

Concept for an Affordable High-Data-Rate VLBI Recording and Playback System

MIT Haystack Observatory
30 September 1999

Summary

Recent and continuing developments of magnetic-tape data-storage systems for the computer industry provide a basis on which to create a high-performance, affordable high-data-rate (HDR) recording and playback system. Such a system would be suitable for VLBI and other applications which demand sustained data rates of a gigabit/sec (Gbps) or higher. These systems should be based on mostly-off-the-shelf hardware which will be readily available from the computer industry, with a minimum of custom design, and which can easily grow to higher data rates as the recording industry advances. Furthermore, they must be capable of unattended operation for at least 24 hours. Based on currently-announced tape products and projected industry roadmaps, the cost *per Gbps* for such systems will drop from a Year 2000 cost of ~\$100K/Gbps to ~\$14K/Gbps by ~2007. This suggests that HDR systems in the range of ~8 Gbps may be available for ~\$110K before the end of the next decade.

Introduction

The world of sustained very-high-data-rate recording, typified by the requirements of Very Long-Baseline Interferometry, has long been a specialty field inhabited by several generations of specialized recording systems developed mostly by the user community itself. This situation has not been by choice, but of necessity, since commercial machines of the necessary capabilities were either simply unavailable or were available only at prohibitive cost. Though these systems have performed often-heroic service over the years, they have generally suffered from predictable drawbacks of specialized systems: high development and purchase cost, difficult and expensive maintenance, and lack of easily and economically expandable capabilities.

The demand for state-of-the-art VLBI data-recording systems is currently at the 1 Gbps threshold, with demands extending to ~8 Gbps over the next decade or so. Though at least two currently-available recording systems (Mark IV and D-6) are capable of supporting the 1 Gbps rate, both the capital and operating expenses of these systems are high, and the systems are not economically expandable to the needs of the coming decade.

Goals for a Next-Generation VLBI HDR System

The goals for a next-generation VLBI HDR system can be stated fairly succinctly:

- Minimum of 1 Gbps data rate (native, uncompressed)
- Economically upgradeable/expandable to ~8 Gbps over the coming decade
- Design based primarily on unmodified off-the-shelf subsystems and components
- Modular, easily upgradeable as better/cheaper technology becomes available
- Robust operation, low maintenance cost
- Easy transportability
- Conformance to emerging VLBI Standard Interface (VSI) specification

- Flexibility to support electronic transfer ('e-VLBI') and computer processing of recorded data
- Easy adaptability to HDR applications other than VLBI
- Minimum of 24-hour unattended operation

Though this list of goals may appear idealistic, we believe they can be achieved in a system which can be designed today, and which will evolve in the following years to ~8 Gbps with minimal additional engineering effort.

The Computer Industry to the Rescue

The demands of the computer industry have been driving several technologies that are key to meeting the stated goals of the next-generation VLBI HDR system. We shall examine each of these key technologies.

Tape Recording Technology

Although developments in tape technology has seemingly taken a back seat to disc drive technology in recent years, that situation is rapidly changing. This change has been brought about by the increased need of users, particularly corporate users, to backup vast and ever-increasing amounts of data stored on disc systems. In today's market of inexpensive 20-50 GB disc drives, it is not unusual for corporate users to maintain many TeraBytes (1 TB=1000 GB) of on-line information, all of which must be backed up at regular intervals. The most economical way to back up these large systems is with magnetic tape. But the amounts of data are so vast that practical tape systems must have both very large storage capacity *and* a commensurately high data rate to match shrinking backup windows.

Figure 1 shows the historical progression of the native storage capacity of a single tape for several commercial systems, both of the helical scan and linear variety, and all of which use a tape cartridge format. Also shown is the industry projection for storage capacity over the next 7-8 years. We note that today's single-tape capacity of roughly 100 GB is expected to rise to ~1 TB by ~2007; capacities of 0.5 TB/tape have already been demonstrated in commercial laboratories. Though projecting future technology is clearly a risky business, some credibility to these projections may be gained by noting that the disc drive industry has handily outstripped nearly all predictions made for that industry over the past 5 years; tape technology is quickly picking up steam and is now beginning to move very aggressively in the same direction.

Figure 2 shows the history of single-drive data rates for the same set of systems, as well as industry projections for single-drive data rates over the next 7-8 years. The currently available data rates of 10-12 MB/sec will rise to 16-18 MB/sec in announced products available within the next 4-8 months, and are expected to rise to the order of 128 MB/sec/drive by ~2007, roughly mirroring the growth rate of data storage capacity.

Current costs of a single tape drive, regardless of type, is on the order of \$5-7K each in small quantities. Based on past experience and on industry projection, the cost per drive will remain relatively constant as the technology progresses over the next 7-8 years.

These tape systems are computer grade systems with high-reliability and extremely low error rates due to extensive use of error-correcting codes and redundancy. Though such low error rates are not needed for VLBI, the costs of circumventing such error-correcting codes to gain higher data rates and capacities is probably not worth the effort or cost.

Whatever the details of the tape technology to be used, whether linear or helical-scan or something else, the implications are clear. Within the next year, 8 *standard commercial tape drives* at 16 MB/sec each can, in principle, be paralleled to achieve 1 Gbps data rate for an investment of <\$60K in tape drives. In 7-8 years, as shown in Figure 2, a *single* tape drive is expected to be capable of 1 Gbps (128 MB/sec), so that a bank of 8 such machines will be able to achieve 8 Gbps for the same <\$60K investment in tape drives!

Industry Standard Buses and Interfaces

Of course, tapes and tape drives aren't the only component of a recording or playback system, and projections based on that aspect alone are obviously not the whole story. In this section we will investigate the state of standardized buses and interfaces that will be needed to support the tape drives.

The PCI Bus and Its Derivatives

The most common bus in use in modern computers is undeniably the PCI bus – almost every modern personal computer implements a PCI bus in one form or another. Its ubiquitous nature is supported by thousands of various boards which interface to the PCI bus, and its capabilities are impressive. Modern PCI implementations are based on a 32-bit-wide architecture operating at a 66 MHz clock speed with data transfer rates to ~250 MB/sec. On the near horizon are major extensions to the PCI bus, particularly the PCI-X specification, which operates on a 64-bit bus at speeds to 133 MHz for data transfer rates to ~1 GB/sec and is backwards compatible to present PCI-bus devices. The PCI-X bus is supported by several major computer companies (Compaq, HP, IBM), with first products expected on the market at the end of 1999.

Another candidate for a foundation bus is the Compact PCI bus, which is functionally identical to the PCI bus but is implemented in a standard 6U-height card with a different connector. Currently, the Compact PCI bus is specified only to 33 MHz clock speed, though this may be upgraded in the future. Standard Compact PCI computer-plug-in and interface cards are readily available. One advantage of the Compact PCI specification is the allowance of up to 8 plug-in cards along the bus, compared to 4 for the standard PCI bus.

In short, the PCI bus or one of its derivative successors appears to be an ideal candidate for a bus to support data rates demanded by multi-Gbps systems. In any case, PCI bus technology is pervasive, cost-effective, and will be around for many years, so it is a safe bet on which to build the foundation of a new VLBI HDR system.

SCSI Bus

The venerable SCSI bus continues to thrive and evolve as the dominant high-speed small-computer interface. Beginning more than ten years ago with SCSI-1 at a maximum data rate of 5 MB/sec, many devices now incorporate the well-supported Wide Ultra2 SCSI interface rated at 80 MB/sec maximum. And on the horizon is the Ultra3 SCSI interface rated at 160 MB/sec maximum. Because a single SCSI bus can support multiple devices and because a PCI bus can support multiple SCSI interfaces, easy choices can be made to optimize virtually any combination of bus and device capabilities.

Based on its wide support and significant capability, the SCSI bus is an obvious choice for a next-generation HDR data system. All of the tape drives mentioned earlier incorporate modern SCSI interfaces and are expected to evolve to higher-speed flavors of SCSI as they become available.

VLBI Standard Interface

The VLBI Standard Interface (VSI) specification is currently being developed by an international consortium of VLBI instrumentation developers drawn from both the astronomy and geodetic-VLBI communities. Progress to date has been gratifying, with the expectation that VSI will be ratified by the end of 1999. The VSI specifies a standard interface to and from a VLBI *Data Transmission System (DTS)* in such a way that *non-homogeneous* recording and playback systems can be used simultaneously in the collection and processing of VLBI data. In any case, any new VLBI DTS system should adhere to the VSI specification.

The VSI specification in its present draft form is modular at the 1 Gbps level in that a single connector will carry up to 1 Gbps of aggregate data-rate. This matches well with current VLBI demands and is easily expandable to any degree with multiple connectors.

Concept for an HDR VLBI Recording System

Given the above discussion, the outlines of a new mostly-off-the-shelf VLBI HDR recording system are undoubtedly obvious. Figure 3 shows a simple block diagram of such a system based on the standardized components and interfaces discussed above. A 64-bit-wide PCI-X bus operating at 133 MHz forms the backbone of the system, supporting a data transfer rate of up to 1 GB/sec (8 Gbps). A standard CPU, probably of the high-speed Intel or Motorola variety, along with its attendant semi-conductor memory, will attach to the bus for control and monitor. Two standard SCSI interfaces are shown attached to the PCI-X bus, each of which supports two standard commercial recording devices; for most VLBI applications these devices would be magnetic tape recording devices of the types described above, but they could also be other types of recording devices such as optical and/or magnetic discs and tapes. The important point is that these are all standard off-the-shelf devices with standard interfaces.

The *only* custom hardware in the HDR recording system is a plug-in VSI formatter card which implements the VLBI Standard Interface specification and acts as a PCI-bus master device to transfer data *directly* to the SCSI recording interfaces with no interaction with the CPU or its memory. Initial investigation suggests that a single VSI formatter card should easily handle at least 1 Gbps and possibly substantially more.

The software operating system for the HDR system will also be a standard off-the-shelf product, chosen from among several potential candidates after careful investigation. Possible contenders in this arena are such systems as Linux, LynxOS, OS-9, QNX, VxWorks and pSOS. The final choice will depend on many considerations, from robustness of implementation to ease of application development. Networking capabilities will be of significant importance, both for future support of e-VLBI and for remote operation and monitoring.

Concept for an HDR VLBI Playback System

Figure 4 shows a block diagram of a complementary HDR playback system. The hardware is identical to the recording system except a custom VSI *de-formatter* card replaces the VSI formatter card for collection of the data from the SCSI devices for transmission to the VLBI Standard Interface and the user. Alternate data paths exist for collection of the playback data into the computer memory for local processing or transmission through a standard network interface (e-VLBI).

What Can Be Done Today for What Cost?

A survey of the computer marketplace today allows us to estimate the performance and cost of an example system that can be built with hardware available today or in the very near future. For a recording device, we can choose one of the soon-to-be-available systems that supports a data rate of at least 16 MB/sec (see Figure 2). These devices, whether using linear or helical scan recording, can sustain a continuous recording data rate of 128 Mbps for at least 104 minutes on a single tape cartridge, at a cost of ~\$7000/device. If we implement the system shown in Figure 3 using these devices, we can make a conservative estimate of the total hardware system cost:

Computer, memory, backplane	\$5,000
SCSI interface cards (2)	\$2,000
VSI Formatter card (1)	\$2,000
Tape drives (4)	\$28,000
Cabinets	\$1,000
Contingency	\$2,000
<u>Total</u>	\$40,000

System capability: 512 Mbps for ~104 minutes

Tape cost: 4 tape cartridges at ~\$70 each = \$280

These cost estimates are based on current *small quantity* prices and should become even more economical at quantity prices.

The costs for the playback system are approximately the same, only with a VSI de-formatter card replacing the VSI formatter card.

The 24-Hour Challenge

We stated at the outset that 24-hour unattended operation is a major goal of a new HDR system. Here again the computer industry has risen to the challenge, with solutions from many vendors. As an example, a 14-tape changer from ATL Products is readily available for ~\$2000 each to support DLT-style tapes; other similar changers are available for other styles of tape cartridges. The changer and its attendant drive will fit into a 7" standard-rack space. We need four changers, so the system (dubbed 'HDR-1') cost becomes:

System 'HDR-1'-	
Basic system (from above)	\$40,000
14-tape changers (4)	\$8,000
<u>Total</u>	\$48,000

System capability: 512 Mbps for 24 hours

Tape cost: 56 tape cartridges at ~\$70 each = \$3920

Possible physical configuration: 3 small (~18" high) racks with changers and electronics

Two such systems operating in parallel will achieve 1 Gbps for 24 hours unattended operation for <\$100,000 (not including tape). And this is just the beginning, since the computer industry is investing hundreds of millions of dollars to help us.

What Does the Future Hold?

Fortunately, we can be fairly certain that, by building on widely accepted computer standards and subsystems, the future is very bright. As shown in Figures 1 and 2, by ~2007 single-tape capacity is expected to reach 1 TB at a data rate of 1 Gbps, *and the per drive cost is expected to remain essentially the same.* This would be hard to believe if we did not have the demonstrated example of the disc drive industry already meeting or exceeding projected price/performance growth rates.

The projected system of ~2007 (dubbed 'HDR-4') looks and costs about the same as the system described above:

System 'HDR-4'-	
Computer, memory, backplane	\$5,000
SCSI interface cards (2)	\$2,000
VSI Formatter card (4)	\$8,000
Tape drives (4)	\$28,000
Cabinets	\$1,000
14-tape changers (4)	\$8,000
Contingency	\$2,000
<u>Total</u>	\$54,000

System capability: 4 Gbps for ~30 hours

Tape cost: 56 tape cartridges at ~\$70 each = \$3920

Possible physical configuration: 3 small (~18" high) racks with changers and electronics

Depending on the capabilities of the interface cards and backplane, it may be possible to add additional interface cards and tape drives directly to this system to reach 8 Gbps; alternatively, two HDR-4 systems operating in parallel will reach 8 Gbps and cost only marginally more since the foundation computer and bus are inexpensive. In either case, the cost would be <~\$110K for a system capable of recording 8 Gbps for 24 hours unattended.

Table 1 illustrates several of the relevant system parameters for current and other proposed VLBI data systems, as well as projections for four generations of HDR systems (HDR-1 through HDR-4) through ~2007. What we see is an astonishing projection of costs declining from ~\$100K/Gbps in the near term to ~\$14K/Gbps around 2007 for systems which can record for at least 24 hours unattended. Even if these projections prove to be too optimistic by a factor of 2, it is clear that the cost of VLBI HDR systems will decline rapidly, *provided they are largely based on standard off-the-shelf technology.*

Other Possible Uses of the HDR System

Though we have focused primarily on VLBI users, there are clearly other disciplines that may benefit from HDR systems of the type proposed. Depending on the particular application, it is possible that the VSI interface is already sufficiently capable since systems based on the VSI standard simply transfer a set of time-tagged parallel bit streams from input to output. We encourage other potential users to suggest ways in which small design modifications or additions would make the HDR system more suitable to their needs.

Proposal for Proceeding

The path to the proposed HDR VLBI system is not believed to be particularly difficult. A commitment of an estimated 2-3 person-years of engineering and development time will be needed to develop the basic system. Following this initial investment, follow-on enhancements in capacities and data rates will become available mostly through the largesse of the computer industry, with only minimal additional engineering.

Haystack Observatory has been active in VLBI data-systems development arena for more than 30 years. The products of these efforts have been the Mark IA, Mark III, Mark IIIA, Mark IV and VLBA data recording and playback systems, many of which are in active use worldwide. During these years of development, Haystack has maintained close interactions with the development laboratories of several major industry players in magnetic data recording (Seagate, HP, IBM, Quantum). We propose to maintain and strengthen these industry ties to ensure, as much as possible, that future commercial recording products provide easy compatibility with VLBI HDR requirements.

References

- SCSI: http://dmoz.org/Computers/Hardware/SCSI/SCSI_Bus/
<http://www.scsita.org/terms/scsiterms.html>
<http://www.kudra.com/rs/wisdom/handbook/handbook167.html>
- PCI: <http://www.pcisig.com/specs.html>
- PCI-X: <http://www.compaq.com/products/servers/technology/pci-x-enablement.html>
- Compact PCI: <http://www.picmg.com/gcompactpci.htm>
- DLT: <http://www.dltp.com/>
- LTO: <http://www.lto-technology.com/>
http://www.seagate.com/support/tape/lto/lto_tech_paper.shtml
- AIT: <http://www.aittape.com/ait.htm>
- Exabyte: <http://www.exabyte.com/index.html>
- Tape changer: <http://www.atlp.com/products/1500/1500.html>

	Mark IV	Mark IV Thin-Film	S2	S3	K4 (Sony)	D6 (Toshiba)	HDR-1 (buildable Y2000)	HDR-2 (projected)	HDR-3 (projected)	HDR-4 (projected)
When available?	now	?	now	2002?	now	now	(2000)	(2003)	(2005)	(2007)
Technology	linear	linear	helical	helical	helical	helical	linear or helical	?	?	?
Max data rate (Gbps)	1-2	2-4	1/8	1	1/4	1	1/2	1	2	4
#tape transports	1	1	8	8	1	1	4	4	4	4
Single tape capacity (TB)	0.5	1.0-2.0	.04	.14	0.1	.46	0.1	0.25	0.5	1.0
Est. system cost (\$/Gbps): w/o changer w/changer	\$250-\$140K -	\$130-70K -	\$235K -	\$100K ? ?	\$300K ?	\$500K ?	- \$100K	- \$50K	- \$30K	- \$14K
Unattended operation (hrs at max rate): w/o changer w/changer(s)	1.1-0.5 -	~1.1 -	5.7 -	2.5 24	1.0 24?	1.0 >24	1.7 24	2.2 30	2.2 30	2.2 30
Tapes/24 hrs (at max rate)	~22-44	~22	~32	~80	24	24	56	44	44	44
TB/24 hrs (at max rate)	11-22	22-44	1.4	11	2.8	11	5.5	11	22	44
Media cost (\$/TB)	\$1950	\$975-490	\$360	\$312	\$2000	\$1090	\$710	\$280	\$140	\$70
Shipping weight (kg/TB)	21	10.5-5.2	9.3	2.7	14.4	3.6?	3.8	1.5	.7	.35
Heads: Est lifetime (hrs) Replacement cost \$/hr head cost	3-5K \$20-40K \$5-8	>10K? \$2.5K <\$0.25	3-5K?	3-5K? ~\$8K ~\$2	? ? ?	1K \$30K \$30	>10K ? <\$1	>10K ? <\$1	>10K ? <\$1	>10K ? <\$1
Speedup on playback for data-rates < max	Y	Y	N	N	Y	Y	Y	Y	Y	Y

Table 1: Characteristics of various current and proposed VLBI recording systems (best estimates)

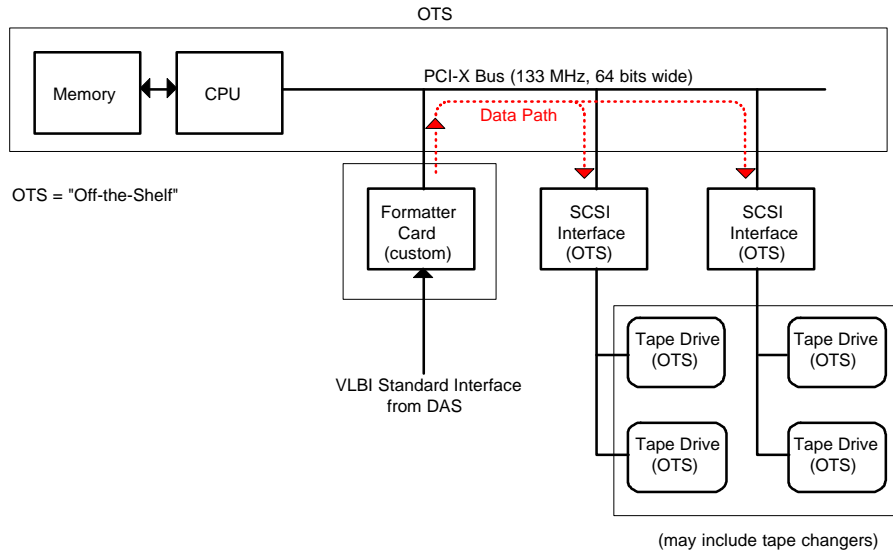


Figure 3: Example Block Diagram of HDR VLBI Data-Recording System

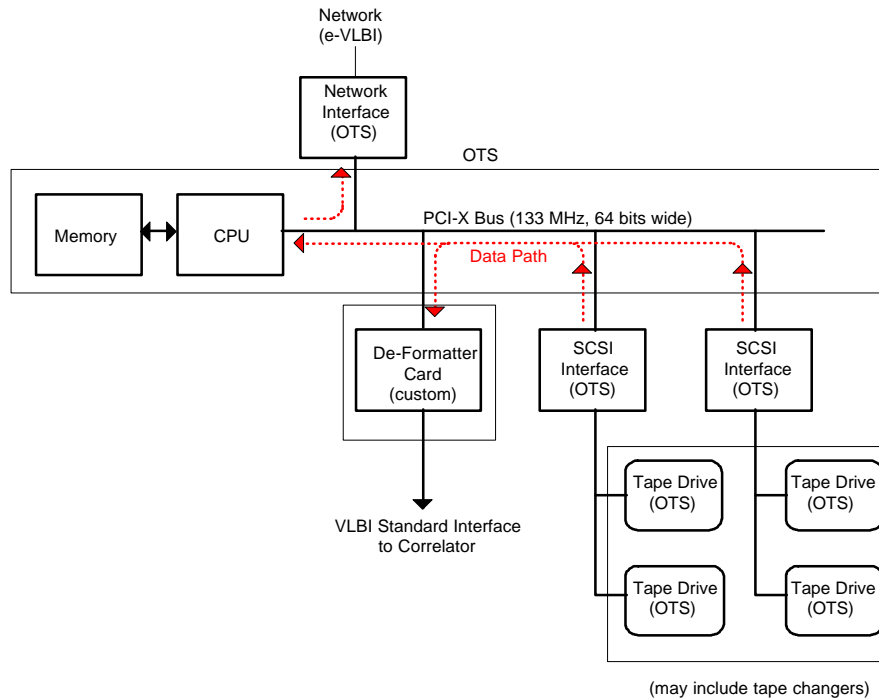


Figure 4: Example Block Diagram of HDR VLBI Playback System

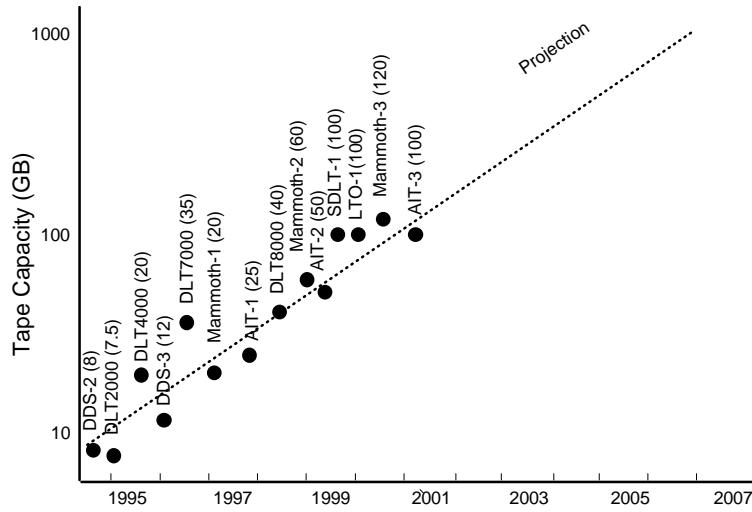


Figure 1: Tape Capacity vs. Time

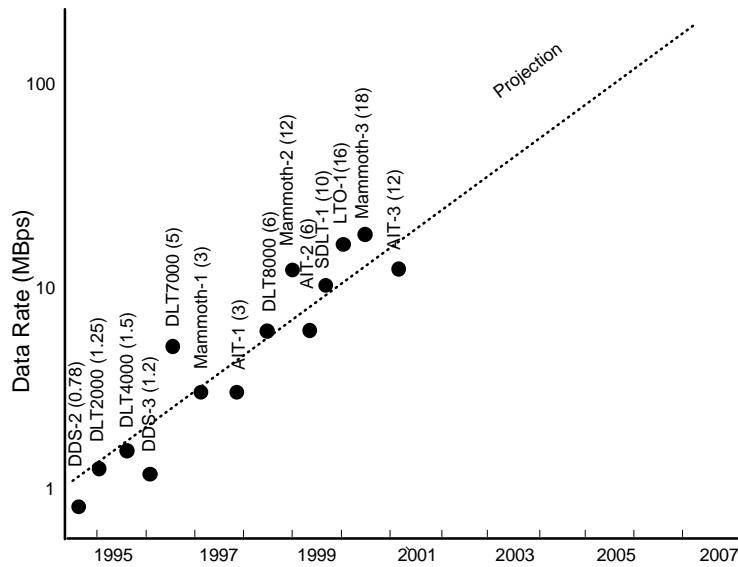


Figure 2: Tape Data-Rate vs. Time

Note:
 AIT: Sony, helical scan
 DDS: Sony, helical scan
 DLT: Quantum, linear
 LTO: HP, IBM, Seagate, linear
 Mammoth: Exabyte, helical scan