Mark 6 16Gbps VLBI Data System

Alan Whitney & David Lapsley
for the Mark 6 development team
MIT Haystack Observatory

5 December 2011
DiFX workshop
MIT Haystack Observatory

Mark 6 project

- 16 Gbps COTS-based VLBI data system based on COTS-hardware, Linux OS, open-source software
- Immediate targets – VLBI2010 and mm-VLBI
- Mark 6 is a collaborative development effort between:
  - Haystack Observatory – all software and software support; hardware specification
  - NASA/GSFC High-End Network Computing group – consultation on high-performance COTS
  - Conduant Corp –Mark 6 disk module, disk-module power management

Mark 6 features & goals

- 16Gbps sustained record and playback capability
- >=32Gbps burst-mode capability
- General Ethernet packet recorder
  (can be straight-forwardly adapted to other interfaces as well)
- Based on inexpensive high-performance COTS hardware
- Easily upgradeable on Moore's Law curve
- Linux OS (Debian Squeeze 6.0.3) w/**fully open-source software**
- Playback as standard Linux files
- VLBI Disk-Adaptive Format (VDAF) file system to manage slow and failed disks (so you can write 'VDIF' on 'VDAF' ☻)
- e-VLBI support
- Smooth transition from Mark 5
- Preserve as much investment in existing Mark 5 systems and disk libraries as possible
- Extensive stress testing in real-world operational environment (systems currently deployed at Westford and GGAO antennas)

Basic Mark 6 System

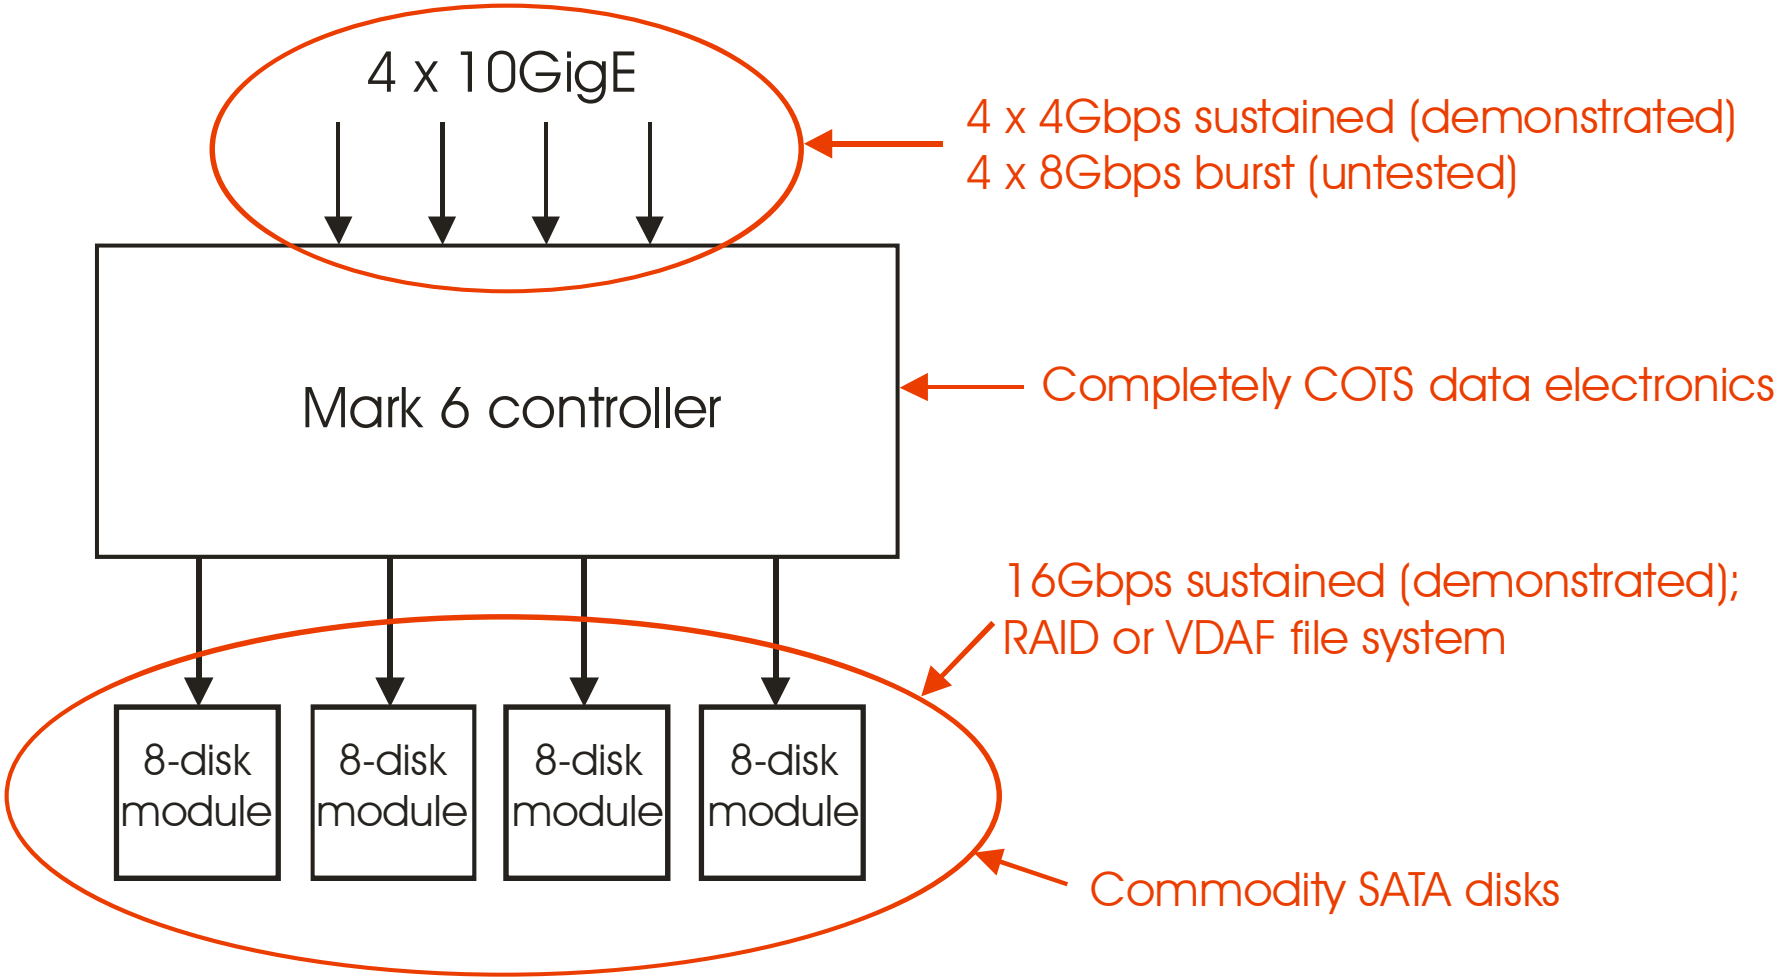

Mark 6 Physical Layout

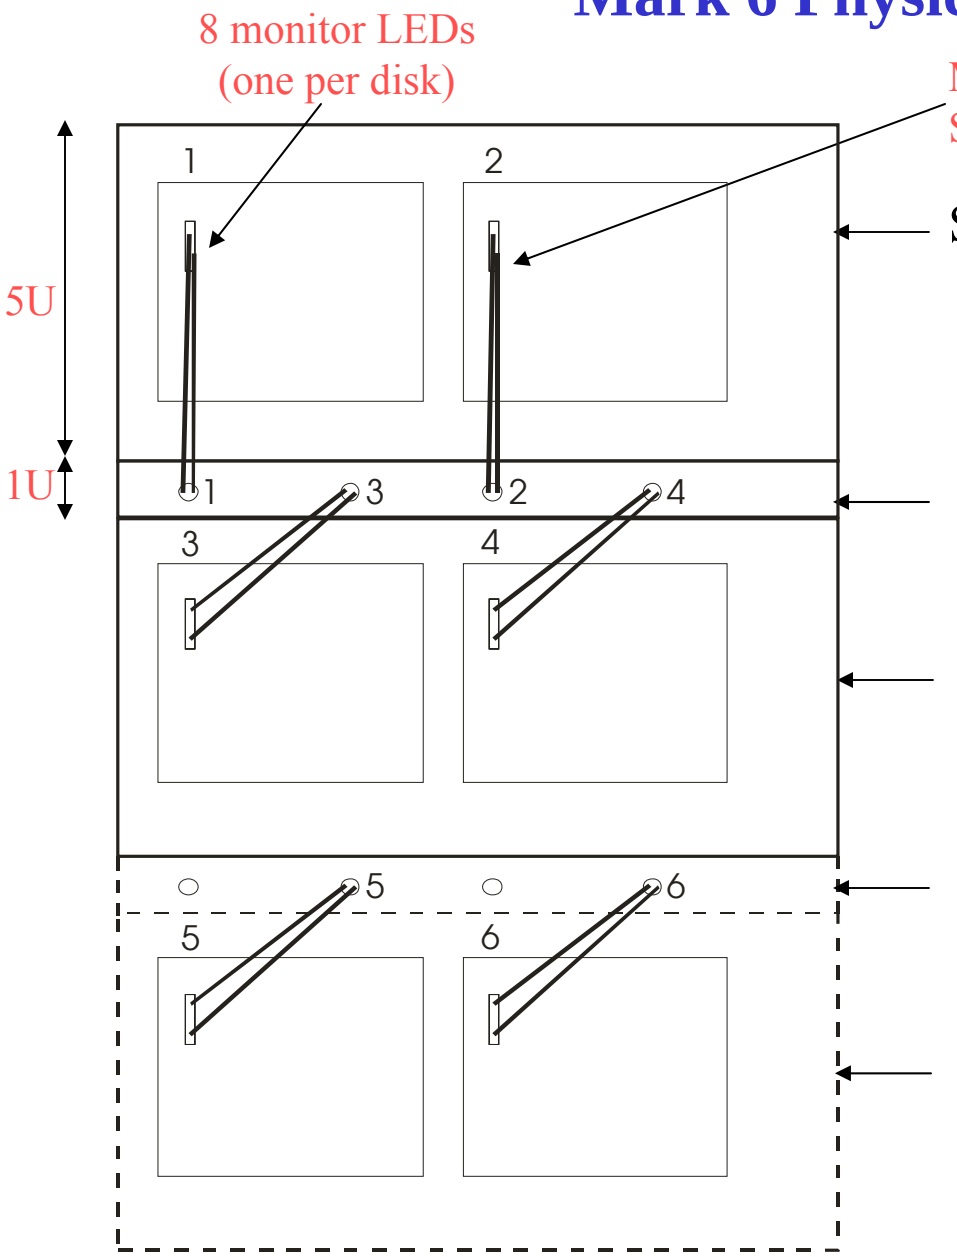

System chassis (supports 8Gbps by itself)
- Mark 5/6 enclosure
- Power supply
- Chassis backplane kit
- Data-electronics hardware

Retractable cable panel
(for easy management of eSATA data cables)

Expansion chassis (needed for 16Gbps)
- Mark 5/6 enclosure
- Power supply
- Chassis backplane kit

(Optional) 2nd cable panel

(Optional) 2nd Expansion chassis

Mark 5 SATA Drive Module Upgrade to Mark 6

New Front Panel

Connectors for two eSATA cables

8x LED (1 per drive)

Re-use Handle from old Module

New Latch provided (pre-installed on new Panel)

Cooling slots

New PCB and power connector

Easily removable disks

Front Panel

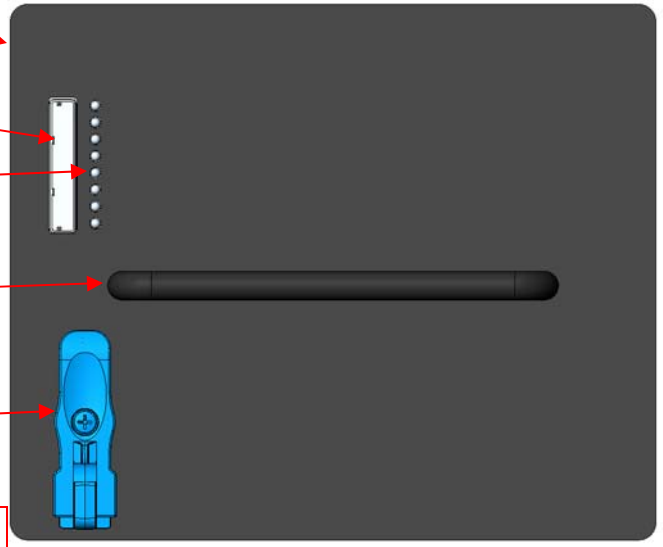

Rear Panel

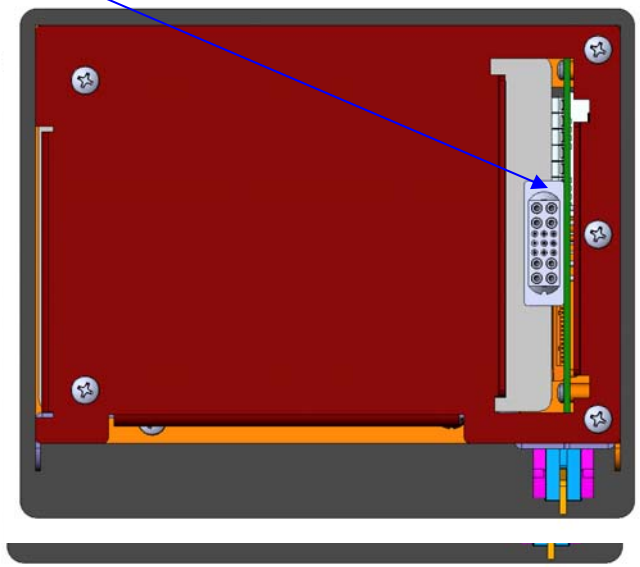

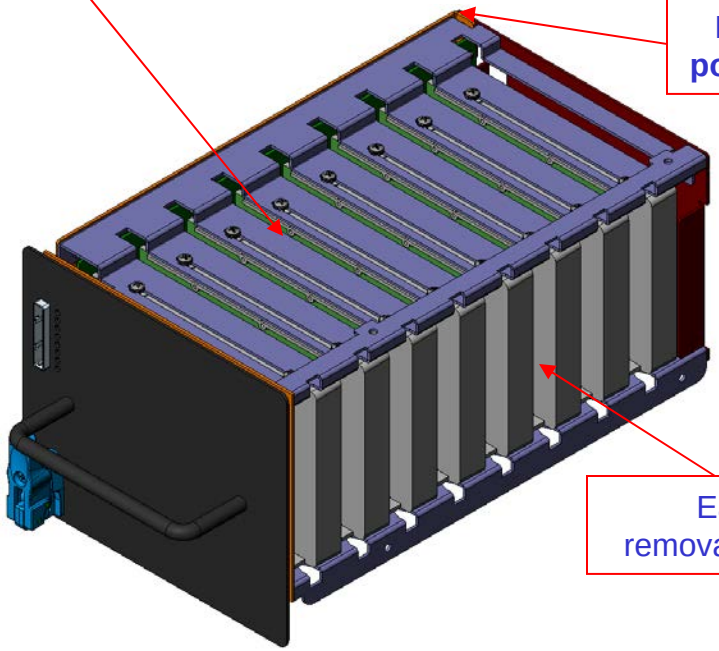

Mark 5 Chassis Backplane Upgrade

New Drive Module Backplane (x2):
-Sequences power to disks
-**Regulates voltage at disk power pins**

New Connector Board:
-simple disconnect to allow easy removal of Module Tray from chassis

Module guide rails

Module Tray

Cooling fans

Connections to chassis Power Supply

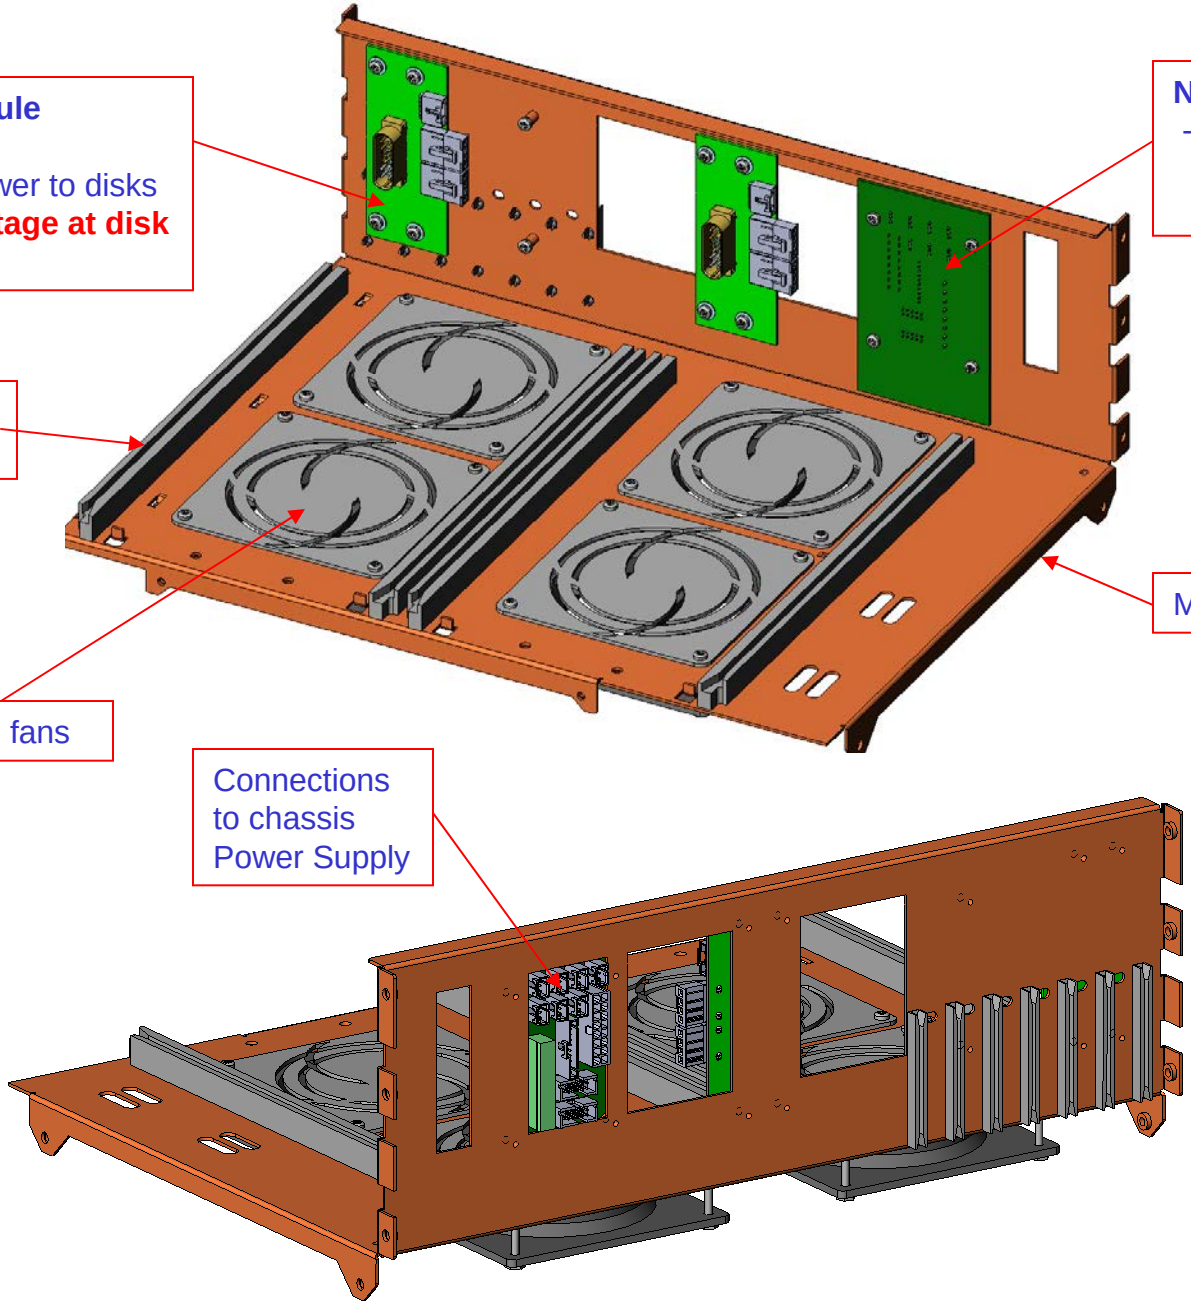

Mark 6 Challenges

- Choose best hardware (our partners at NASA High End Computer Networking generously provided the entire hardware specification based on extensive NASA/GSFC testing)
- Optimize settings such as interrupt-to-processor mapping and process-to-processor mapping
- Control-plane integration
  - Implement full-set of operational controls
  - Minimize stress of transition from Mark 5 to Mark 6
- Thorough testing in real-world environment

Open-Source Software

- Mark 6 code is fully GPL'ed and available to the community
- Advantages:
  - No proprietary software
  - Many eyes to understand the system, add capabilities, and integrate into VLBI operational systems
- Plan to create an Mark 6 open-source support and development group

Have a look at the Mark 6 source web sites:
http://www.vdas.org - for general documentation
http://code.vdas.org - for code repository

No New IP in Mark 6

- Much of Mark 6 is based on publically available software libraries:
  - Boost C++ library
  - CPPUnit testing framework
  - FIO disk benchmarking software
  - PF_RING libraries
- ….and uses standard languages:
  - Python
  - C++
  - C
- VDAF file system is an extension of a standard file system to deal with slow and/or failed disks and uses concepts used for many years
- Bottom line: Writing software for 16Gbps data capture and recording is not rocket science, but
  - Does require considerable attention to details of choice of hardware, system configuration, and software optimization
  - Large part of challenge is create a highly reliable, flexible system that integrates well into the normal VLBI/e-VLBI operations

Mark 6 M&C and concepts

- VSI-S command set
- Recording units are defined as ‘volumes’, each of which consists of one or more physical disk modules
  - Multi-module volumes are required for recording rates >~4Ggps
  - Multi-module volumes retain identity thru correlation processing, then are returned to single-module volumes
- Volumes are managed on an ordered ‘Volume Stack’ that allows multiple volumes to be mounted simultaneously
  - Allows volumes to be queued in specific order for usage
  - Supports automated switchover to next volume in Volume Stack when current volume becomes full; switchover takes place between scans
- Disk statistics gathered during recording allow easy identification of slow/failing disks by disk serial number

Mark 6 Project Status

- Sustained 16Gbps from four 10GigE interfaces to disk has been demonstrated for more than 10 minutes with no packet loss
- Code is extensively instrumented for performance evaluation
- To be done:
  - Complete the VSI-S command set (operational skeleton exists)
  - VLBI Disk-Adaptive Format (VDAF) file system
  - Playback as standard Linux files (current implementation is RAID-based)
- Prototype deliveries expected Q1 2012 from Conduant:
  - Mark 6 chassis-backplane boards
  - Mark 5-to-Mark 6 SATA disk module upgrade kit

16 Gbps VLBI demonstration with Mark 6

24 October 2011

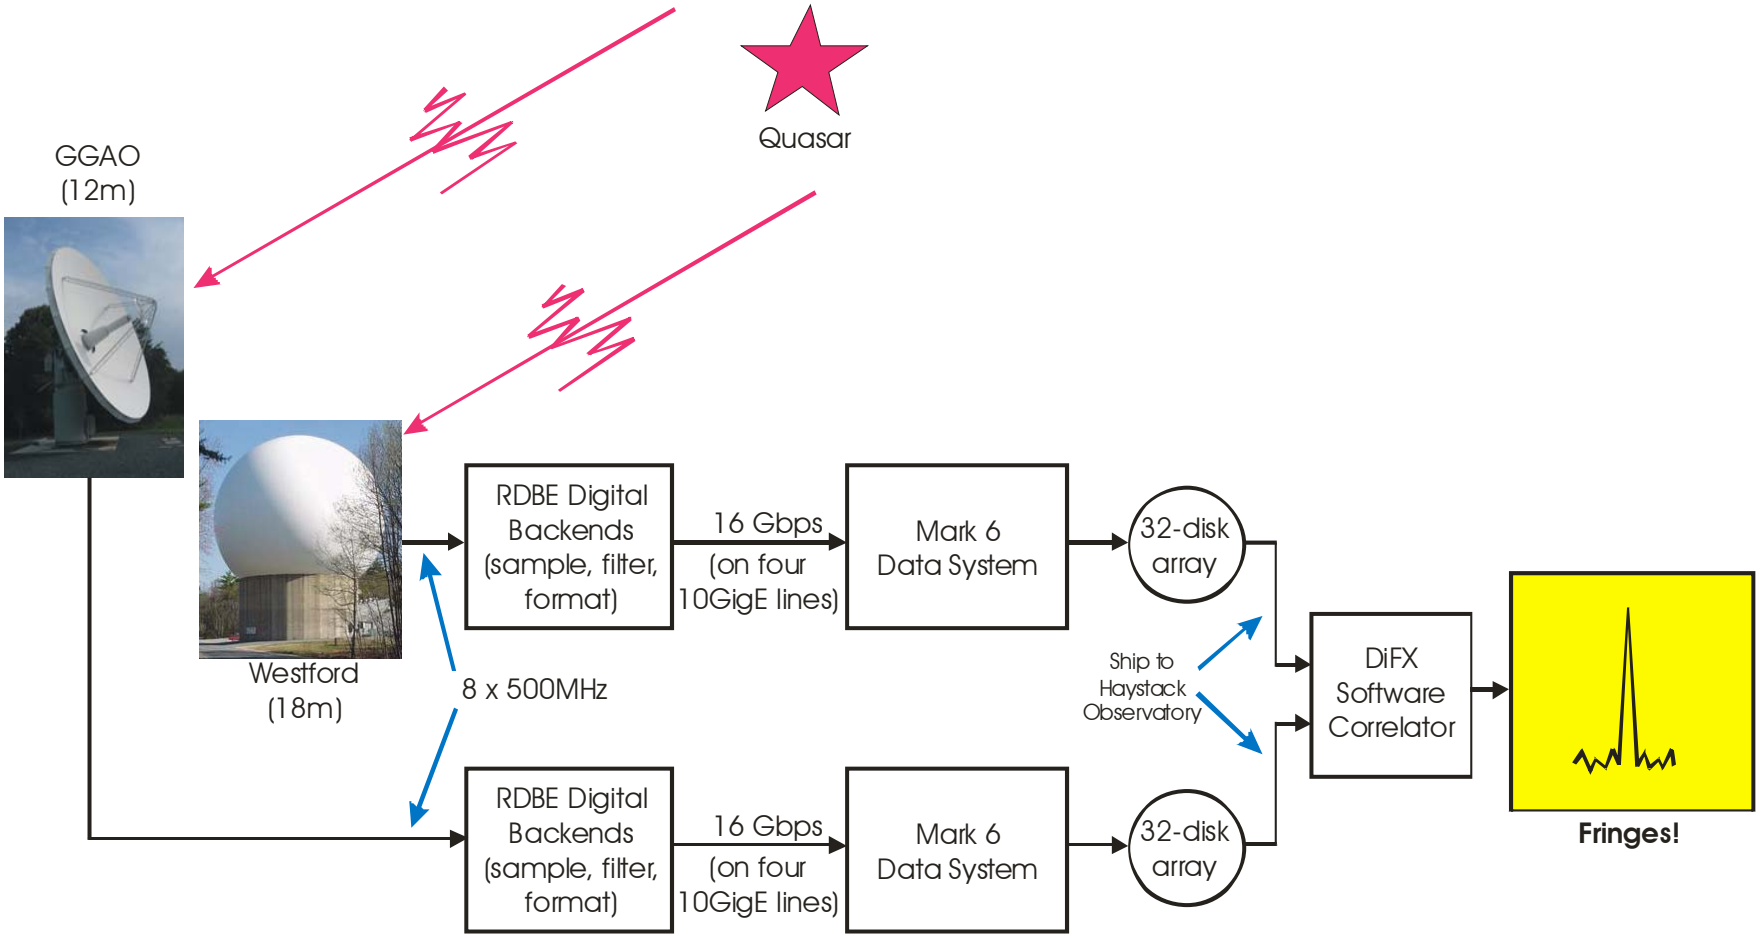

Mark 6 16Gbps demonstration system

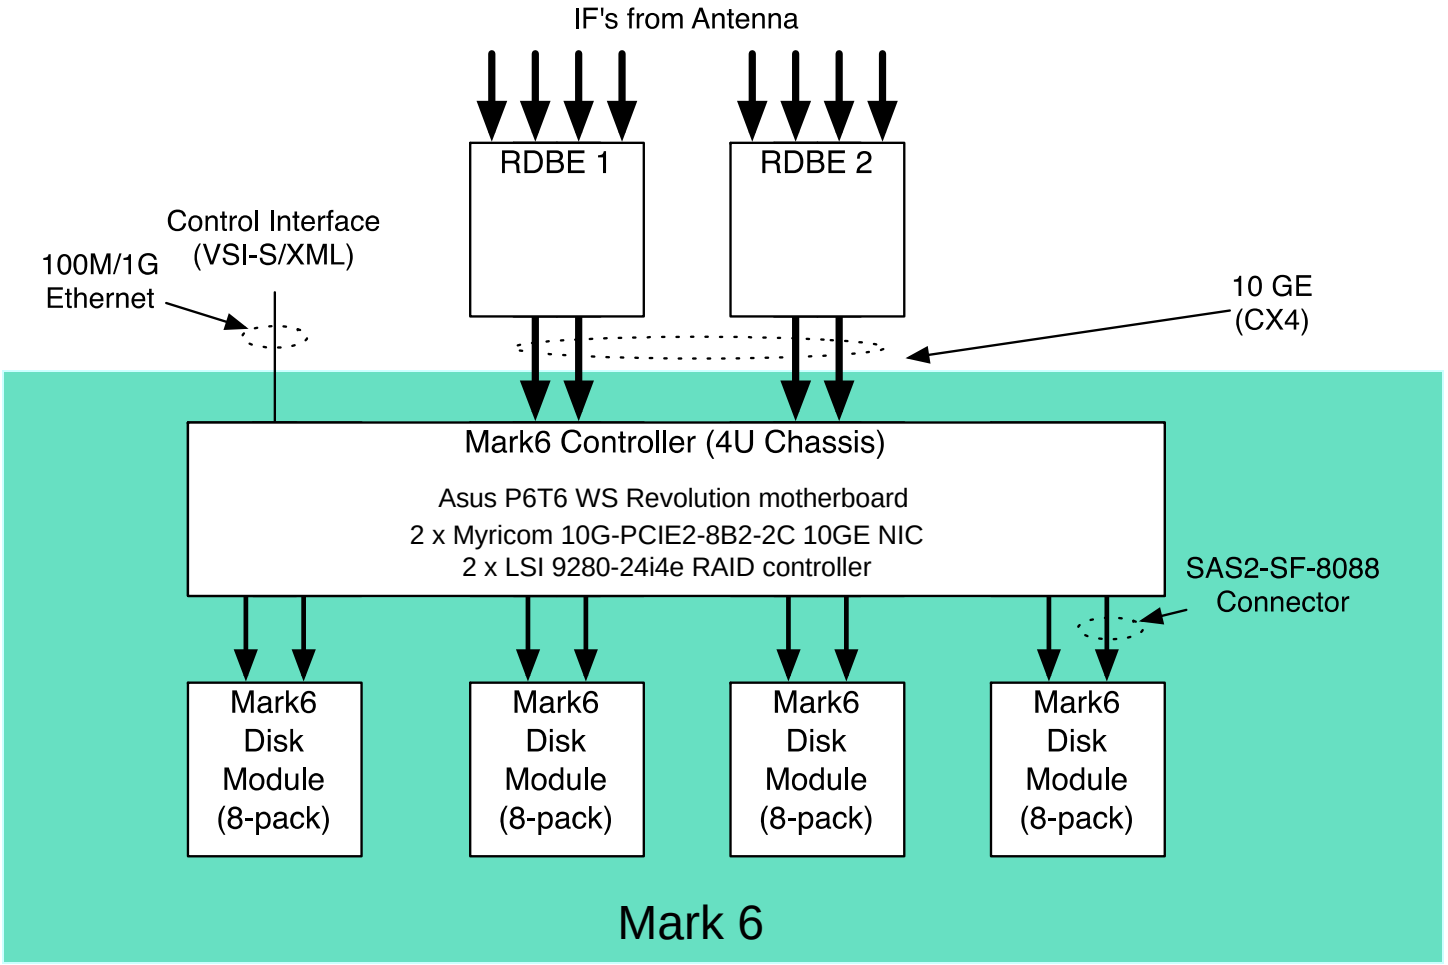

Mark 6 Ethernet read rates

24 October 2011 (GGAO)

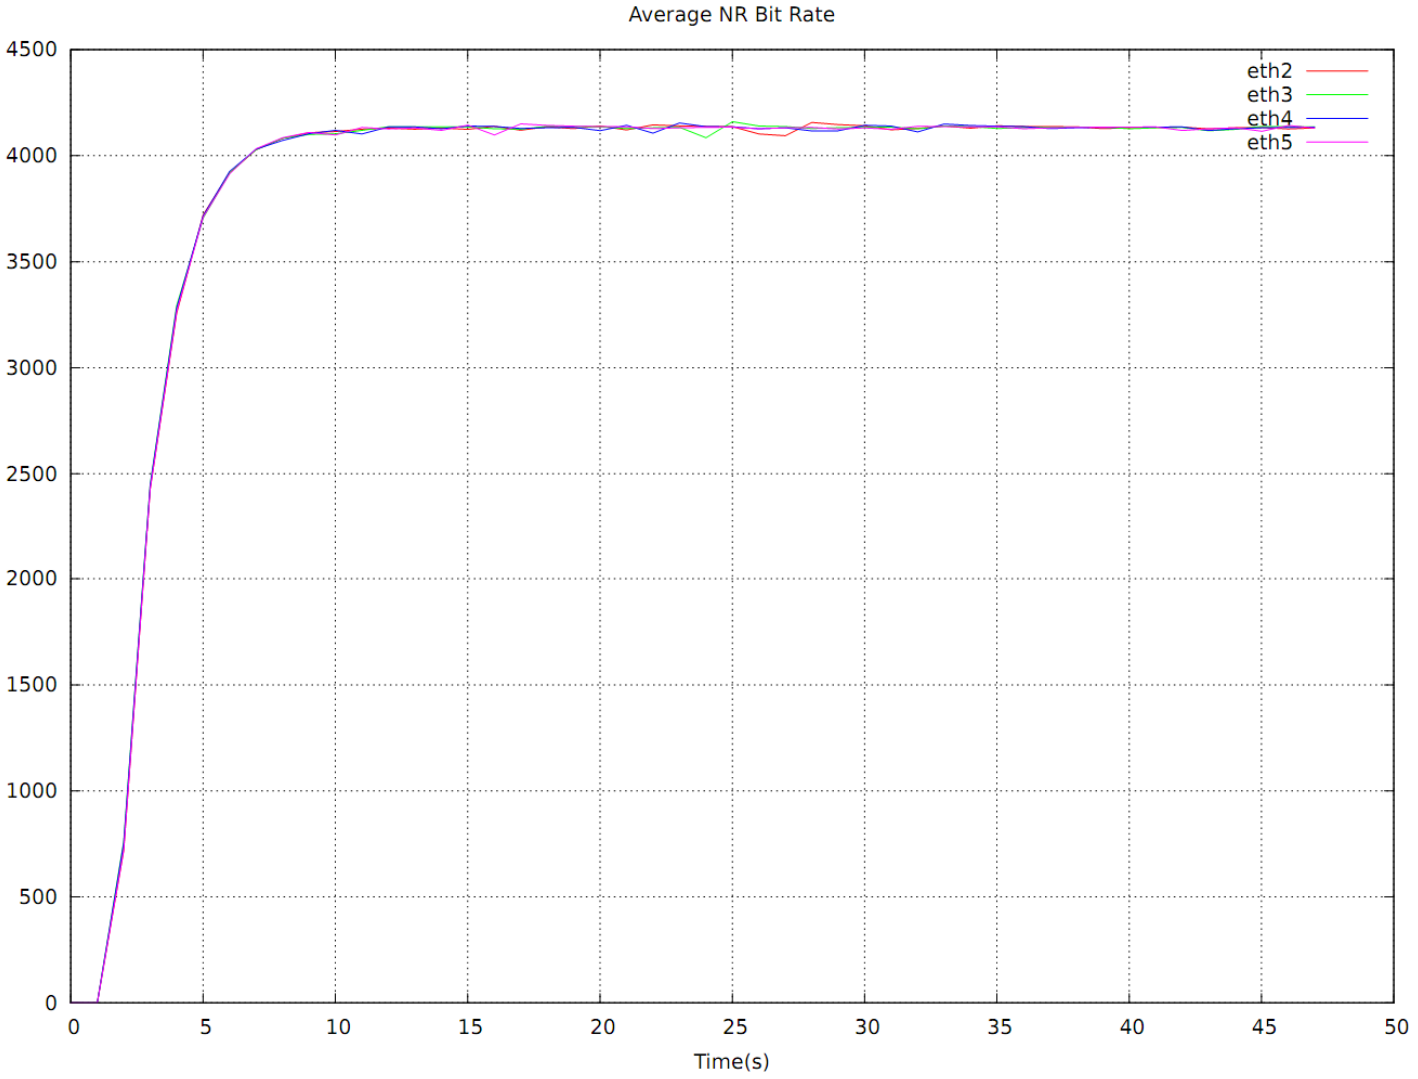

Mark 6 file-write rates

24 October 2011 (GGAO)

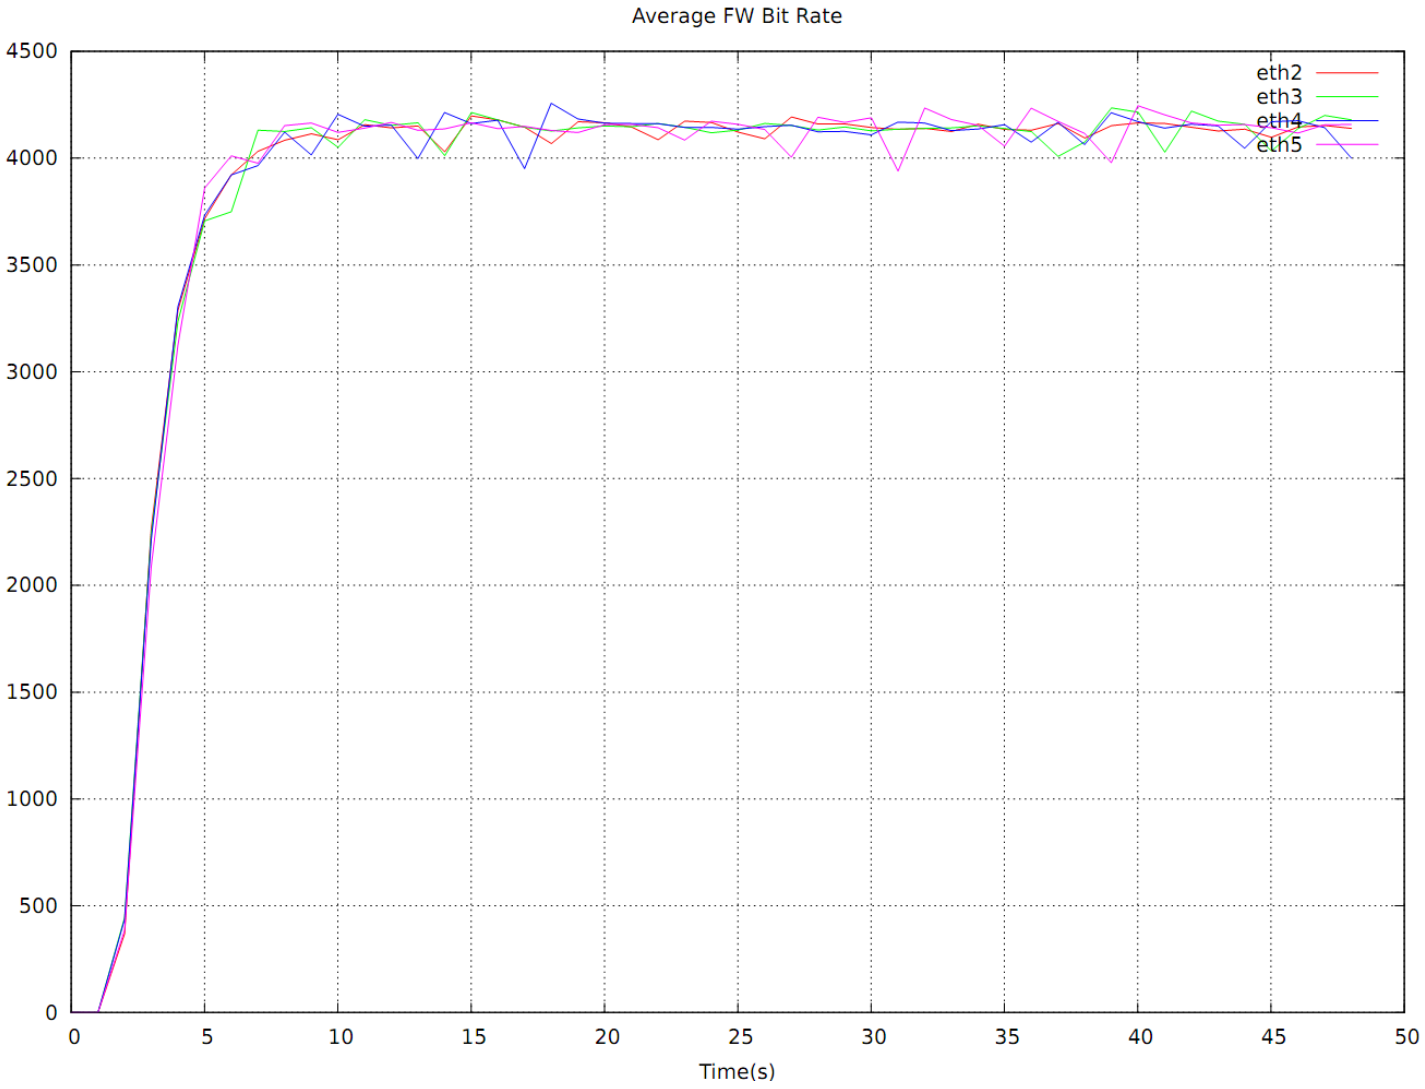

Correlation results (single 500MHz channel)

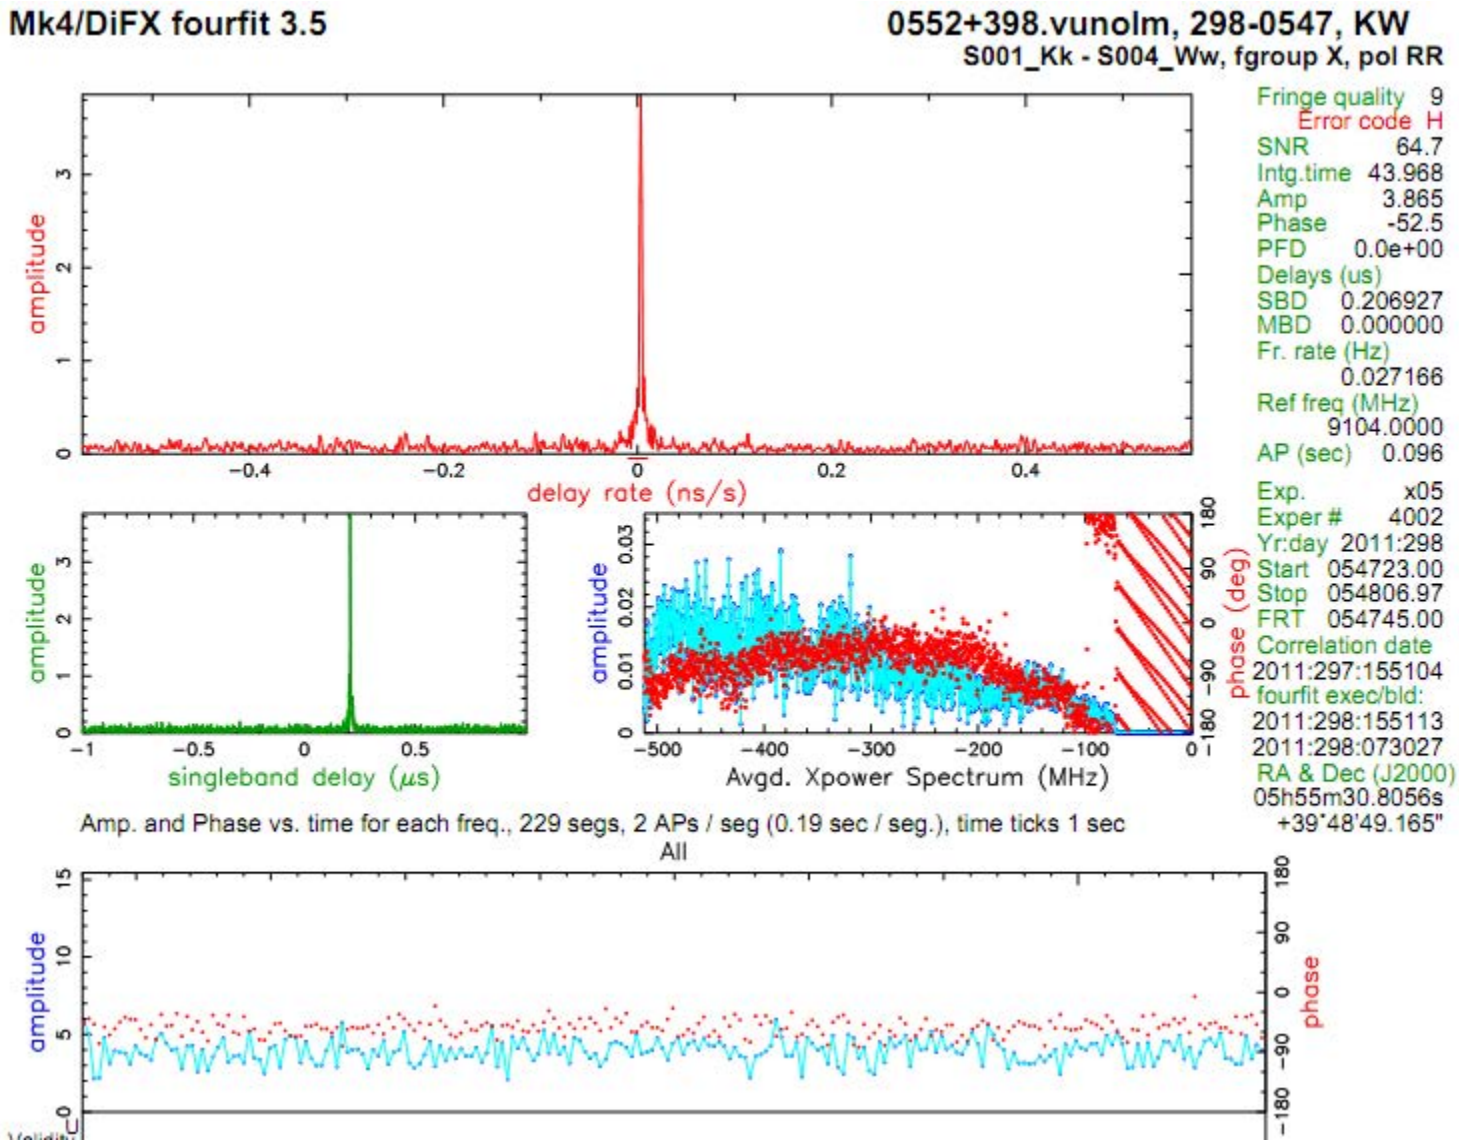

Projected schedule

- Nov 2011 – GGAO/Westford test with broadband system (dual-pol with 2GHz BW/pol)
- Dec 2011 – Test Conduant prototype hardware; integrate complete hardware system
- Jan 2012 – Mark 6 systems orderable
- Jan 2012 – Begin integration with Field System
- May 2012 – System complete and fully tested; ready to deploy
- Ongoing - Integrate Mark 6 with DiFX correlator

How will Mark 6 be available?

- Several options:
  - Purchase new Mark 6 system from Conduant
  - Have Conduant upgrade existing Mark 5
  - Upgrade existing Mark 5 yourself
    - Purchase chassis-upgrade kit from Conduant
    - Purchase electronics yourself (must be exactly as prescribed)
  - Purchase new expansion chassis; add own electronics
  - Purchase Mark 5 SATA-module upgrade kits as needed
  - Purchase Mark 6 modules (with or without disks)

Beyond 16 Gbps.....

- mm-VLBI
  - Habitually starved for sensitivity due to weak sources, short coherence time, and small collecting areas; system bandwidth capabilities are being expanded at telescopes
  - Phased-ALMA project will bring 16GHz bandwidth (64Gbps) potential to mm-VLBI by ~2015 and will almost certainly be exploited
- Geodetic-VLBI
  - Inherent VLBI2010 bandwidth of >10GHz allows much higher sensitivity with larger captured bandwidth
  - Higher sensitivity greatly expands number of suitable sources for more uniform sky coverage
- cm-VLBI
  - Larger bandwidths more problematic due to RFI, but flexible digital backends may to avoid (or deal with) RFI, particularly at higher RF frequencies

Thank You's

- Haystack/Westford –
  Chris Beaudoin, Pete Bolis, Roger Cappallo, Shep Doeleman, Geoff Crew, Rich Crowley, Dave Fields, Alan Hinton, David Lapsley, Arthur Niell, Mike Poirier, Chet Ruszczyk, Jason SooHoo, Ken Wilson
- NASA/GSFC VLBI Group –
  Tom Clark, Ed Himwich, Chopo Ma
- NASA/GSFC GGAO –
  Roger Allshouse, Wendy Avelar, Jay Redmond
- NASA/GSFC High-End Computer Networking Group –
  Bill Fink, Pat Gary (recently deceased), Paul Lang
- Conduant –
  Phil Brunelle, Greg Lynott, Ken Owens

Questions?

Backup slides

Recording rate capability vs. time

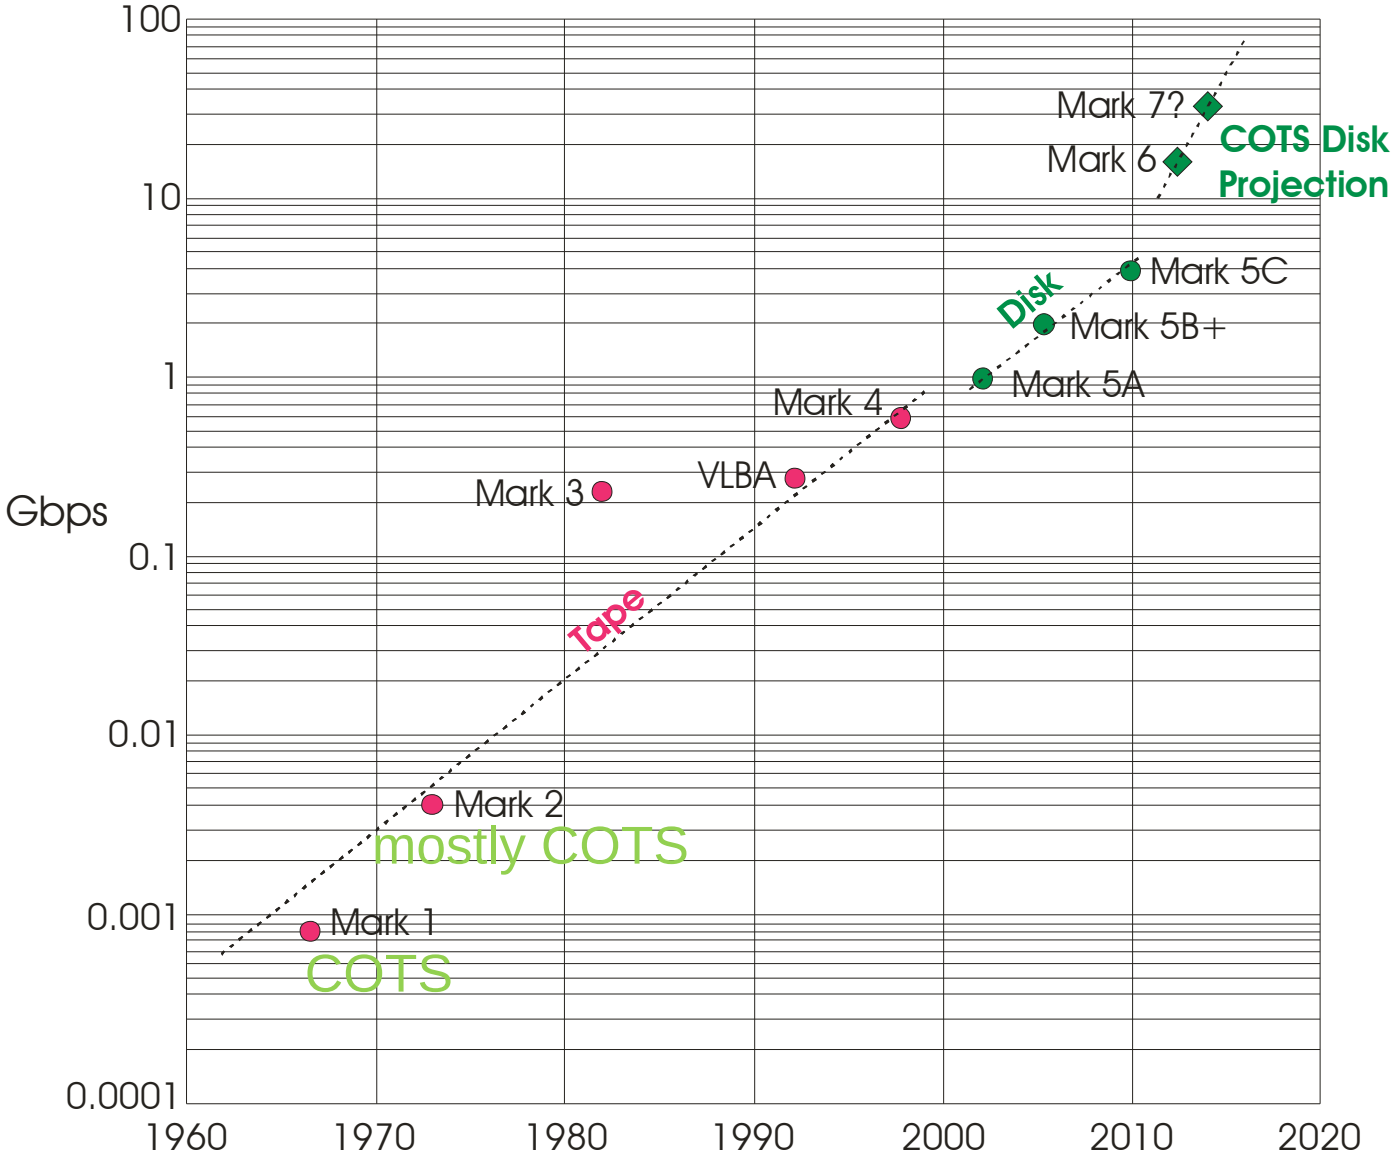

Recording-rate cost vs. time

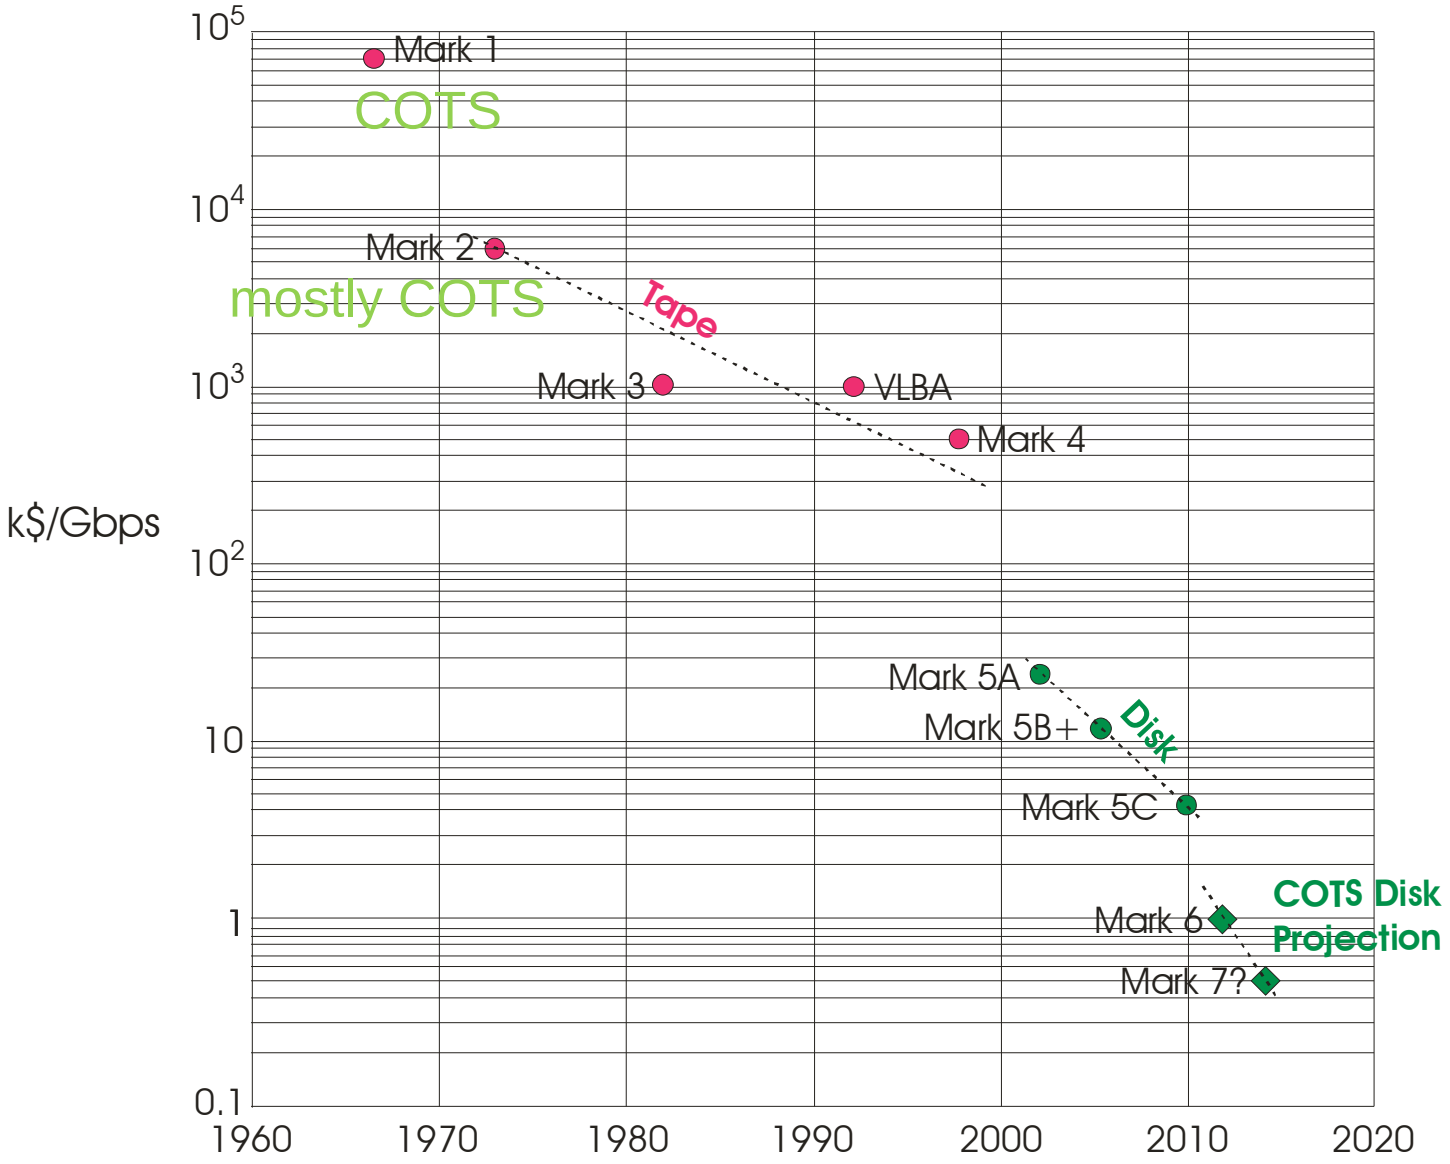

Backend-bandwidth cost vs. time