs will sustain 4 Gbps for ~9 hours

Disk Price Timeline

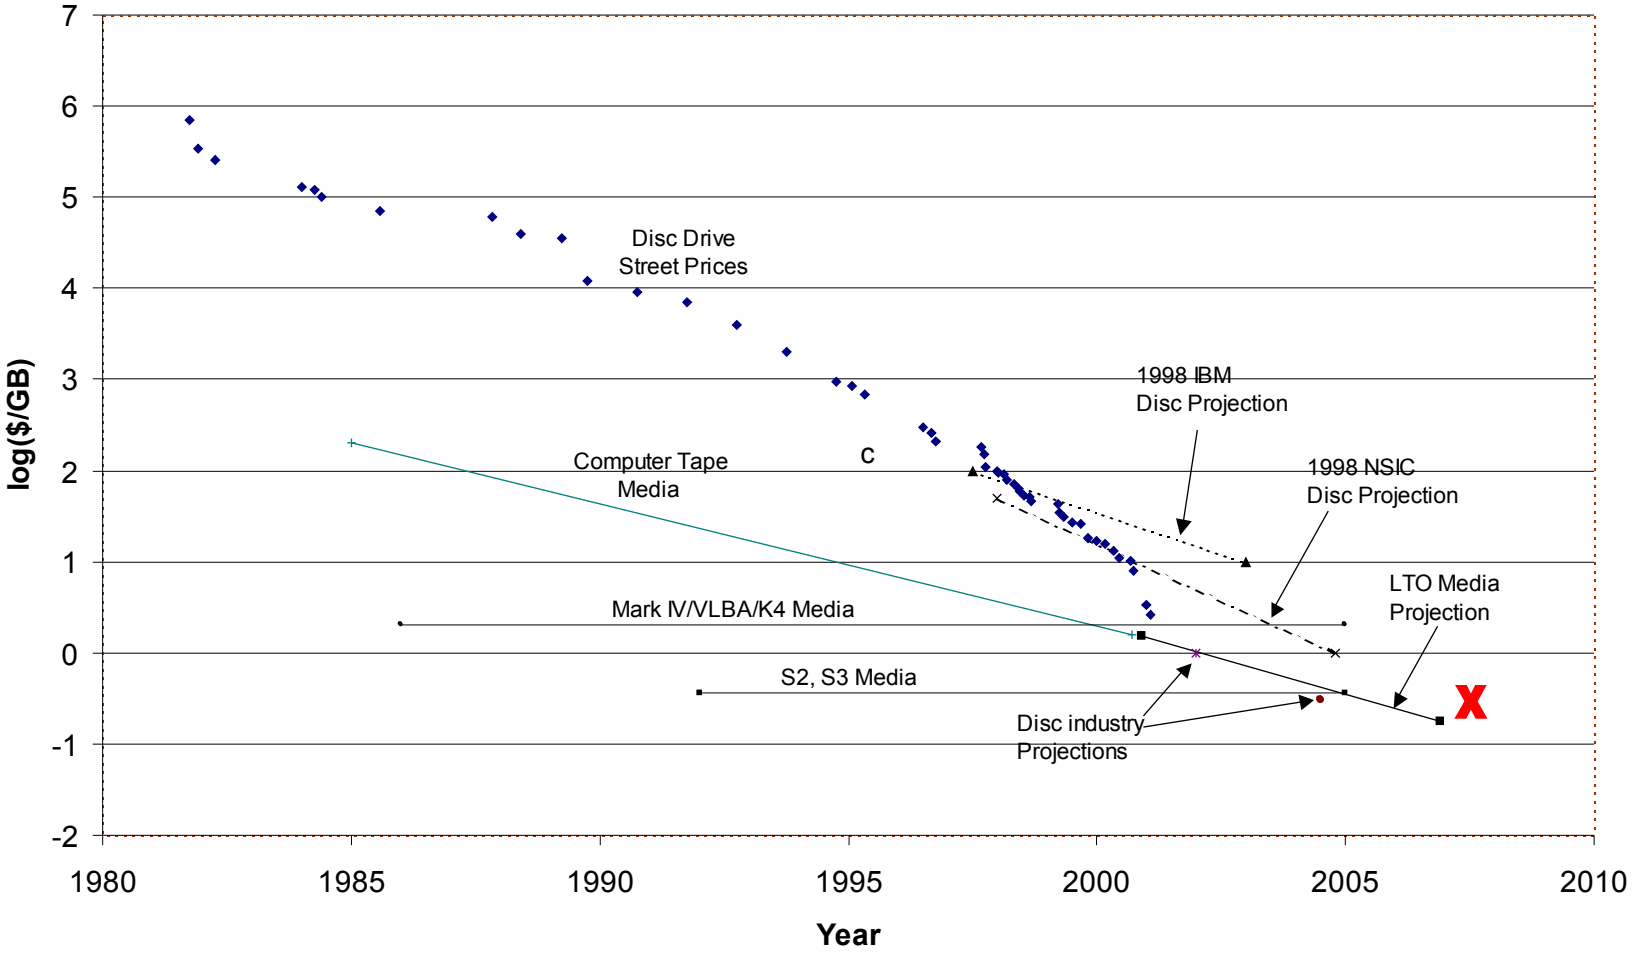

SATA disk module

- Now available from Conduant
- Interchangeable with PATA disk module in Mark 5A/A+/B/B+/C
- New mechanical design allows very easy access to insert/remove individual disks; increased module stiffness for better mechanical stability

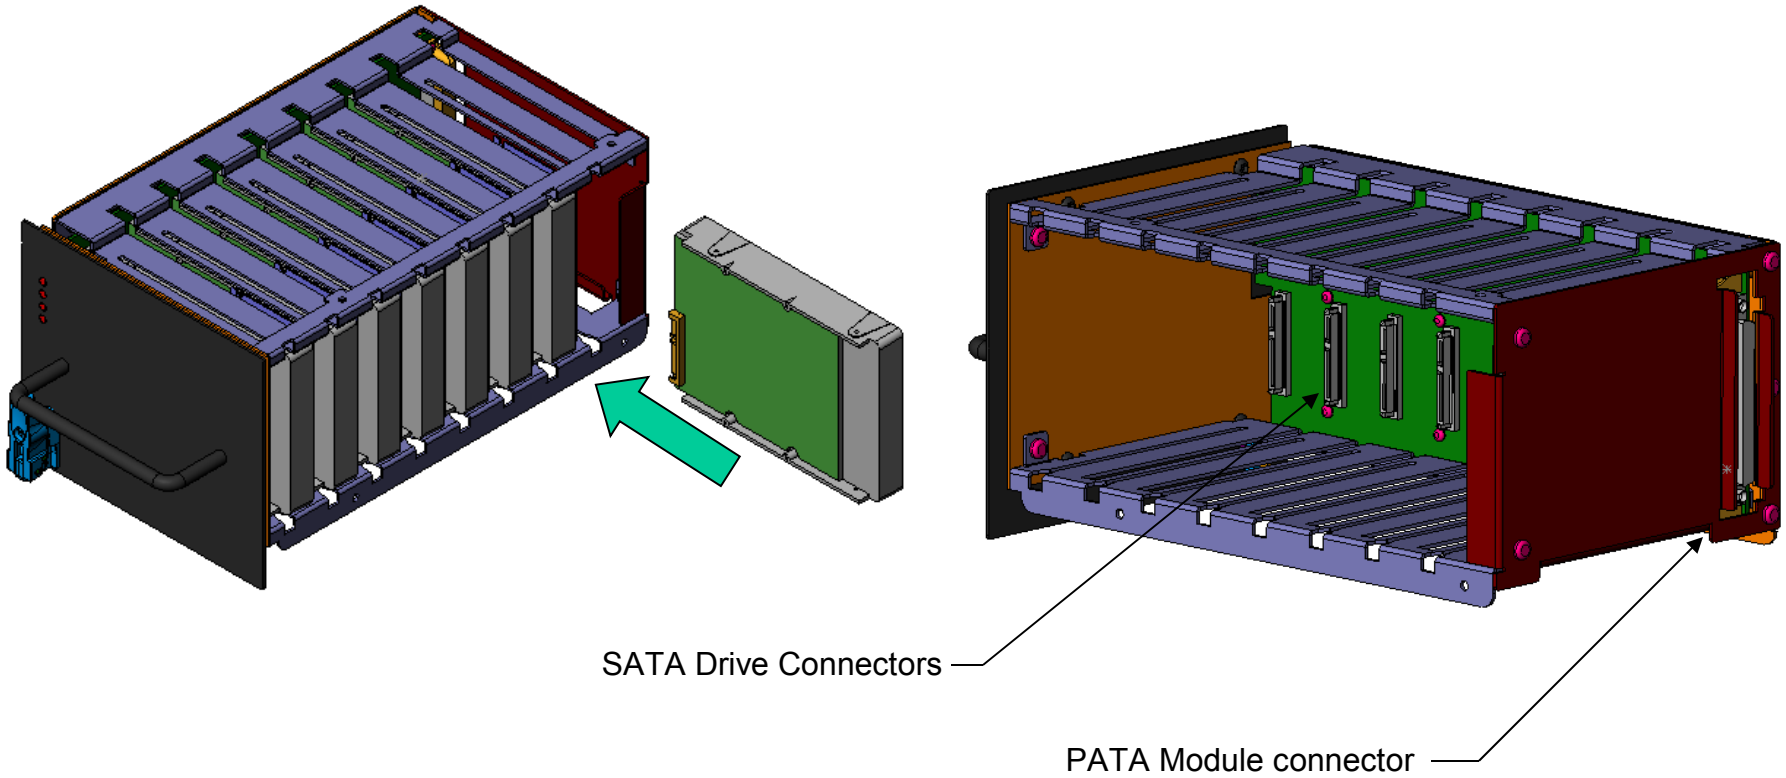

Summary

- Mark 5C is first of Mark 5 series to adopt industry-standard 10GigE data interface
- Will be able to take advantage of standard commercial 10GigE switches to re-organize and re-route data as desired
- 10GigE is natural interface to software correlator systems
- Estimated upgrade cost from Mark 5A or Mark 5B: ~$12-14K
- Estimated upgrade cost from Mark 5B+: ~$2-4K
- Expect first prototype Mark 5C systems in early/mid 2008

Questions?