


The Mark 5C VLBI Data System

Alan Whitney
MIT Haystack Observatory

25 September 2008
EVN Symposium
Bologna, Italy

Mark 5 Data Acquisition System

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



Mark 5A introduced in 2003, Mark 5B in 2005, Mark 5B+ in 2006
Mark 5C expected early 2009

Original Goals of Mark 5 VLBI Data System

(c. 2002)

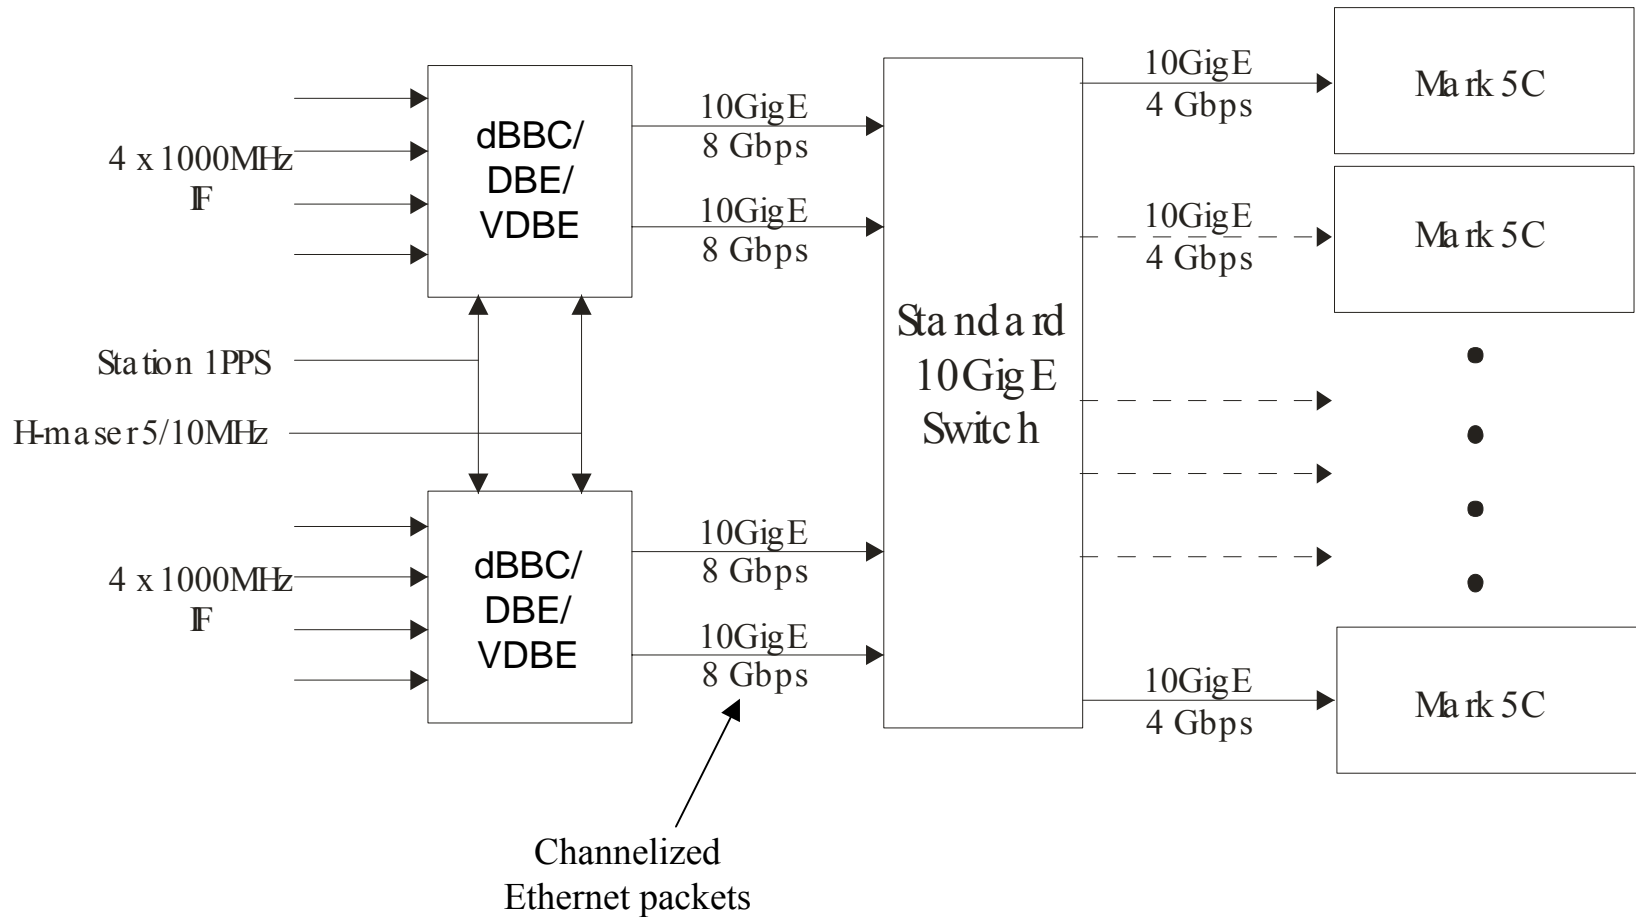
- ≥ 1 Gbps data rate
- Low cost (at least compared to tape)
- Based primarily on unmodified COTS components
- Easy-to-use
- Convenient media management
- Reliable operation, low maintenance cost
- Easy transportability
- Conformance to VSI specification
- Compatibility with existing VLBI systems during transition
- e-VLBI support
- 24-hour unattended operation at 1 Gbps

Mark 5A/5B achieved all of these objectives by ~2006,
More than 150 Mark 5 systems now in place worldwide,
displacing almost all existing magnetic tape systems in ~4 yrs,
at a cost $< 15\%$ of the tape systems.

Mark 5A/B/B+/C Thumbnail Sketch

	Mk5A	Mk5B	Mk5B+	Mk5C (under development)
Data Interface	Emulates Mk4/VLBA tape transport	VSI-H (64MHz max clock rate)	VSI-H (64MHz max clock rate)	10 GigE (OSI Layer 2)
Max data rate	1024 Mbps	1024 Mbps	2048 Mbps	4096 Mbps
Record modes	8, 16, 32, 64 “tracks”	1,2,4,8,16,32 bitstreams	Same as Mk5B	—
Disks	Mk5 “8-pack”	Same	Same	Same
Chassis	Mk5	Same	Same	Same
I/O card	Mk5A	Mk5B	Mk5B	None
SS card	XF2	XF2	Amazon	Amazon
I/O-SS intf	Modified FPDP	FPDP	FPDP2 (clocks on both edges)	10GigE Direct to Amazon
Max playback rate	1024 Mbps (to VSI-H)	1024 Mbps (to VSI-H)	1024 Mbps (to VSI-H)	~3Gbps (to host)
Playback Compatibility	Mk5A (to VSI-H)	Mk5A+ (to VSI-H)	Mk5B Mk5A+ (except 2Gbps)	—

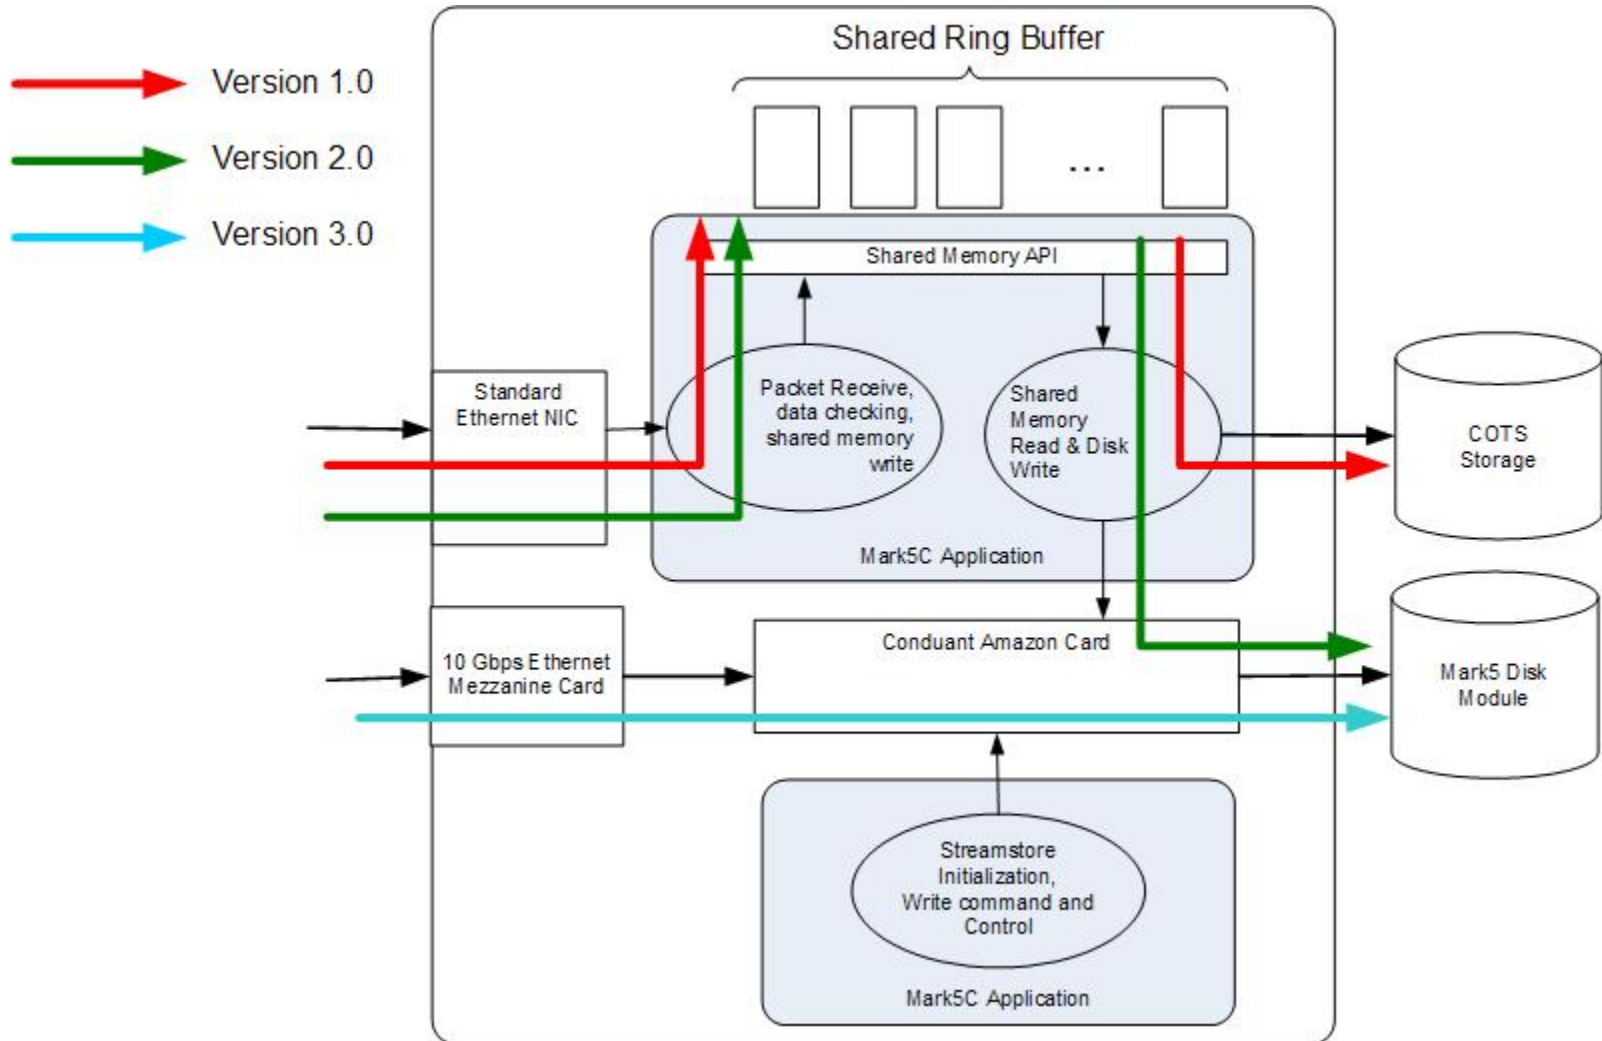
Generalized 10GigE Data Distribution Concept



Mark 5C basics

- Mark 5C is basically a ‘dumb’ asynchronous Ethernet packet recorder that simply records the payload of each Ethernet packet that it receives
- Data source is responsible for creating Ethernet packet, including time-tagging and data-formatting
- ‘Mark 5C’ packet characteristics
(but Mark 5C itself doesn’t really care):
 - Will support new VDIF data-format standard
 - One freq channel per Ethernet packet
 - Supports an arbitrary # of channels
(i.e. not constrained to 2^n channels, unlike Mk5A/B/B+)
 - Supports 1-32 bits/sample
- At 4Gbps, records to 16 disks (2 Mark 5 disk modules) simultaneously
- Backwards-compatible Mark 5B mode will write disks in Mark 5B data format (best way to do this is under discussion)

Evolving Mark 5C Functionality




‘VDIF’ (VLBI Data Interchange Format)

- Panel discussion at Shanghai e-VLBI meeting led to formation of VDIF Task Force (Mark Kettenis, JIVE; Chris Phillips (ATNF), Mamoru Sekido (NICT), Alan Whitney (MIT, chair)
- Intent is to allow easy transfer of data among heterogeneous systems (Mark 5C, PC-EVN, K5, LBADR, etc) using real-time e-VLBI, e-transfer and/or disk files
- Addresses data format only; e-VLBI transfer protocol will be addressed in a separate specification, yet to be written; combination will replace original VSI-E proposal, which was judged too complex
- Draft specification is available; final ratification expected soon!

VDIF Characteristics

- Data packaged into Data Frames
 - Compatible with e-VLBI, e-transfer, disk file storage
 - Standardized low-overhead Data Frame header
 - Standardized time-sampled data format
 - Length is user selectable according to application
 - Supports arbitrary number of channels through use of ‘Data Threads’ (breaks 2^n legacy constraint)
 - Strongly encouraged for future developments
 - Highly compatible with emerging software correlators
 - Supports 1 to 32 bits/sample
- Supports data rates in excess of 100Gbps
- Simple time-keeping algorithm supports seamless observations over leap seconds and year boundaries; requires no user knowledge of past or future leap seconds

\$/GB vs time



\$/GB declining at about same rate as increase of recording rates

Mark 5 Recording Durations

- Sustained 1 Gbps recording → ~11TB/24 hrs
 - Supported with two 8-packs of 750GB disks
 - comparable to ~24 VLBA/Mark 4 tapes!
 - Sustained 4 Gbps recording → ~44TB/24 hrs
 - Two 8-packs of 1.5TB disks → can sustain ~13 hrs
 - comparable to ~48 VLBA/Mark 4 tapes!
- Conclusion – disk-capacity increase not sustaining same rate as \$/GB decrease

Mark 5 Upgrade Costs

Target \ Existing	Mk5A	Mk5B (requires VSI-H data source)	Mk5B+	Mk5C (not yet available; rough estimates)
0	Unavailable	\$20.8K	\$22.3K	\$20.4K
Mk5A	-	~\$3.5K (Mk5B I/O)	~\$13K (Amazon plus Mk5B I/O)	~ \$12K (Amazon plus 10GigE DB)
Mk5B	-	-	~\$9.6K (Amazon)	~ \$12K (Amazon plus 10GigE DB)
Mk5B+	-	-	-	~\$3K (10GigE DB)

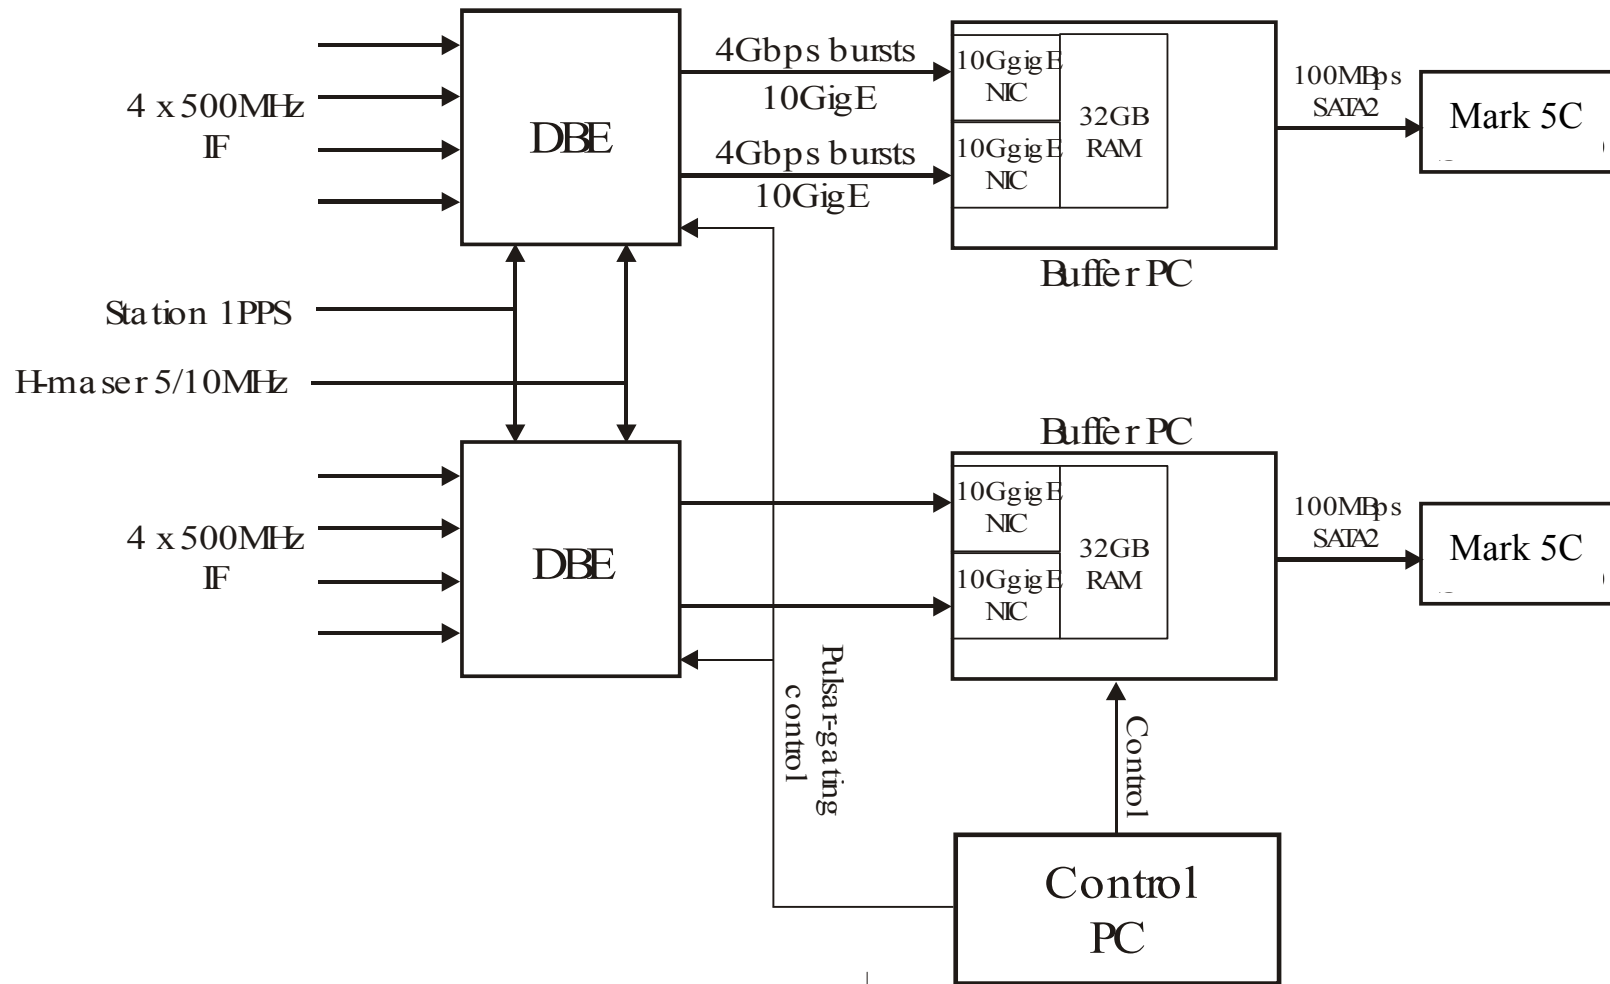
SATA Disk Support

- SATA now fully supported
- Must have proper firmware and software
- Significant difference noted between disks from different vendors
(Seagate best, WD OK, Maxtor poor)
- Many of newer, larger disks now available only with SATA

Where do we go from here?


- e-VLBI becoming popular, but...
All data from all stations must be recorded if any station unconnected!
- Disk capacities and speeds continue to increase while \$/GB drop
- 8 Gbps sustained recording rate quite feasible in next 1-2 yrs with next-generation StreamStor disk-interface card → probably write to 24 disks
- Observation rates significantly beyond 8 Gbps will be demanded by VLBI over the next few years – mm-VLBI, geodetic-VLBI

16 Gbps Burst-mode System



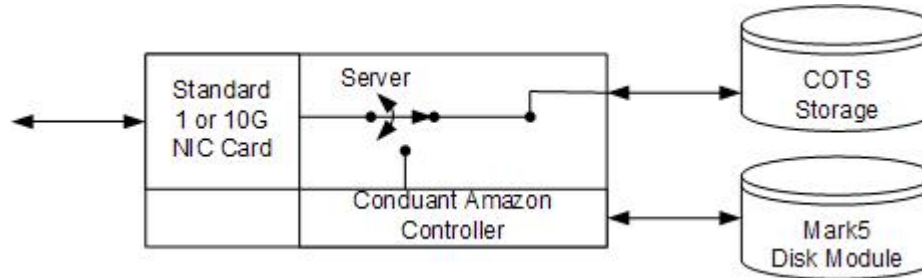
Thank you!

Mark 5C Version 1 (aka “Mark 5c-”)



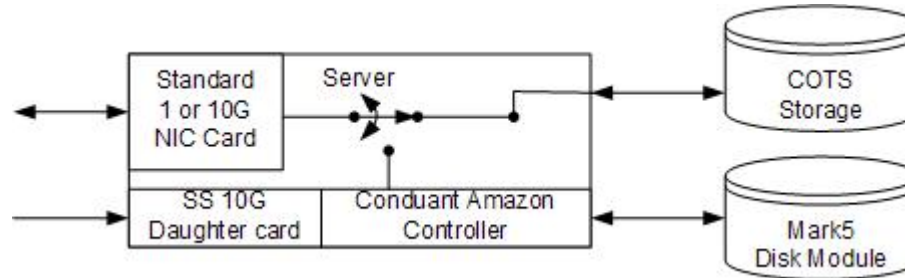
- Record path
 - Receive data from a standard NIC
 - Store data to COTS storage
- Playback path
 - Read data from COTS storage
 - Outputs data via standard NIC
- Provides the framework of application
- Limited performance capabilities < 1Gbps
- Interface for command / control
 - Command line

Mark 5C Version 2



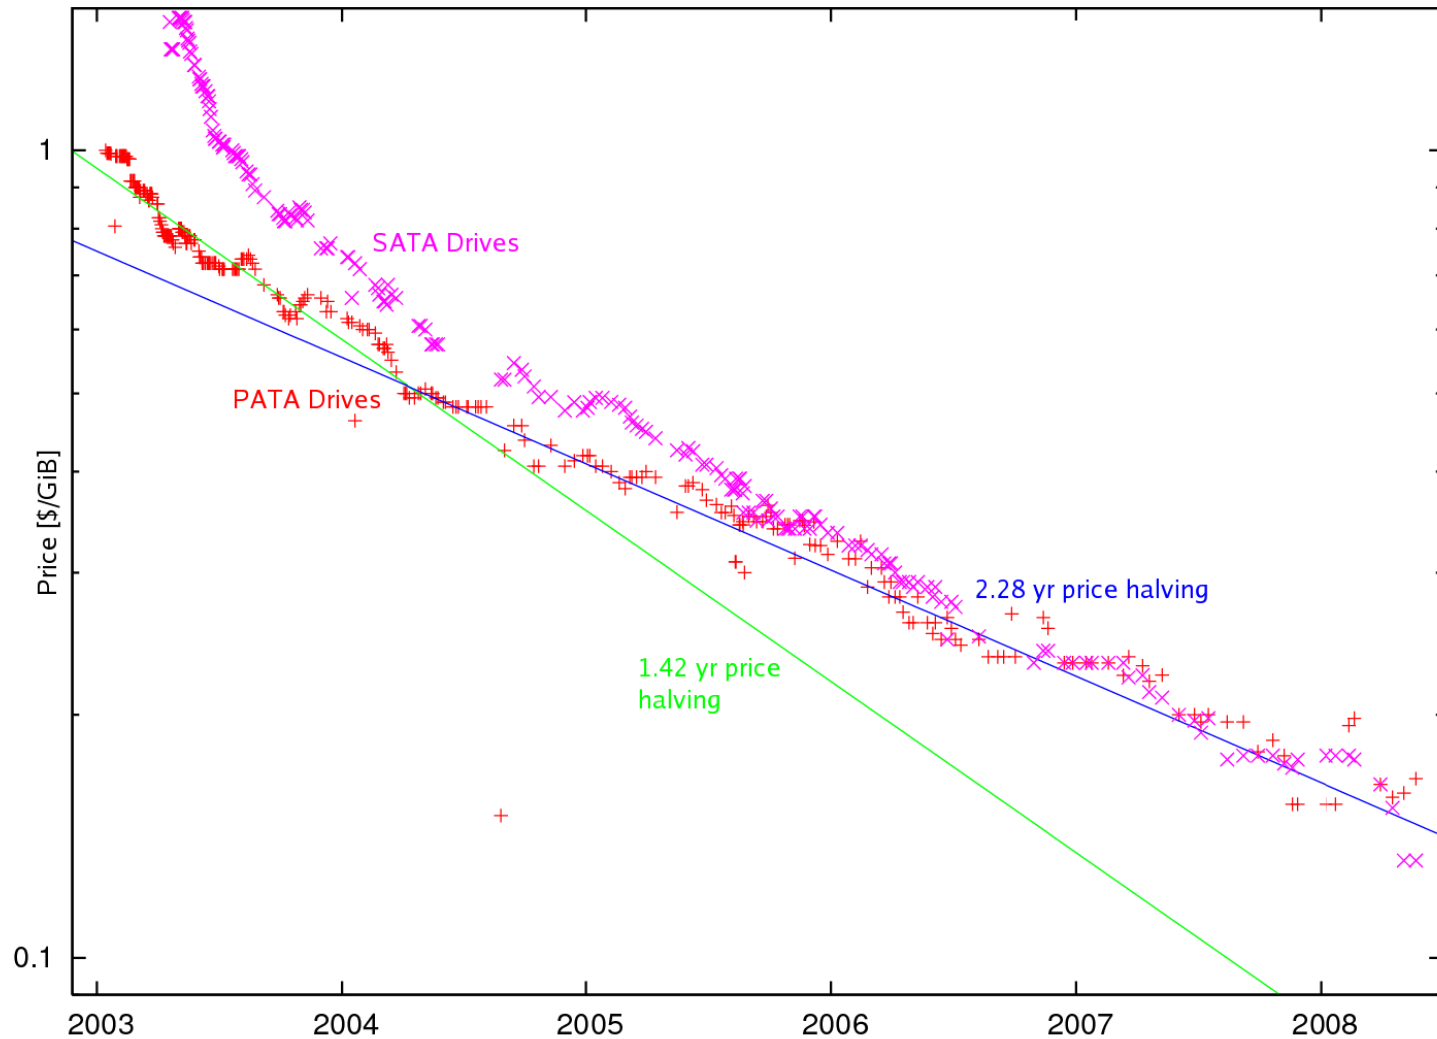
- Added capabilities
 - Record path stores data to Conduant disk module
 - Playback path reads data from Conduant disk module
- Interface for command / control
 - Full command line capabilities
 - for Mark5C-
 - Initial version of Graphical User Interface

Mark 5C Version 3



- Added capabilities
 - Support for Conduant's
 - Hardware - 10Gbps daughter card
 - Software Design Kit (SDK)
 - Data received / stored by Conduant controller card
 - Initialization of controller card
 - Packet acceptance criteria
 - Status of hardware / receive process

PATA & SATA: \$/GB vs. time



courtesy Walter Briskin; data based on pricewatch.com monitoring