

Mark 5C VLBI Data System

Alan Whitney

MIT Haystack Observatory

Chet Ruszczyk

MIT Haystack Observatory

Jon Romney

National Radio Astronomy Observatory

Kenneth Owens

Conduant Corporation

10 Feb 2010


IVS General Meeting

Hobart, Tasmania



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	2010	4096	10GigE	21	-

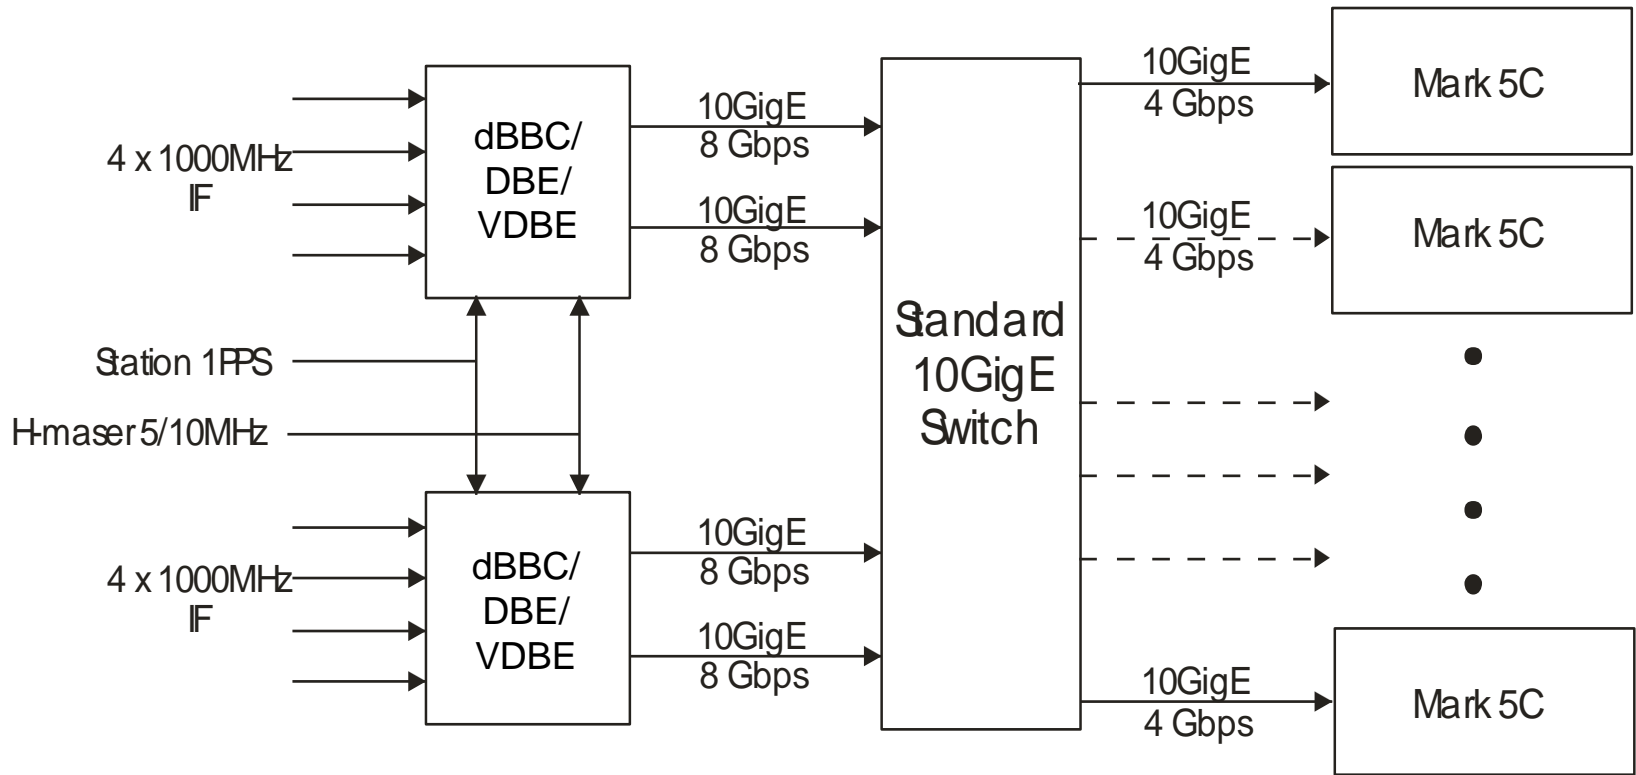


Mark 5C Characteristics

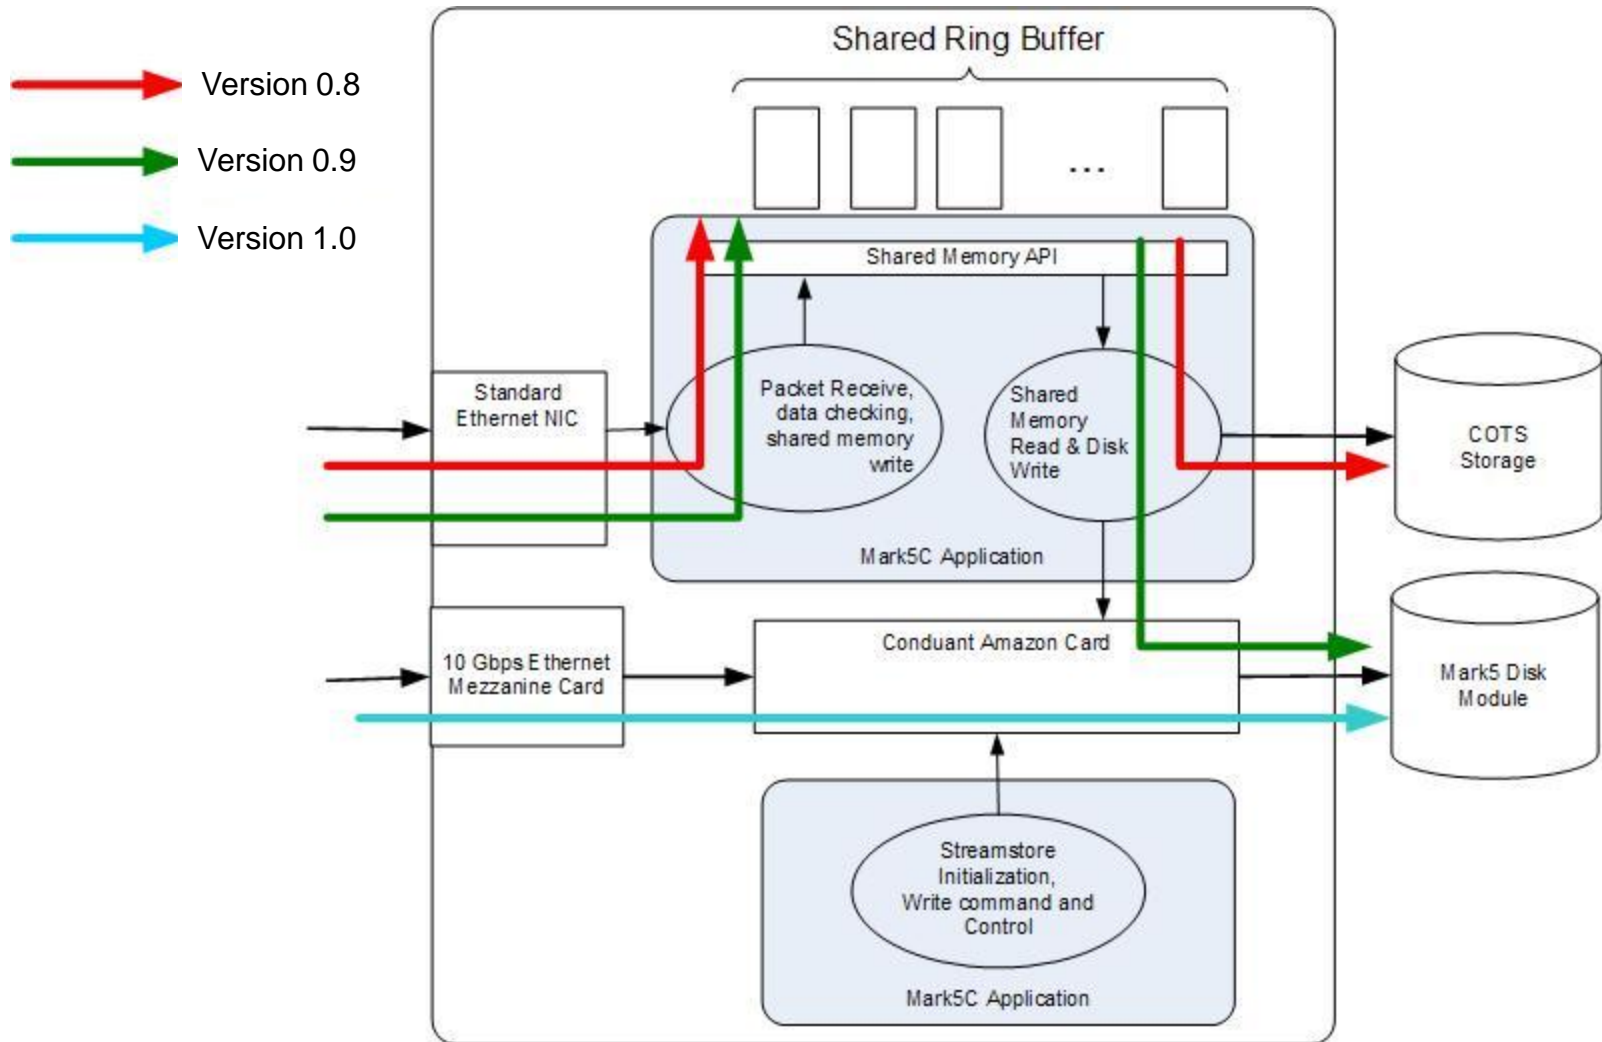
- Mark 5C specification developed jointly by Haystack, NRAO & Conduant Corp
 - 4096 Mbps max data rate to two standard Mark 5 disk modules in ‘non-bank’ mode
 - 10GigE receive interface on I/O daughterboard for Amazon SS card; CX4 physical connector
- A “dumb” Ethernet packet recorder (accepts OSI Layer 2 or higher)
 - Not VLBI specific
 - Plan to use VDIF format
- Data source is responsible for data formatting and Ethernet packet creation
 - Expect most applications will create VDIF-compliant data streams
- Playback of Mark 5C will be through host computer accessed as standard Linux files using FUSE
 - Natural for software correlators



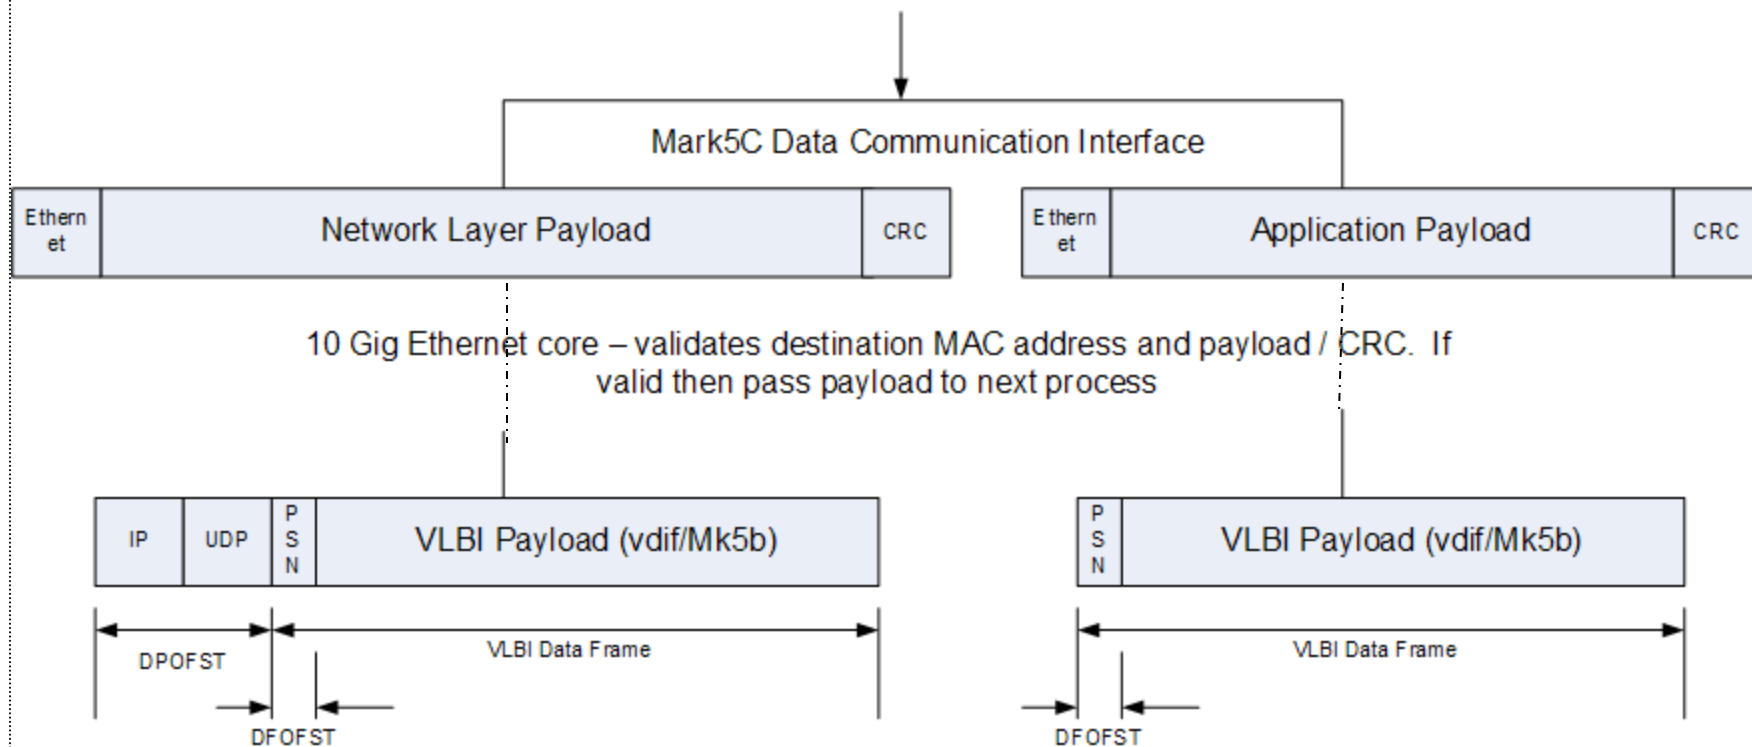
Generalized 10GigE Data Distribution Concept



Evolving Mark 5C Functionality



Physical Media



10 Gig Ethernet core – validates destination MAC address and payload / CRC. If valid then pass payload to next process

- Using PSN Mode 1: guarantees order and corrects for missing packets
- DPOFST – (Data Payload Offset): Number of bytes from received packet to find VLBI Data Frame.
 - For UDP/IP Packets number of bytes of standard header ~46 bytes
 - For Ethernet encoded packets = 0 bytes
- DFOFST – (VLBI Data Frame Offset): Number of bytes within payload to start of recording to disk.
- PSNOFST (Packet Serial Number Offset) : Since the PSN does not have to be the first word but embedded in the packet – the location of PSN.



VLBI Data Interchange Format (VDIF)

- Problem: Variety of VLBI data formats used internationally complicates international data transfer
- VDIF Task Force appointed in Shanghai in June 2008 to study problem and create transport-independent time-sampled data-format standard (designed for VLBI, but non-VLBI specific)
- VDIF specification ratified in June 2009 at Madrid e-VLBI wkshp; see Poster at this meeting
- e-VLBI transport standard (VTP) being addressed separately; task force led by Chris Phillips (ATNF) has been appointed



Mark 5C Status

- Mark 5C prototype hardware has been delivered
 - Software development and hardware testing underway
 - Tested in 2Gbps, bank mode (i.e. 2Gbps to one disk pack)
 - Waiting for SKD 9.0 from Conduant (in Beta test)
 - Will support 4Gbps in ‘non-bank’ mode (i.e. writing to two disk packs simultaneously))
 - Expect first 4Gbps demonstration soon
 - Read capability: FuseMk5
- 2nd-generation digital backends currently being designed as 10GigE VDIF data sources
 - dBBC with Fila10 board (Europe)
 - RDBE: based on ROACH board; 2nd generation iBOB (US)
 - Currently under test in 2Gbps Mark 5B+ emulation mode
 - Also possible developments in Australia, China, Japan
 - Prototypes expected 2010



Mark 5 Upgrade Costs

Target	Mk5A	Mk5B (requires VSI-H data source)	Mk5B+	Mk5C (not yet available; estimates)
Existing				
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)



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 ≥ 8 Gbps with single module;
 ≥ 16 Gbps 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-2010
- VDIF data format standard has been developed in parallel as international standard
- VDIF-compatible backends are being developed




Thank you



Mark5C Data System

- Joint development effort of MIT Haystack, NRAO and Conduant Corporation
- Designed to meet the Mark5C specification
 - MIT Haystack Memo #57
 - VLBA Sensitivity Upgrade memo #12
- 4 Gbps recording capability
- Payload format independent
- Two main components
 - Hardware
 - Application



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-of-order data support
- Support for aggregate data rates up to at least 100Gbps

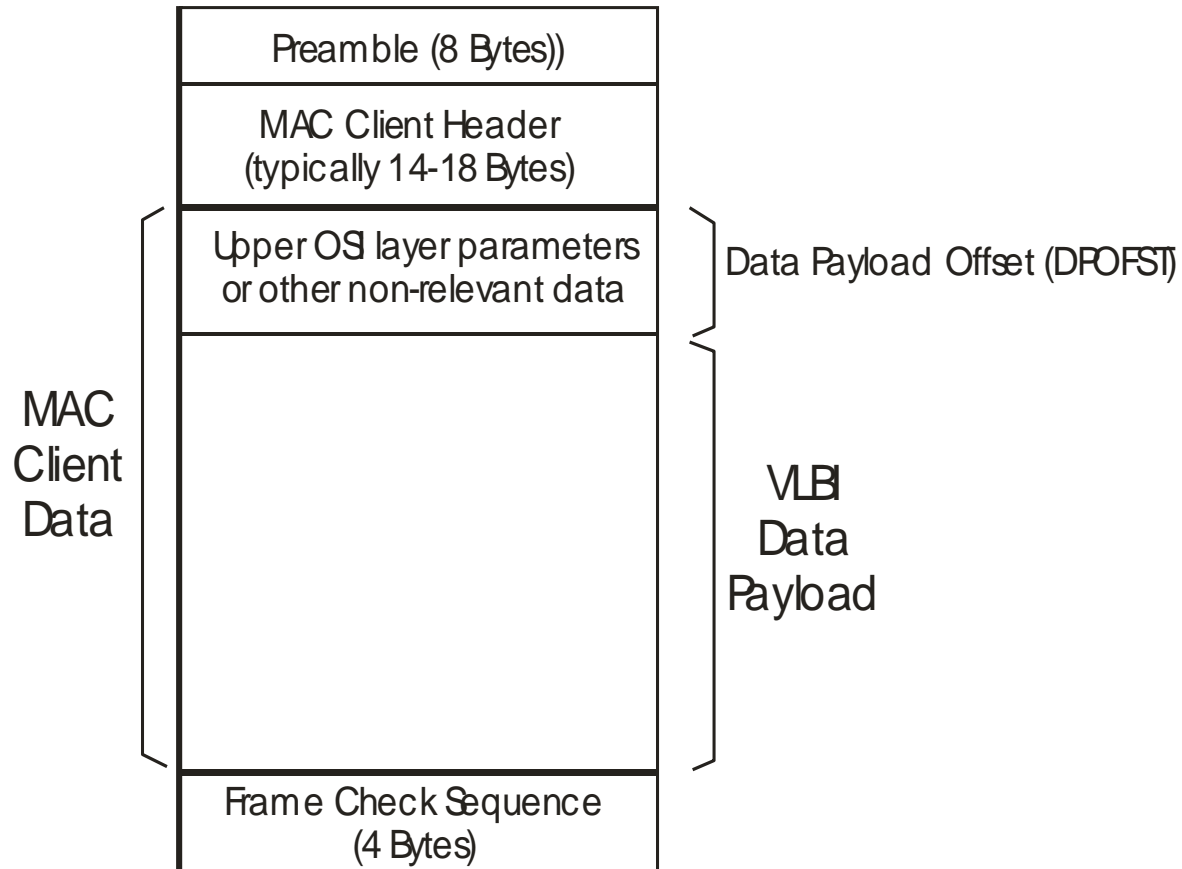


VDIF Data Structure

- 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



Ethernet Packet Structure



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



SDK 9.0

- Under beta testing at Conduant Corporation
- Key features
 - Standard function calls to access 10G hardware
 - Presently these are through direct addressing
 - Using diagnostic function calls
 - Supports 32 and 64 bit Linux kernels
 - 64 bit support not expected in initial release
 - Supports > 1TB disk drives

