Mark 5C VLBI Data System

Alan Whitney
MIT Haystack Observatory

Jon Romney
National Radio Astronomy Observatory

Kenneth Owens
Conduant Corporation

27 April 2009
TOW meeting
MIT Haystack Observatory

Mark 5 Data Acquisition System

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

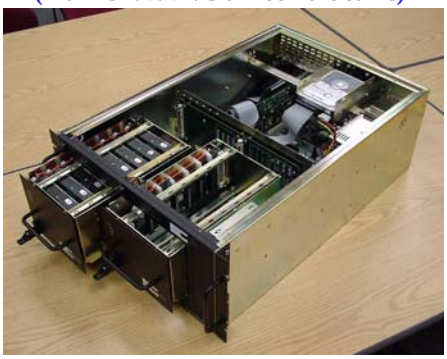

	Year introduced	Record rate (Mbps)	Interface	Cost (USk$)	#deployed
Mark 5A	2002	1024	Mk4/VLBA	21	~130
Mark 5B	2005	1024	VSI-H	22	~40
Mark 5B+	2006	2048	VSI-H	23	~30
Mark 5C	2009	4096	10GigE	21	-

Mark 5C Characteristics

- Mark 5C specification developed jointly by Haystack & NRAO; comments solicited from global community in early 2007
  - 4096 Mbps max data rate to two standard Mark 5 disk modules
  - 10GigE data interface
- A "dumb" Ethernet packet recorder (accepts OSI Layer 2 or higher)
  - Not VLBI specific
- Data source is responsible for data formatting and Ethernet packet creation
  - Expect most applications will create VDIF-compliant data streams
- Record through real-time hardware 10GigE interface
- Playback of Mark 5C will be through host computer accessed as standard Linux files using FUSE
  - Natural for software correlators

Generalized 10GigE Data Distribution Concept

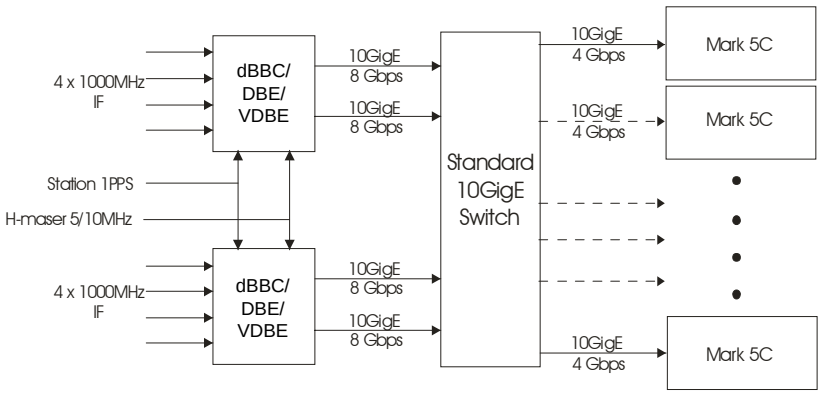

Evolving Mark 5C Functionality

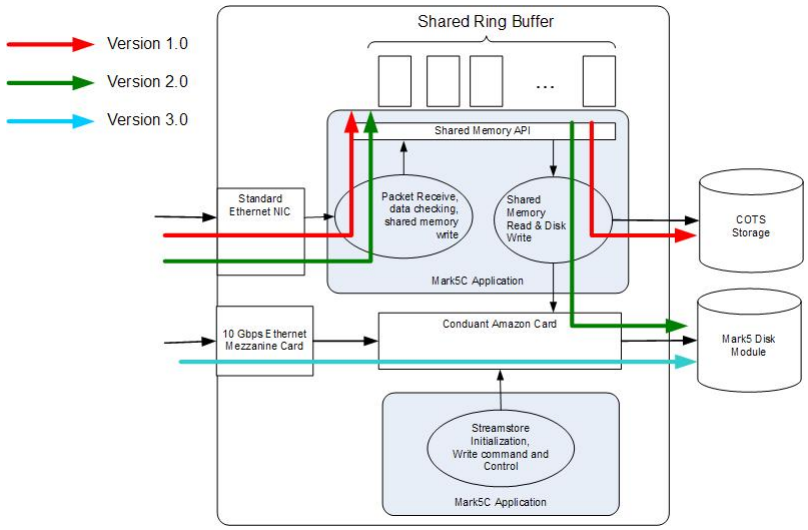

5

Mark 5 Upgrade Costs

Target / Existing	**Mk5A**	**Mk5B** (requires VSI-H data source)	**Mk5B+**	**Mk5C** (not yet available; estimates)
0	Unavailable	$21K	~$23K	~$21K
Mk5A	-	~$4K (Mk5B I/O)	~$13K (Amazon plus Mk5B I/O)	~$11K (Amazon plus 10GigE DB)
Mk5B	-	-	~$8K (Amazon)	$11K (Amazon plus 10GigE DB)
Mk5B+	-	-	-	~ $3K (10GigE DB)

Mark 5C Status

- First Mark 5C prototype hardware has been delivered
  - Software development and hardware testing underway
  - Expect first 4Gbps demonstration soon
- 2nd-generation digital backends currently being designed as 10GigE VDIF data sources
  - dBBC with Fila10 board (Europe)
  - DBE2 – based on ROACH board; 2nd generation iBOB (US)
  - Also possible developments in Australia, China, Japan
  - Some prototypes expected mid-late 2009

Disk-Media Status

- Hard disk price vs capacity/performance continues to drop
  - Now below ~$0.10/GB and continues to drop (Mark 4/VLBA tape was ~$2.00/GB)
- 750 GB disks – Two 8-packs of 750GB disks comparable to ~24 VLBA/Mark 4 tapes; ~26 hours @ 1 Gbps unattended!
- 2 TB disks – two 8-pack modules will sustain 4 Gbps for ~18 hours
- Currently ~1500 Mark 5 data modules with >2PB capacity

PATA & SATA: $/GB vs. time

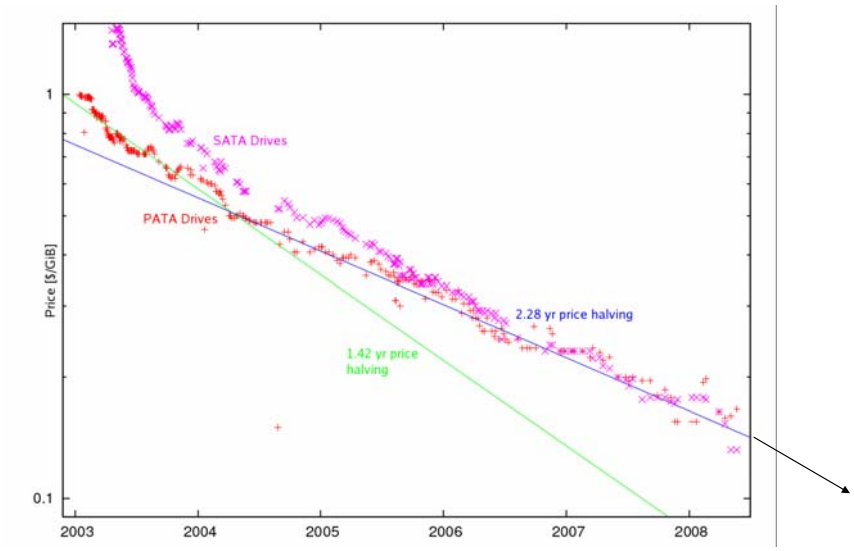

courtesy Walter Brisken; data based on pricewatch.com monitoring

Ethernet Packet Structure

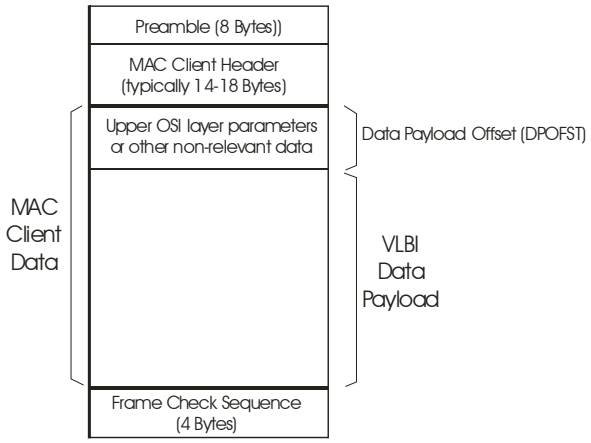

Rule: All packets must have same length within a single scan.

VLBI Data Interchange Format (VDIF)

- Problem: Variety of VLBI data formats used internationally complicates international data transfer
- Internationally constituted VDIF Task Force appointed in Shanghai in June 2008 to study problem and create a recommended uniform transport-independent time-sampled data-format standard (designed for VLBI, but non-VLBI specific)
- Network data-transport standard will be addressed separately

VDIF Assumptions

- Data are assumed to be one or more time series of uniformly time sampled data
- Each time series may have own sample rate, bits/sample and station of origin
- Data from all time series are organized into a single serial 'Data Stream'

VDIF Requirements

- Data are self-identifying wrt UTC time tag and data identification
- Data may be single bit or multi-bit samples
- Data can be decoded with no external information
- Number of sampled time series (or ‘channels’) may be arbitrary (i.e. not confined to 2^n)
- Suitable for on-wire transfer as well as disk file storage
  - Modest out-or-order data support
- Support for aggregate data rates up to at least 100Gbps

VDIF Data Structure

- A ‘Data Stream’ is organized into self-identifying ‘Data Threads’
- Each Data Thread contains of a serial set of ‘Data Frames’
- Each Data Frame contains a 16/32B header plus Data Array
- Data Frame length may be chosen by user (with some restrictions)
- Data Frame may contain single-channel or multi-channel data
- Each Data Thread may have individual # of channels, sample rate and bits/sample

Possible Future Directions

- Higher data rates (8-16 Gbps)
  - New module connector with many high-speed serial data streams
    - Suitable connector has been identified
    - Would not be backwards compatible with current modules
  - Move to 16 or more 2.5” SATA disk drives in current module geometry
    - Data rates to >=8Gbps with single module; >=16Gbps with two modules
    - Possibility to use solid-state disk drives for higher reliability (currently too expensive, but prices dropping rapidly)
- Possible data interface upgrades
  - Multiple 10GigE interfaces
  - 100GigE

Summary

- Mark 5C is first of Mark 5 series to adopt industry-standard 10GigE data interface
- Mark 5C is a general-purpose Ethernet packet recorder
- Takes advantage of standard commercial 10GigE switches and data-transport hardware
- 10GigE is natural interface to software correlator systems
- Relatively low cost: ~$21k
- Expect first prototype Mark 5C systems mid-2009
- VDIF data format standard has been developed in parallel as international standard
- VDIF-compatible backends are being developed

Thank you