Low-Cost Disc-Based Real-time Gbps Data System

Alan R. Whitney
MIT Haystack Observatory
awhitney@haystack.mit.edu

- Application
  - Radio astronomy data collection for Very Long Baseline Interferometry observations
  - Data collected at individual antennas, then combined to form single antenna
  - Signal-to-noise ratio of observations is proportional to data rate, so highest data rates are sought
  - Typically 1 Gbps continuously for 24 hrs at 10 stations (110 TB/day), repeated several times a month; incompressible data
  - Data recycled after processing (~1 month cycle time); initial processing reduces data volume by factor 10^3 to 10^4.
- Problem
  - Currently using in-house-designed magnetic-tape equipment for continuous 1 Gbps, but expensive, aging, unreliable, difficult to transport
  - Media cost ~$1000/station/hr
  - Need to develop new low-cost system based on commercial technology
- Goal of this short presentation
  - Briefly describe problem, work to date
  - Connect with others who may have similar needs

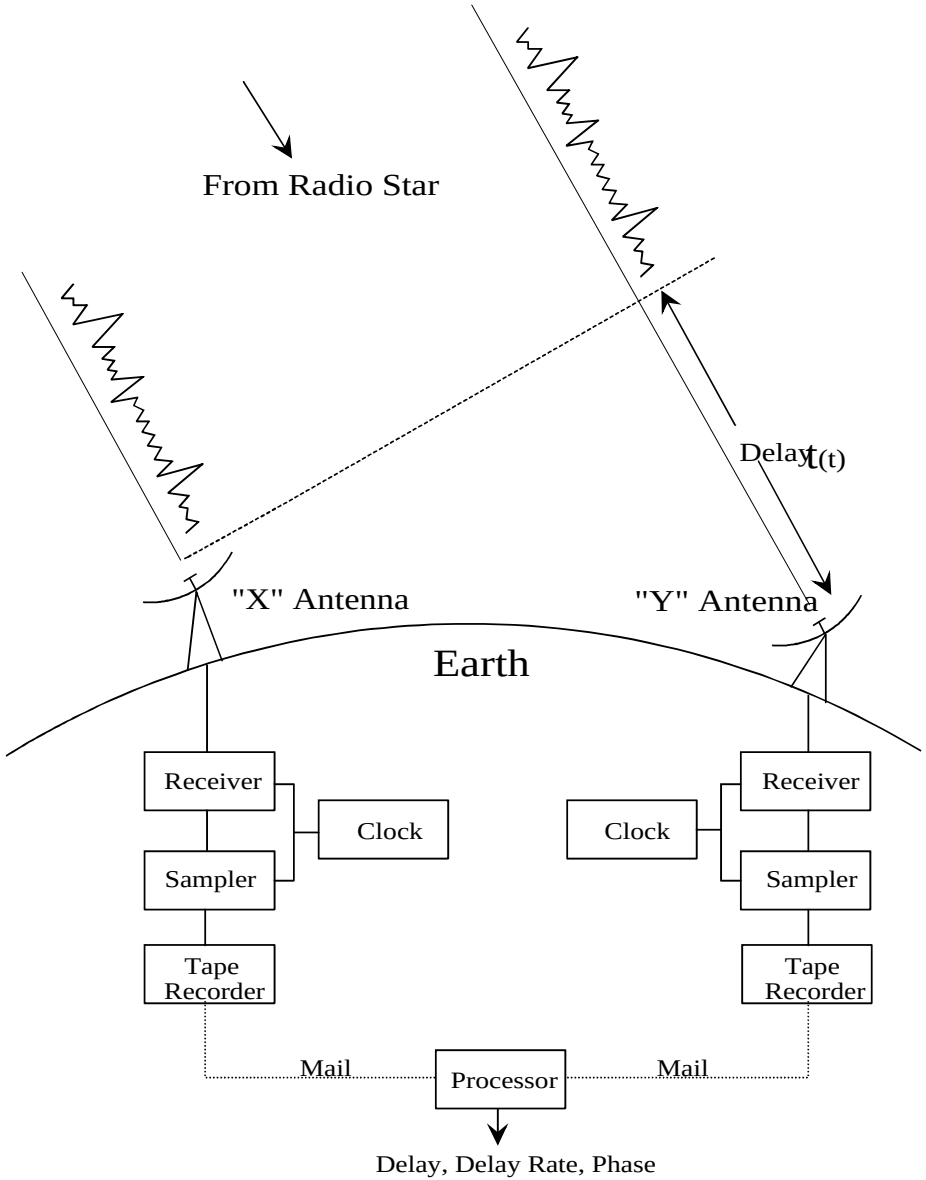

Figure 1: Basic VLBI Block Diagram

Goals of System Design

- Minimum of 1 Gbps data rate
- Based primarily on unmodified COTS components
- Low-cost
- Modular, easily upgradeable
- Robust operation, low maintenance cost
- Easy transportability
- Conformance to VSI interface specification
- Gb Ethernet interface
- Minimum 24-hour unattended operation at 1 Gbps

COTS Candidates for System Design

- COTS Buses and Interfaces
  - VME
  - PCI
  - SCSI
- COTS Computer Platforms
  - PC
  - Motorola 68000 family
- COTS Recording technologies
  - Magnetic tape (initial focus)
  - Magnetic disc
  - Optical disc
- COTS OS's
  - Linux
  - Windows
  - Various real-time OS's (Lynx, VxWorks, etc)

Tape vs. Disc Price Comparison

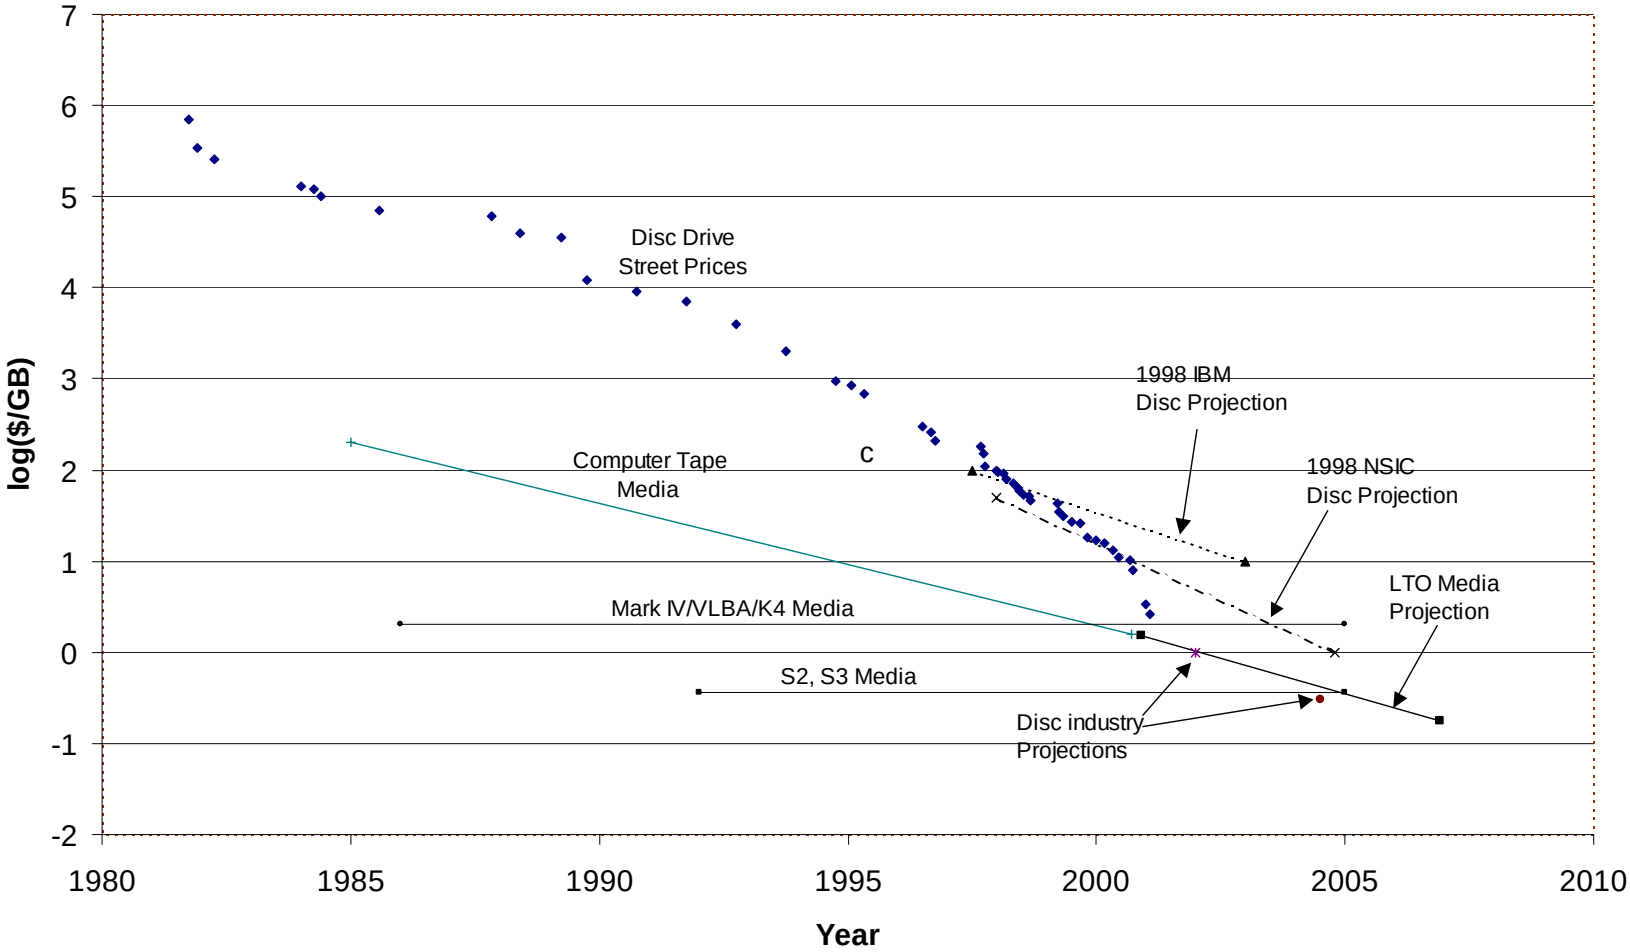

Media Selection

- Expect IDE-interface magnetic hard-disc prices will fall below computer magnetic-tape prices within ~1 year based on $/GB; SCSI-based discs remain significantly more expensive
- IDE hard disc price vs capacity/performance will continue to drop rapidly
  - Now near $3/GB, expected to drop to ~$1/GB within ~12-18 months, to ~$0.30/GB by ~2004-5
  - More likely to find COTS-based solution with magnetic discs than with magnetic tape or other media
- *Conclusion: Next generation systems should be based on magnetic disc technology using standard IDE (aka EIDE, ATA) interfaces*

Other Advantages of IDE Magnetic Discs

- Readily available inexpensive consumer product; continually improving in price/performance
- Self contained; don’t have to buy expensive tape drives, so host system can be inexpensive
- Rapid random access to any data
- Essentially instant synchronization on playback
- Unlike tape, no headstacks to wear out or replace – ever!

Mark V Hardware Components

- Standard PC in rack-mount case
  - 256 MB main memory
  - Standard PCI bus
  - 800 MHz Pentium 3
- General Standards HPDI32 digital I/O card
  - 32-bit wide I/O at up to 20 MHz clock rate
  - 1 MB FIFO buffer
- Boulder Instruments StreamStor multiple-disc interface
  - Support up to 16 IDE discs at streaming data rate to 800 Mbps
- 2 – StorCase 9-disc chassis (8 used)
- 16 – IBM Deskstar 46-GB disc drives
  - Each mounted in a hot-swappable DataPort V carrier
- Total system cost <$25K

Mark V VLBI Demonstration System

Mark V System Block Diagram

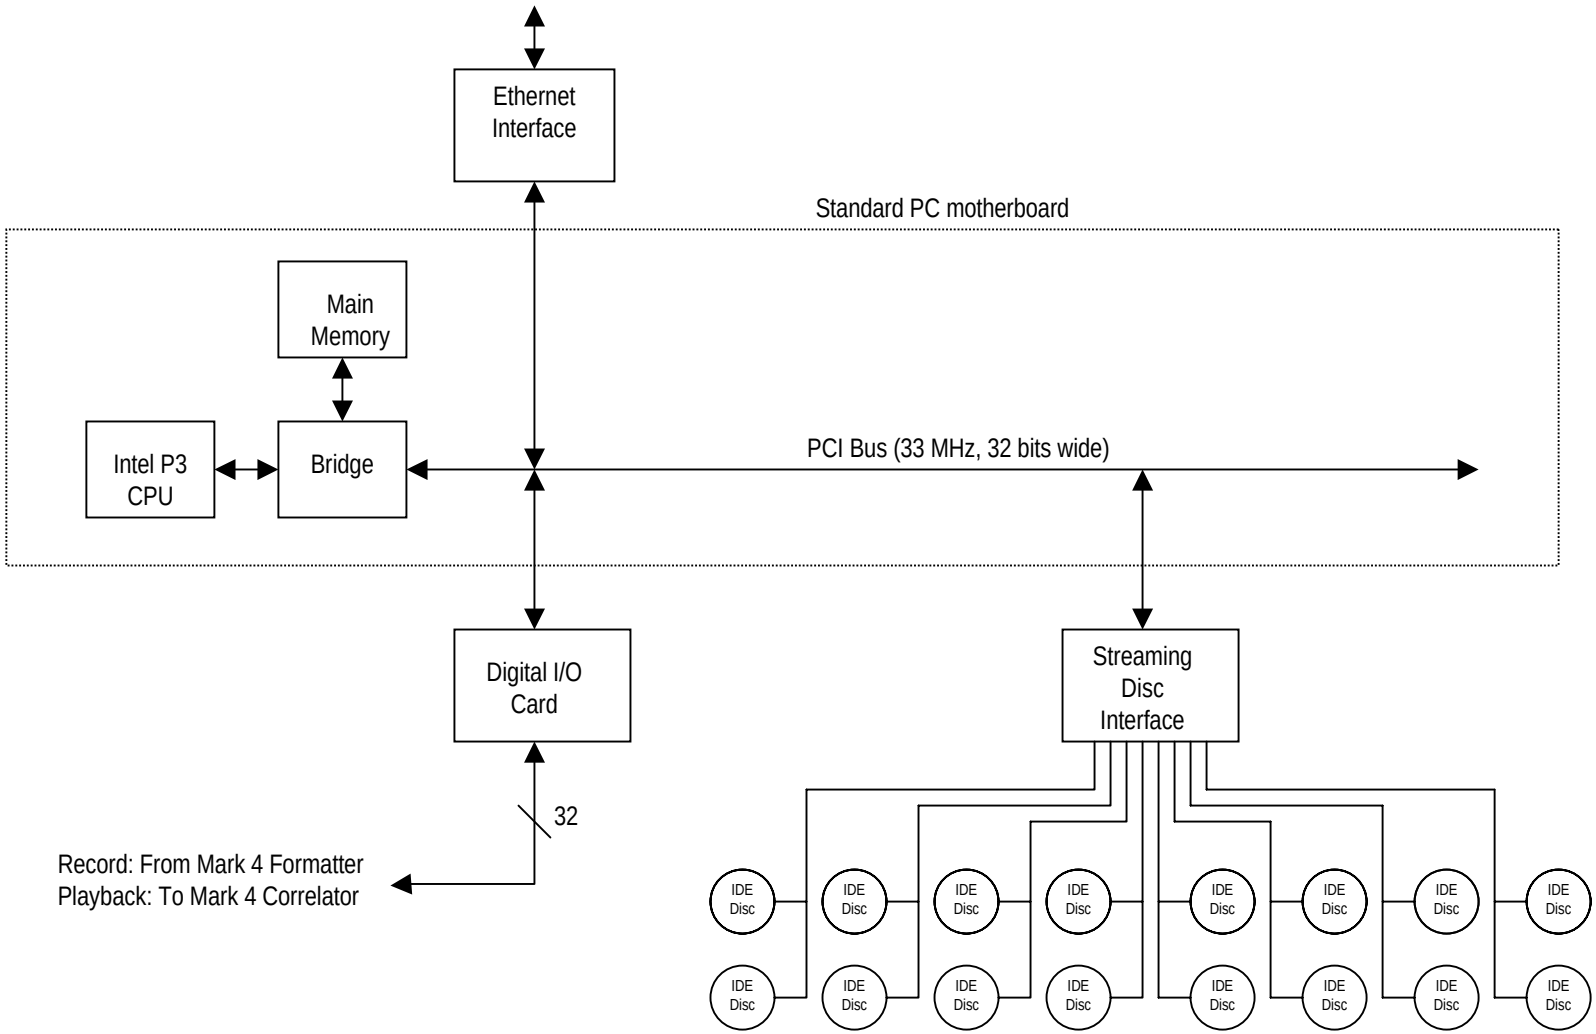

System Characteristics

- 736 GB disc storage (with current 46 GB discs)
- Recording/playback – 32 parallel bit-streams at up to 18 Mbps/bit-stream (576 Mbits/sec)
  - Direct PCI-to-PCI transfer from I/O card to disc interface card
- Only four simple manual functions implemented
    - *Record* – start recording from beginning of media
    - *Append* – append recording to end of existing recording
    - *Play* – reproduce starting at specified byte position
    - *Fetch data to file* – fetch specified data to standard Windows file(s)
- All functions fully tested
- Conception to completed system in <2 months

Status and Plans

- Upgrade to 1 Gbps within 6-8 months
- Add VSI interfaces for general-purpose real-time streaming with full time resolution (to the bit level)
  - VSI interface has been internationally adopted by VLBI community, Japanese high-energy physics community
- Costs
  - Within ~1 year, expect disc costs to drop to ~$1/GB with ~125 GB single-drive capacity (tape media cost is ~$2/GB)
  - Industry projection ~1 TB drive for ~$300 ($0.30/GB) by ~2004-5. A single Mark V with 16 such drives will record 1 Gbps for ~34 hours unattended!
  - Cost of system is expected to remain <~$25K.

Towards ‘Real’ Real-Time

- Real-time data transmission over dedicated fibers is being pursued, but:
  - Most antennas are deliberately remote and ‘last mile’ costs are high
  - Current costs for Gbps international networks are far too high (the cost/bandwidth of a 747 loaded with tapes still can’t be beat!)
- Data-buffering at both antenna and central processor is likely to be required for some time
- Nevertheless, we expect to begin limited demonstrations of real-time data transmission during the next year.

Conclusions

- Gbps real-time COTS-based system is readily achievable today
- Magnetic-disc technology is rapidly becoming cost competitive with magnetic tape and has many advantages
- Choice of magnetic discs allows use of most modern, least expensive disc technology at all times
- Cost of host system is much cheaper than magnetic-tape system (i.e. no tape drives!)
- Multi-Gbps operation is economical with simple parallel operation of multiple systems
- Comments, questions welcomed at awhitney@haystack.mit.edu