Areal Density Growth:

Is it the manifest destiny of the hard-disk drive?

Giora J. Tarnopolsky

TarnoTek - www.tarnotek.com
P.O.Box 519, Palo Alto, CA 94302-0519
Phone: (650) 823-6852 Fax: (650) 322-3160
gjtarno@tarnotek.com

Presented at the THIC Meeting at the Sony Auditorium,
3300 Zanker Rd, San Jose CA 95134-1940
March 4-5, 2003

Single-Slide Summary

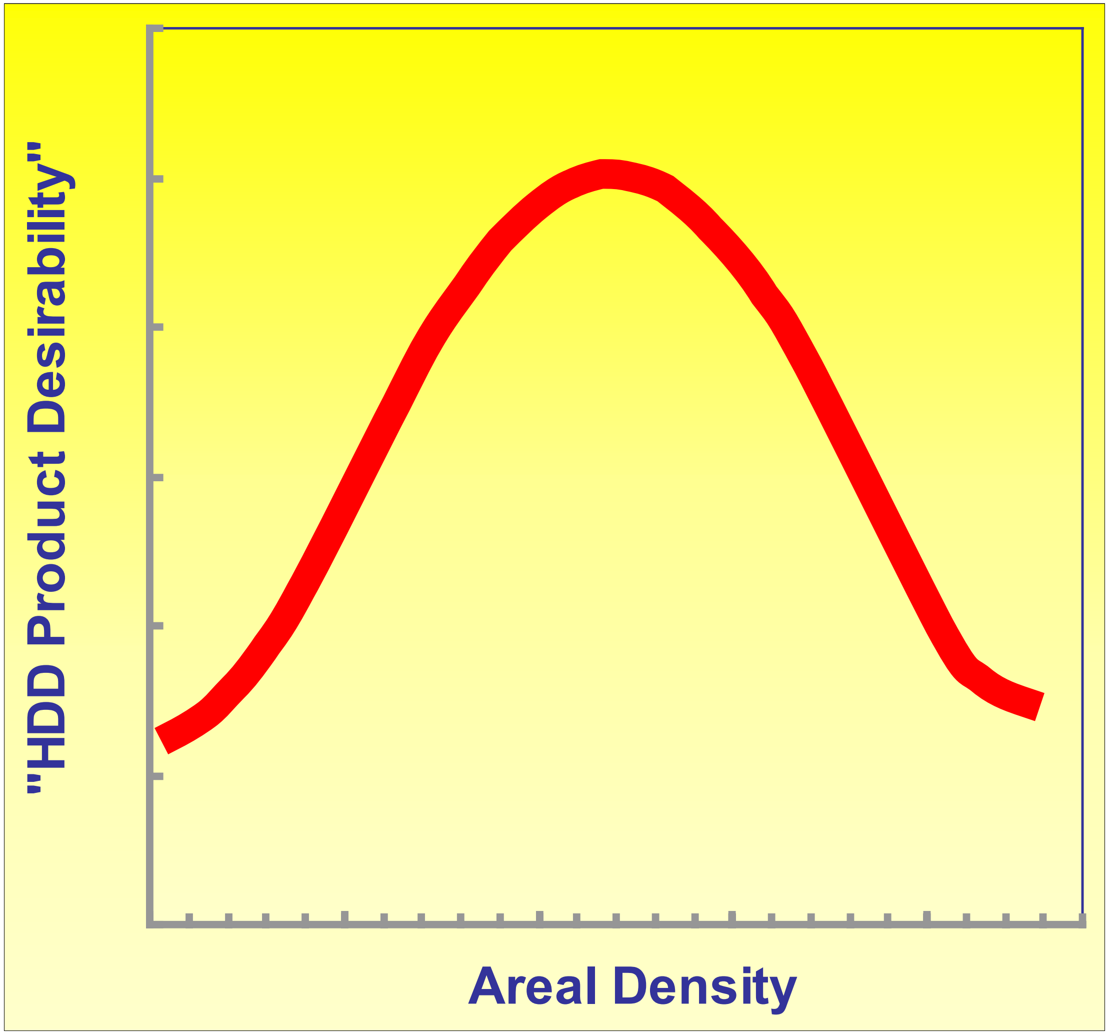

- There is an optimum Areal Density for which the HDD Product Desirability is highest.
- At the optimum AD, the capabilities of the HDD as a high-capacity, high-data-rate device that assures data perman-ency, ruggedness, and low manufacturing and development costs are fully manifested

Two-Slide Talk: Conclusions

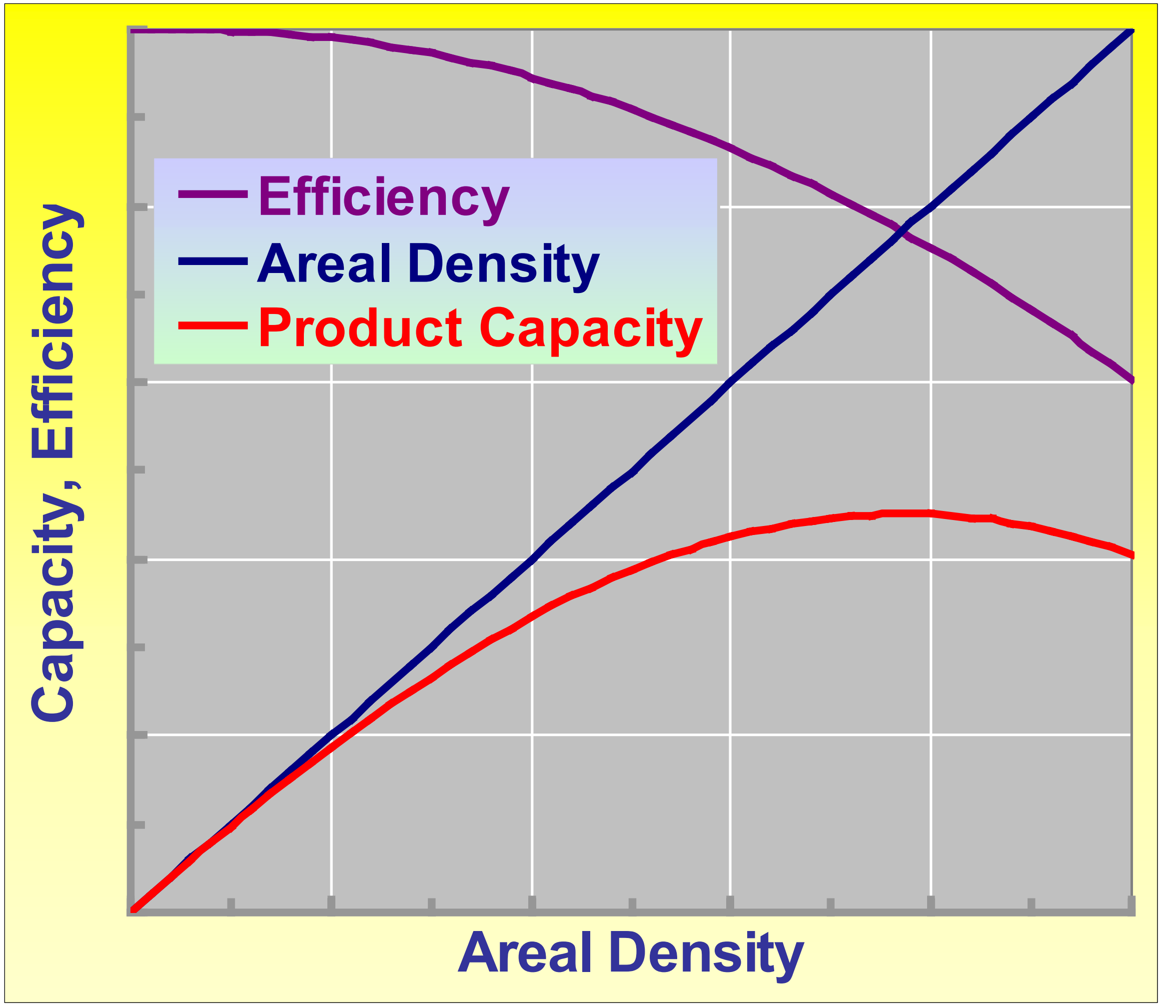

- As the **data recording density** increases, **storage efficiency** is compromised.
- ECC, servo overhead, mechanism
- The users' **capacity** will not grow with the "native" areal density of 1/(bit dimensions)
- Diminishing performance returns

Outline

- Introduction - *at a brisk pace*
  - Evolution of useful products
  - Tradeoff: feasibility vs. desirability
- Developments 1990 - 2002
- 1 Terabit/in^2 proposals
- Capacity scaling
- The value add of A.D.

Introduction

- Areal density has grown from 1 Gb/in^2 in 1990 to ~150 Gb/in^2 by November 2002
- Feasibility assessments of 1 Tb/in^2 have been published, stated INSIC Goal
- 10, 50 Tb/in^2 have been suggested for HAMR

HDD future predicated on continuing growth of the areal density

Issues

- Data storage at 1 Tb/in^2 is possible
- Materials and systems research for 1 Tb/in^2 is scientifically sound

These are not in question!

➔ Is it desirable to accomplish $\geq$ 1 Tb/in^2 products?

➔ What constraints on user features does 1 Tb/in^2 impose?

➔ What is the return on investment?

265 Tb/in^2 virtually demonstrated … and

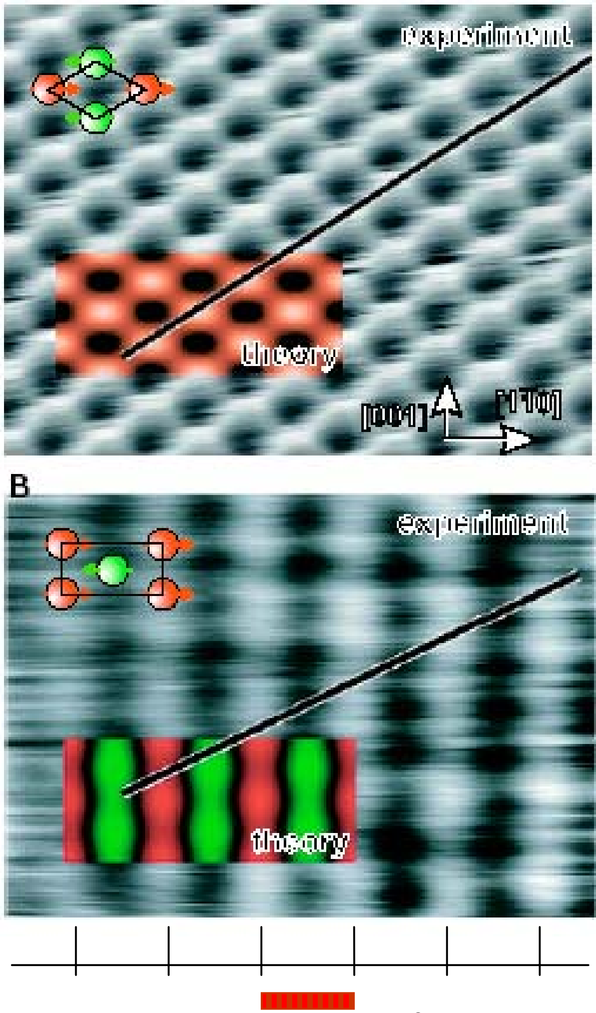

4.5 ± 0.1 Å

non-magnetic W tip

Full image size is 2.7 x 2.2 nm^2

magnetic Fe tip

- Single atomic layer of Mn on W(110) forms a two-dimensional antiferromagnet.
- These are magnetic monolayers of chemically equivalent atoms, where adjacent atoms at nearest-neighbor sites have opposite magnetic moments.
- The ***single-atom spin has been resolved*** by Spin-Polarized Scanning Tunneling Spectroscopy
- AD ≥ 265 Tb/in^2 for a 12-atom bit

- **S. Heinze, M. Bode, A. Kubetzka, O. Pietzsch, X. Nie, S. Blügel, R. Wiesendanger, Science 288 (2000) 1805-1808: Real-Space Imaging of Two-Dimensional Anti-ferromagnetism on the Atomic Scale**
- **R. Wiesendanger & M. Bode, Solid State Communications 119(2001) 341-355: Nano- and atomic-scale magnetism studied by spin-polarized scanning tunneling microscopy and spectroscopy**

… the drive is … “just engineering”

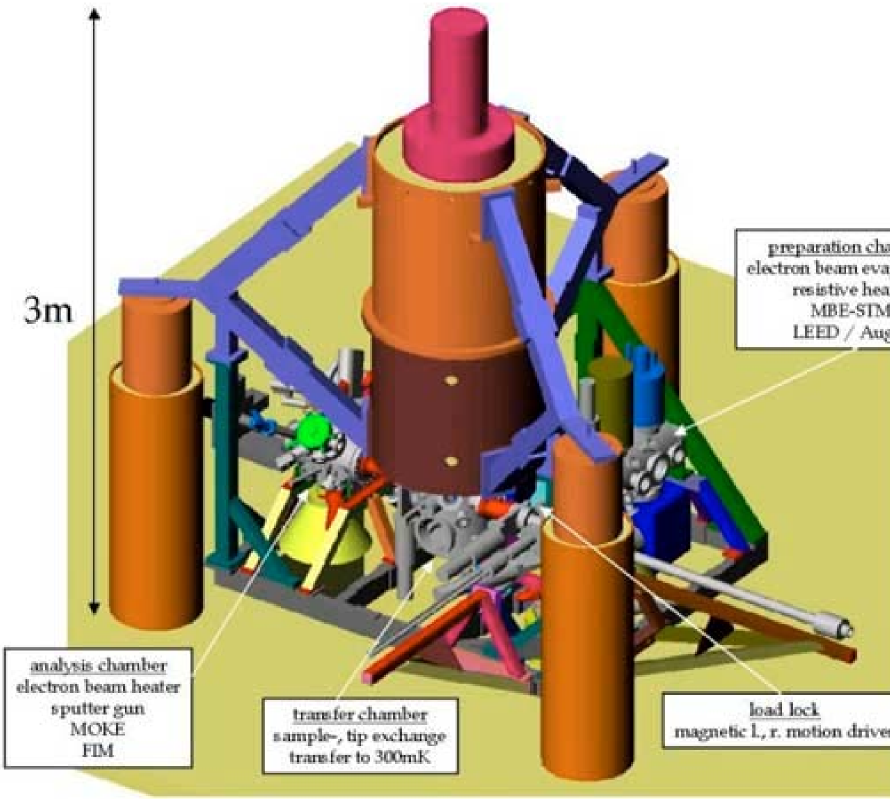

Ready for Fry’s?

Nearly Perfect Inventions

- Some inventions are born perfect
  - This assures their permanency …
  - … and defines the domain of development
- Examples of perfect inventions are the bicycle, the umbrella, the book, and the disk drive

Nearly Perfect Inventions: Bicycle

- Gyroscopic effect assures stability of the rider
  - Under torque T, the bike turns but does not fall
- Low ratio of vehicle mass to rider mass
  - ~ 15 % (as compared to ~2,200% for car)
- Efficient
- Rugged
- Mass-produced
- Affordable

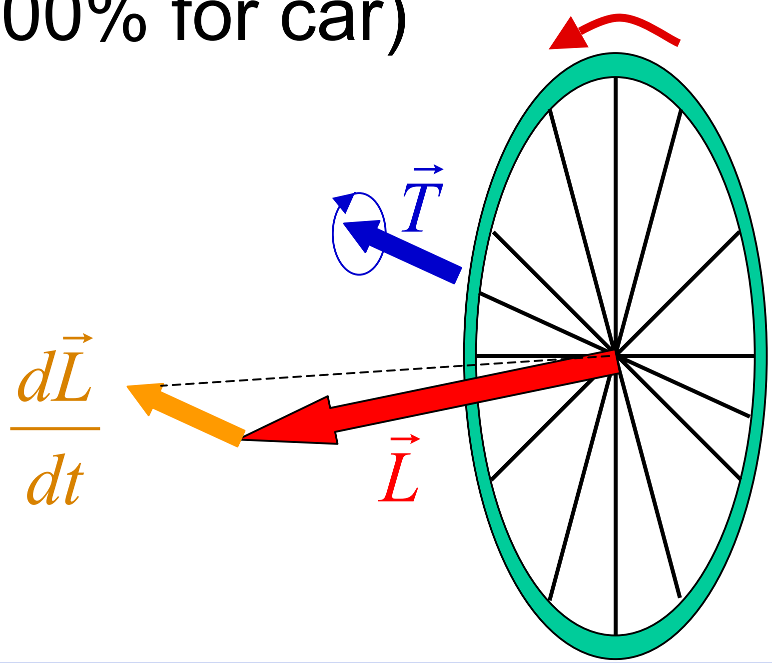

Disk Drives - Nearly Perfect Invention

- 2-D travel with only one linear motion
- High volumetric density
- Random access
- Mass-produced
- Non-volatile
- Affordable
- Rugged

→ Few-hundred $/box

→ No vibration isolation, no T stabilization

→ These properties define drives

Areal Density and Storage Costs

Areal Density Benefits

- Fewer components
  - Heads
  - Disks
  - Spindles
- Volumetric density

No AD Benefit

- Data storage management (s/w & h/w)
- Data permanency
- Ruggedness
- Electronics
- Interfaces
- Capital & Test Equip.

Cost vs. Capacity, 7200 rpm

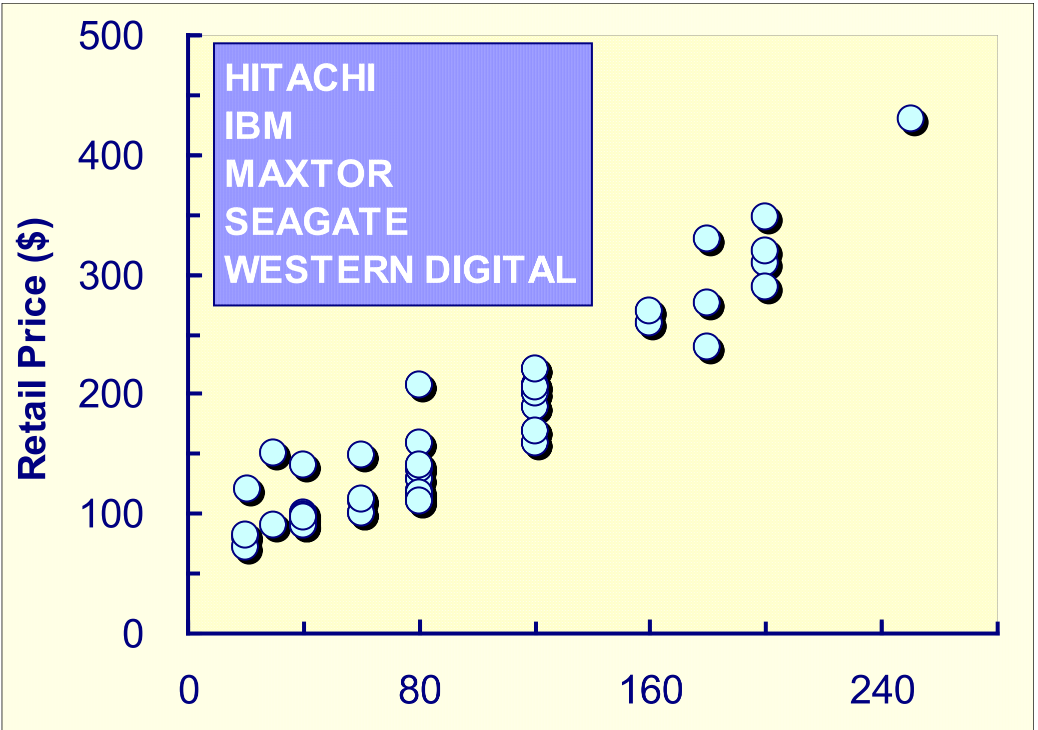

Source CDW, 2/2003

- Chart shows current ***retail*** unit prices (not best price)
- Thus, a 1 Tbyte drive should cost about $2,000. ***Right?***
- ***Not quite!***
- **No. 1: The "box" would not sell above**

$$Price = \$ \, (n \times 100),$$

where $n < 5$

"The price the customer is willing to pay determines allowable costs." P. Drucker

No. 2: For larger capacities, the cost of managing the data exceeds the H/W cost

1 Tbyte Storage Example - early 2003

Fantom Drive 960 GB G-Force RAID LX

- 8 x 120 GB EIDE drives, 7200 rpm
- Host: Any SCSI host computer, 160 MB/s
- Special vibration or temperature handling: *None*
- Hot-swappable everything (RAID 1,3,5)
- MTBF: 500,000 hs

960GB, 8-drive, 8-bay,
Ultra3 SCSI (Ultra160/m),
EIDE backplane,
rackmount RAID
(RAID levels 0,1,0+1,3 and 5)
hot swappable drives

- Retail on 2/24/03: $7,609.44
- EIDE drives' @ Fry's: $960.00
- Non-AD cost: **87%**

O/S independent, data management, redundancy, profit ...

From 1 to 150 Gb/in^2: 1990 - 2002

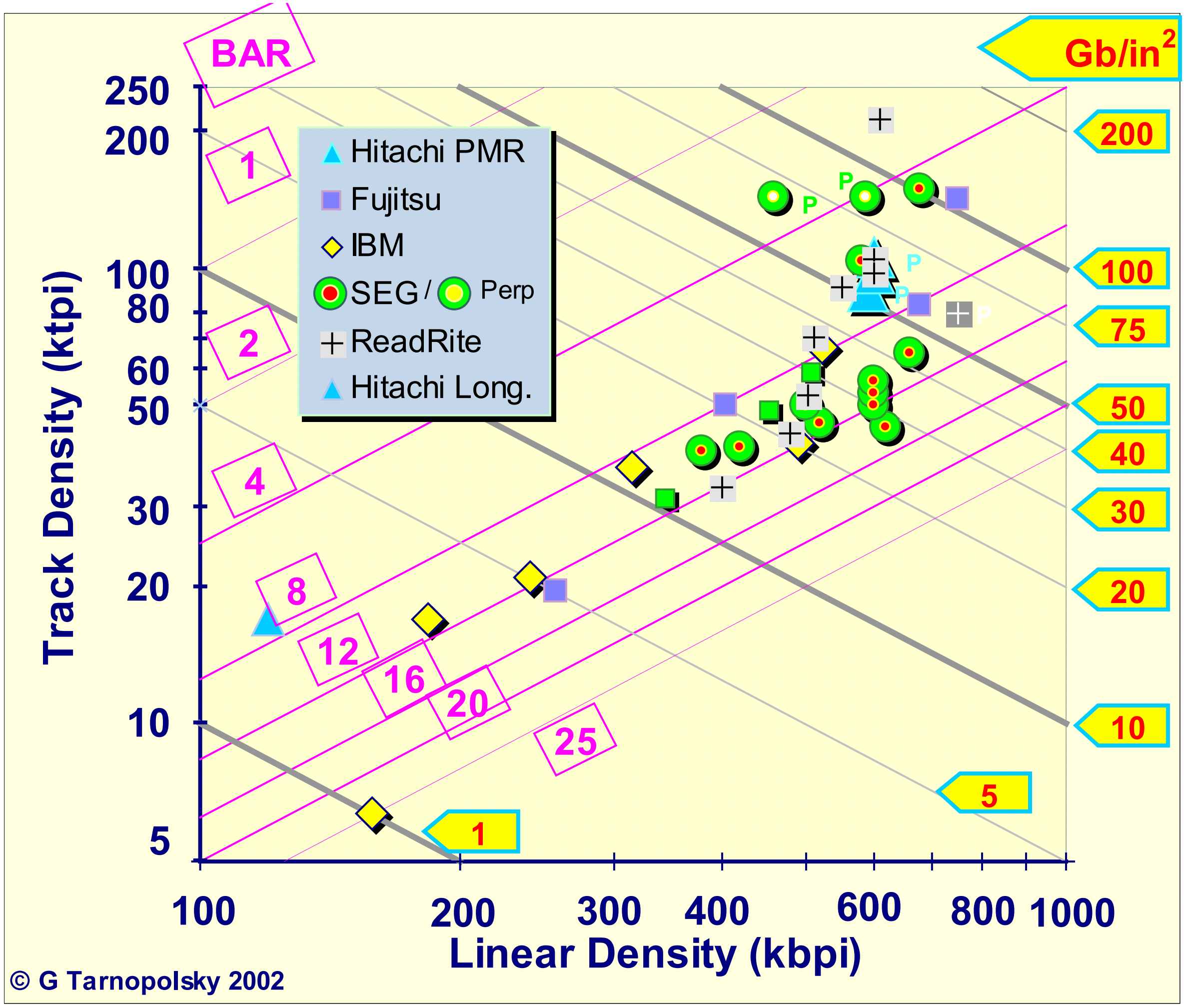

- Extraordinary achievement
- Log- log plot of track density and linear density
- The hyperbolae are lines of constant areal density
- The rays from the origin are lines of constant bit aspect ratio
- Low AD demos had better BER than recent demos
- From about 1 Gb/in^2, the BAR has changed from ~ 25 to ~ 3 to 4

From 0.1 to 1.0 Tb/in^2

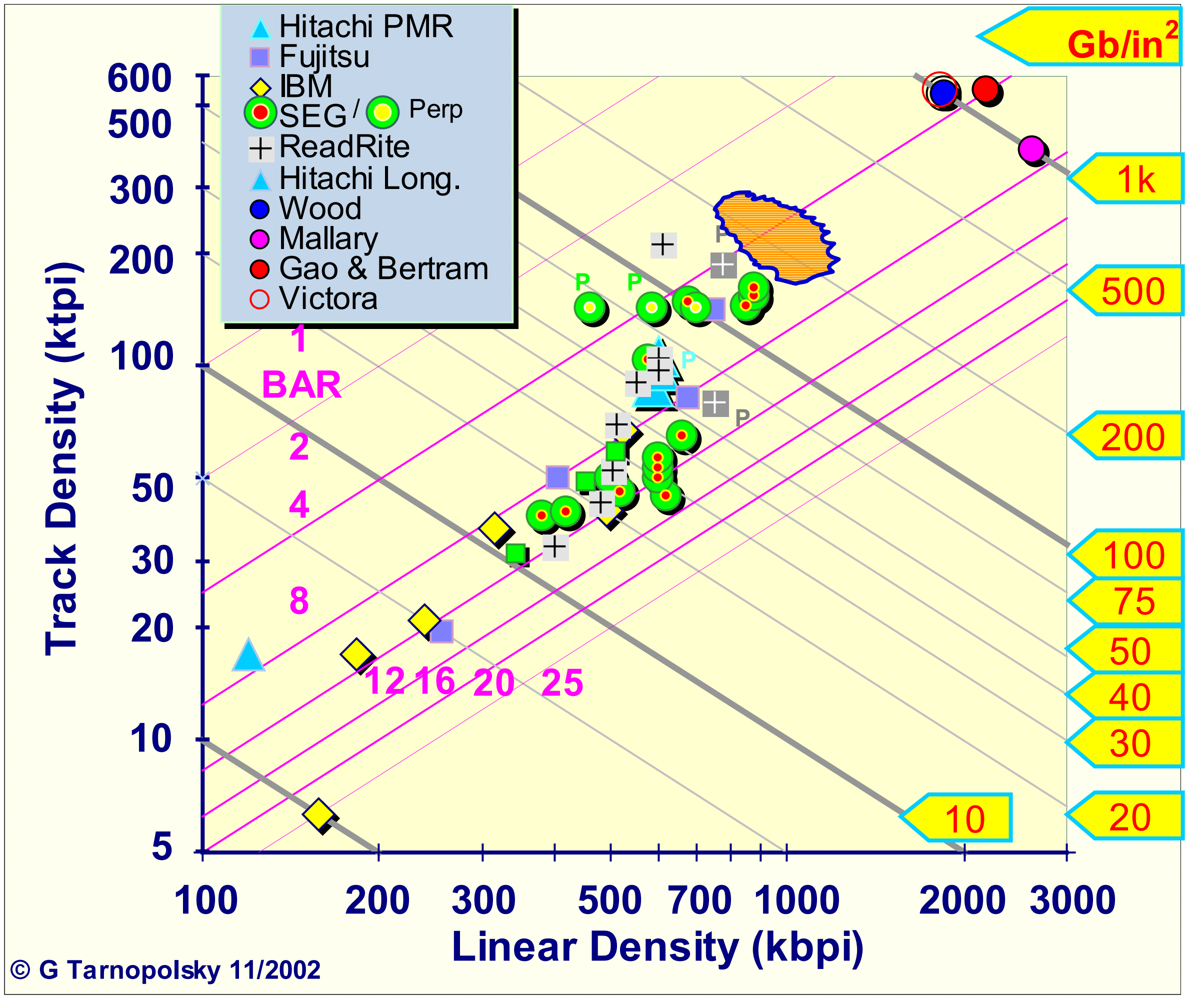

- **At the 2000 MMM Conference, PMR results between 100 to 145 Gb/in^2, and longitudinal results at ~ 150 Gb/in^2 were shown**
- **Four 1 Tb/in^2 studies**
- **R. Wood et al., IBM**
- **M. Mallary et al., Maxtor**
- **Gao & Bertram, UCSD-CMRR**
- **R. Victora, UM - MINT**
- **These studies explore the regime of ~ 2 Mbpi & ~ 500 ktpi**

Early Demos, 1 Tb/in^2 Studies

C. Tsang et al., *IEEE Trans. Mag.*, vol. 26, p. 1689, Sept. 1990. T. D. Howell et al., *ibid.*, p. 2298.

J. Hong et al., *IEEE Trans. Mag.*, vol. 38, p. 15, 2002. (Fujitsu demo)

Z. Zhang et al., *IEEE Trans. Mag*. vol. 38, p. 1861, Sept. 2002. (Seagate demo)

F. Liu et at, presented at InterMag 2002, Amsterdam, The Netherlands, April 2002 (ReadRite-MMC demo.)

IBM Travelstar 80GN, Model No. IC25N080ATMR04, Nov. 2002.

47th MMM Conference 2002, announcements by Seagate & ReadRite.

R. Wood et al., *IEEE Trans. Mag.*, vol. 38, p. 1711, 2002.

M. Mallary et al, *IEEE Trans. Mag.*, vol. 38, p. 1719, 2002.

K. Gao and N. Bertram, to be published *IEEE Trans. Mag.*, TMRC 2002, Santa Clara, CA, USA, August 2002.

R. Victora et al., *IEEE Trans. Mag.*, vol.38, p. 1886, Sept. 2002.

M. Kryder et al., presented at TMRC 2002, Santa Clara, CA, USA, August 2002.

How the Areal Density Was Won

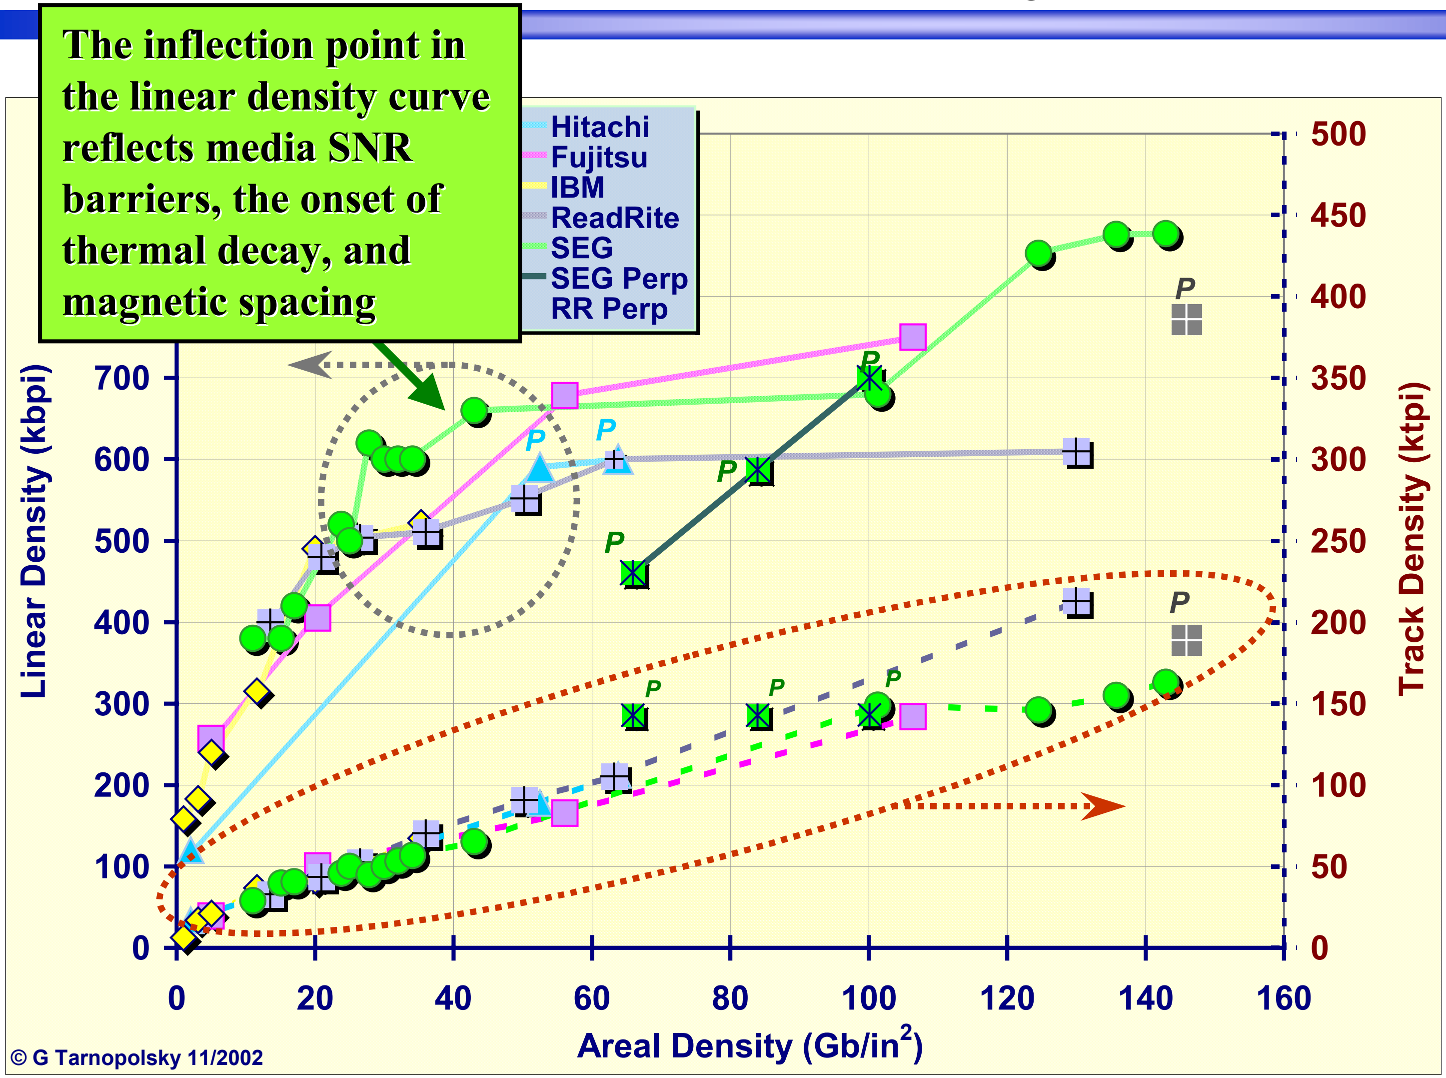

- Linear and track density vs. areal density
- Track density increased by ~ 35, caused most of the areal density gain.
- Enabled by the advent of spin valve and GMR heads, advances in head fabric-ation techniques, media SNR

A Moving BER Target

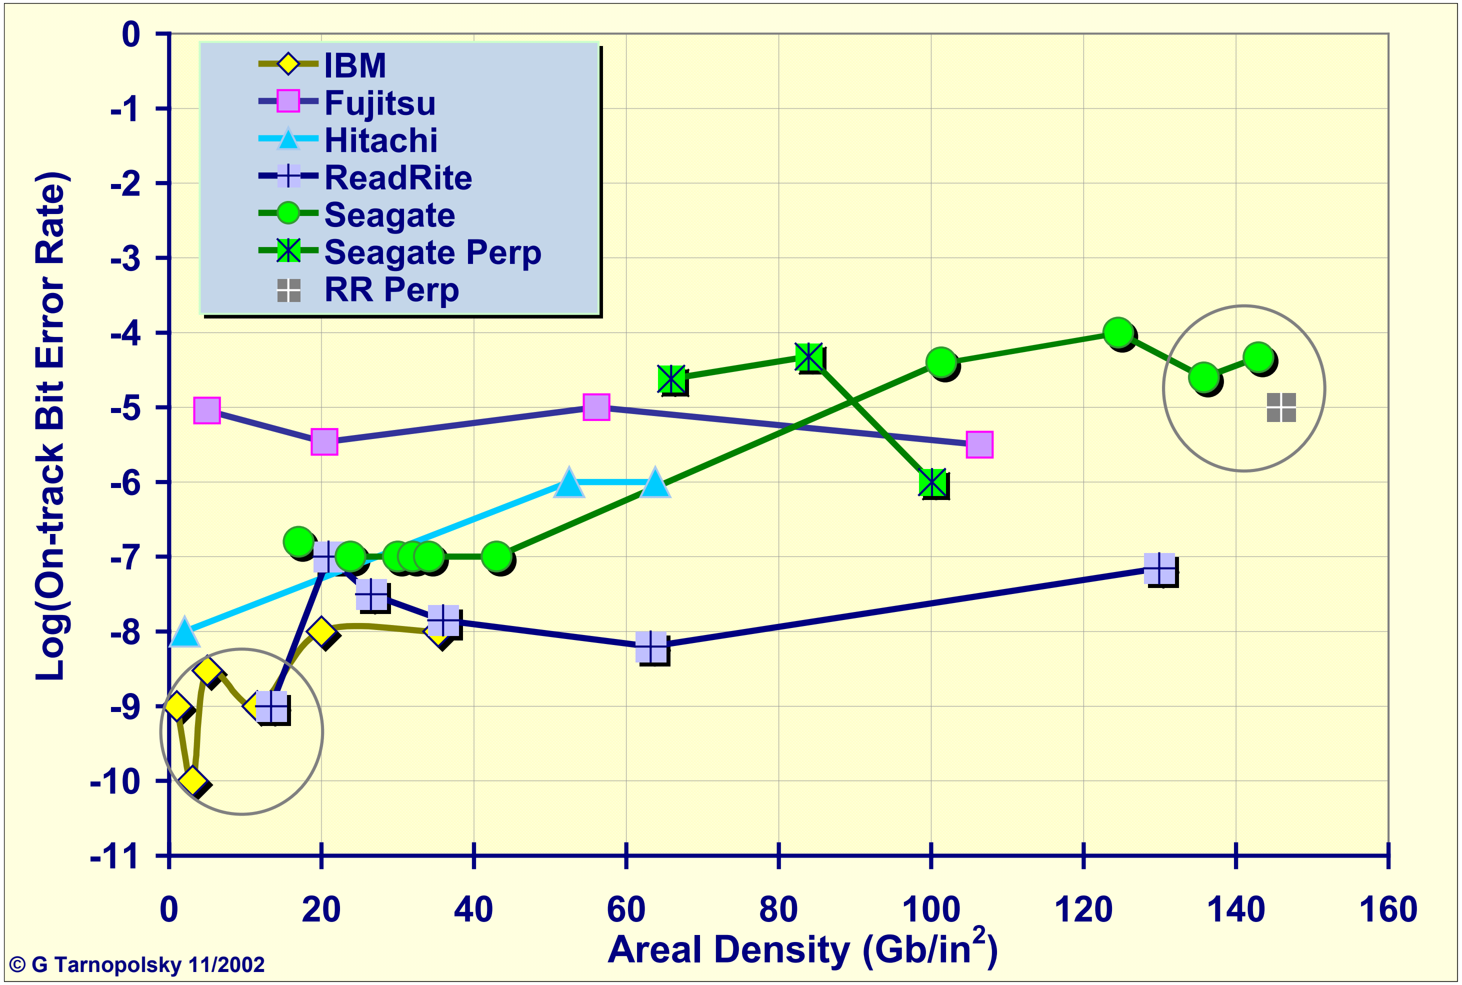

- Areal density demonstrations are rigorous, comprehensive assessments of the technology
- The BER has worsened with increasing areal density of demos
- The 100 ~ 200 Gb/in^2 regime is an important research field

Head-Media Separation: Downtrack Story

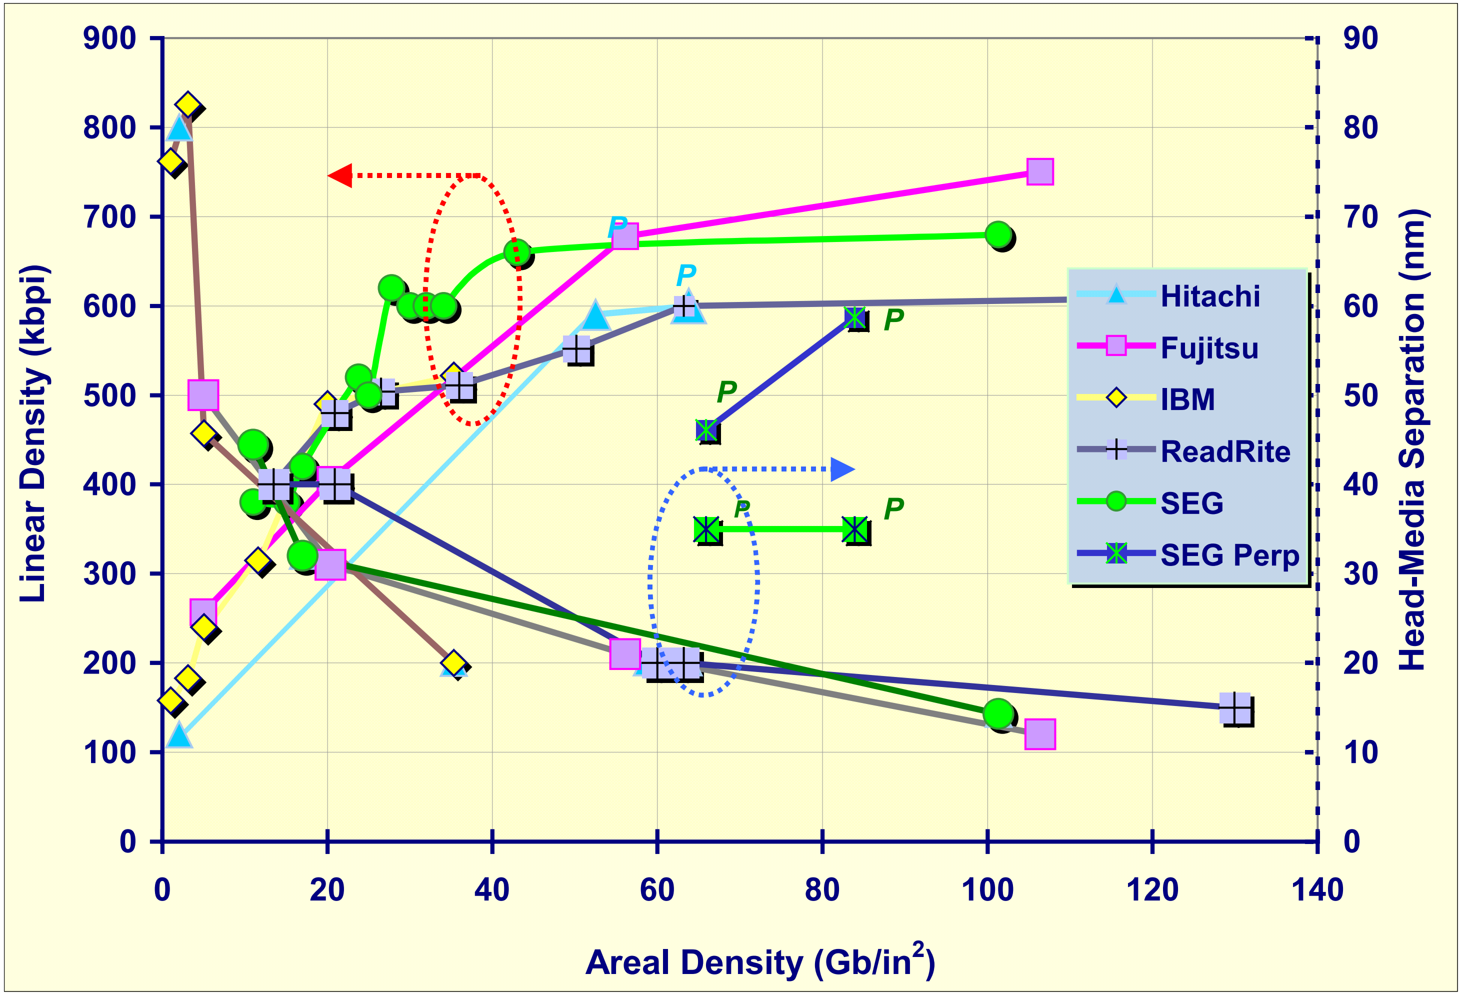

- As the recording wavelength decreased, the HMS shrunk from ~ 80 nm to ~ 10 nm
- Extraordinary efforts put in reduction of pole-tip recession, carbon overcoat thickness, flying height and lube

Flying Height *et al.*: Finite Resources

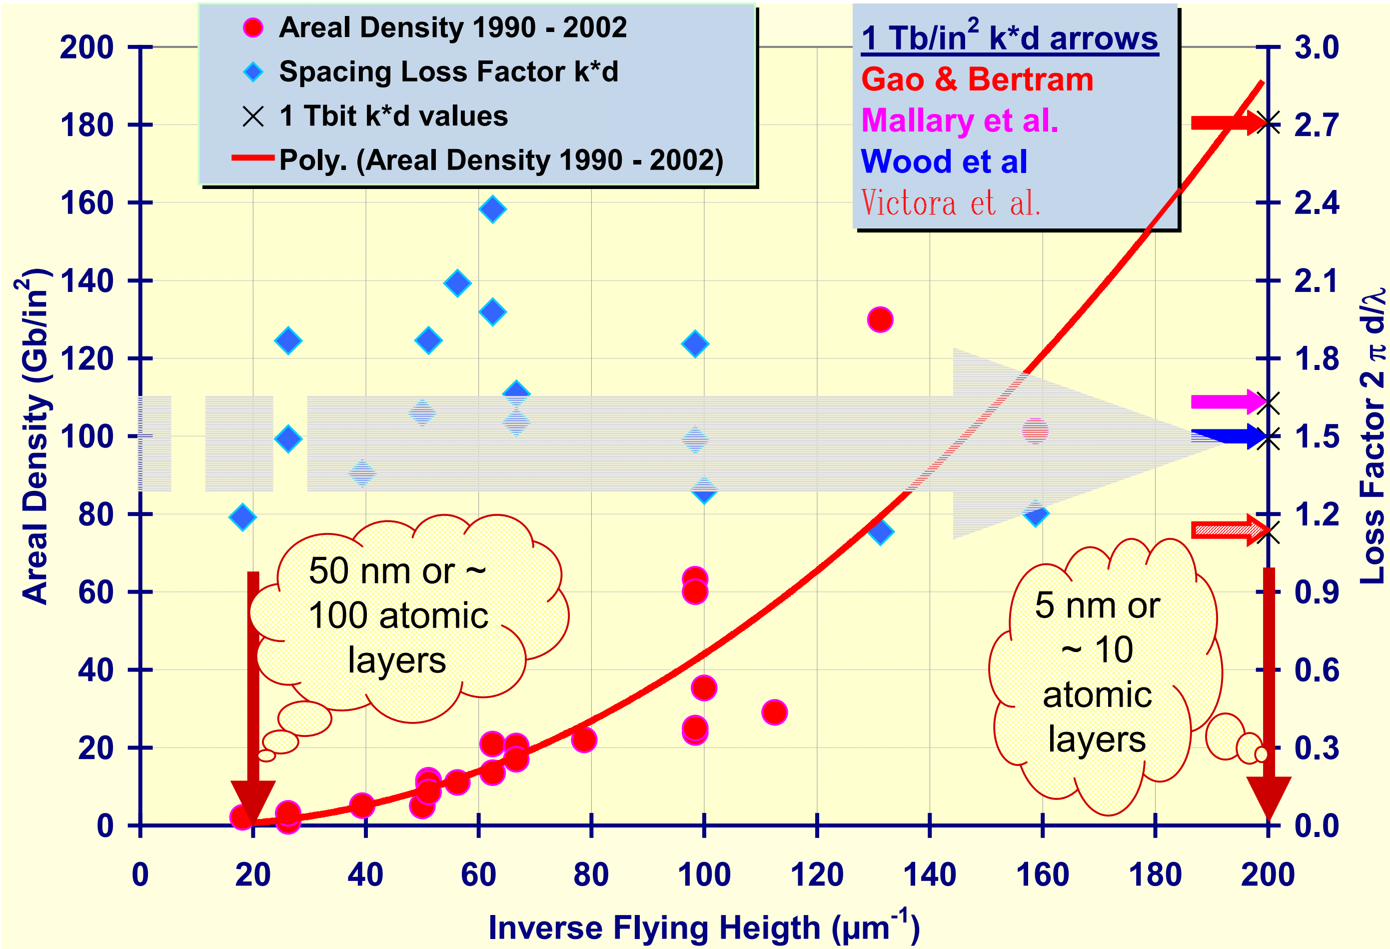

- Abscissae are inverse flying height, an open-ended scale
- 100 Gb/in^2 demo's FH = 60 Å
- Flying height is but one component of the head-media separation
- $e^{-(2\pi \bullet HMS/\lambda)}$ spacing loss factor. *kd* plotted

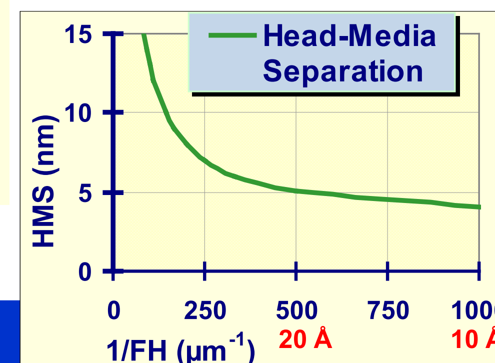

1 Tb/in^2 TMR: A Different Future

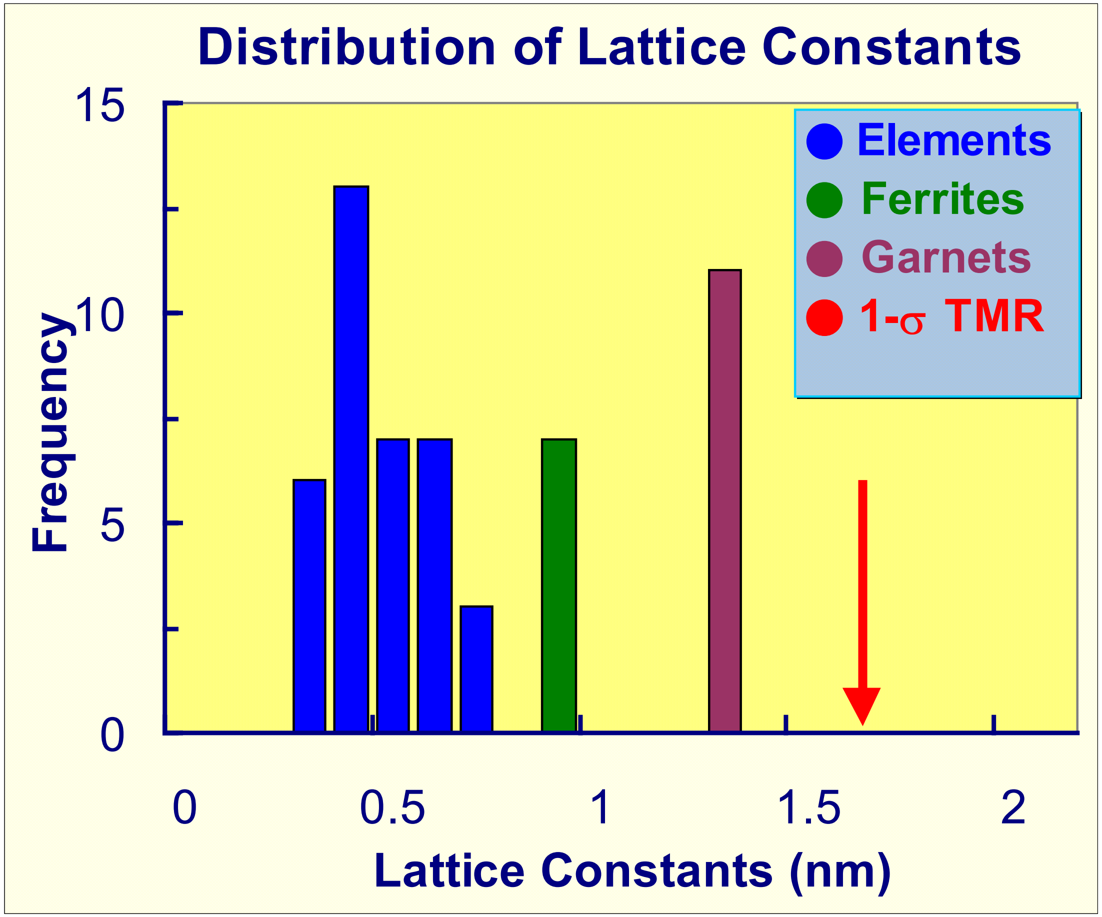

- In the 1 Tb/in^2 regime, 500 ktpi, the 1σ TMR (track misregistration) is about 1.5 nm
- The permissible TMR is of the order of magnitude of lattice constants
- This is not the case at 1, 10, or 100 Gb/in^2

Capacity will not grow linearly with AD

- The capacity of a product will not grow linearly with areal density
  - The limited SNR of high AD demands higher ECC overhead
  - The limited PES (position error signal) SNR requires higher servo overhead
  - TMR of $O($nanometer$)$ requires smaller arm lengths and platter diameters
- The effective user areal density becomes much lower that that given by the bits' dimensions

Scaling

- Capacity = areal density x area x ECC efficiency x servo efficiency and also mechanical limits

$$C = LD \times TD \times 2 \times \pi \left(R^2 - r^2\right) \times (1 - \psi_{ECC}) \times (1 - \psi_{servo})$$

$$C = f\,(LD, TD, R; mechanism)$$

Capacity is a parametric function of the mechanism.

Parametric dependencies: Mechanism

- Constraints: immunity from excessive vibration that causes TMR, and access time
- Number of requests N per unit time from the host to the drive

$$N \sim \frac{capacity}{<user\ data\ file>} \sim \frac{1}{access\ time} = \frac{1}{\tau}$$

- Data throughput leads to access times inversely related to the drive capacity

Otherwise, capacity wasted

Parametric dependencies: Mechanism

- Access time is a function of the applied torque to the actuator arm. Large torque produces large bending moment in the arm and excites deflections from nominal position
- Impose a desirable access time, determine the necessary torque, estimate the magnitude of the amplitude of residual vibration after actuation
- Amplitude of vibration is measured against track pitch

Deflection of flexure

- Torque for actuator sweep of arc $l\alpha$ cm in τ seconds

R:	disk radius
l:	arm length
h:	width at pivot
a:	acceleration at head
ω:	angular velocity
α:	angle
s:	$= l\alpha$ = arc length
m	mass
I_m:	moment of inertia w.r.t. pivot

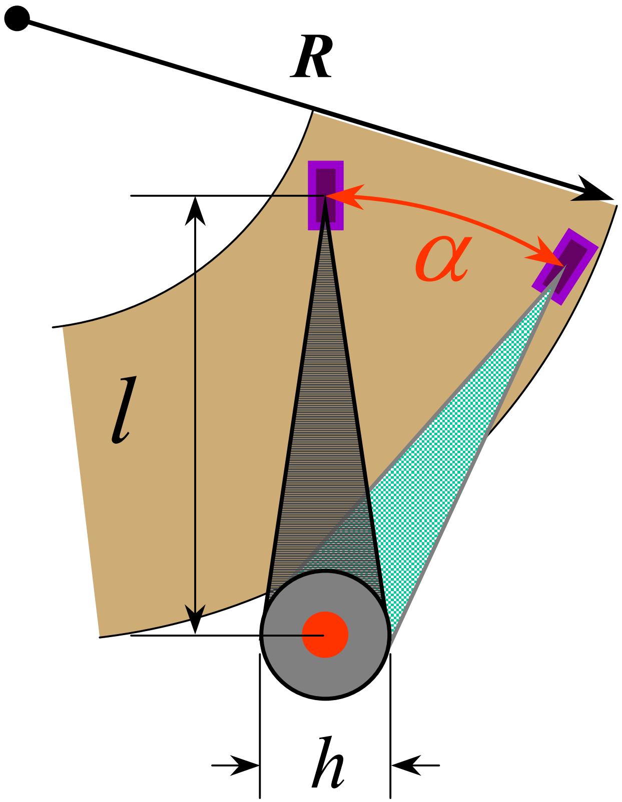

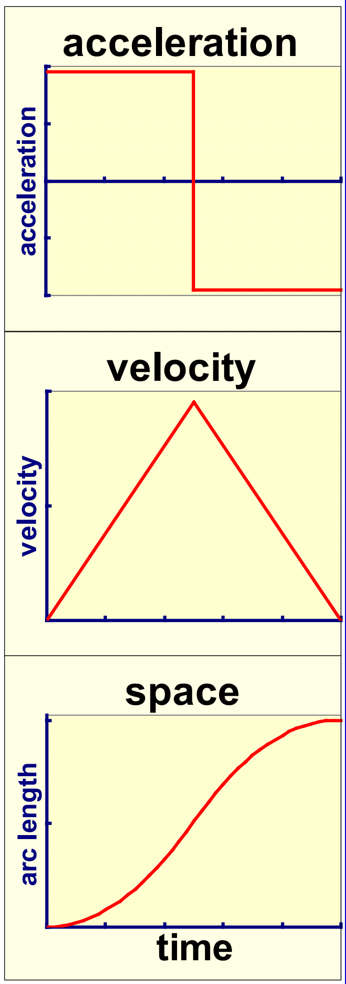

Torque

$$T_{act} = \frac{dL}{dt} = I_m \frac{d\omega}{dt} \quad ;$$

$$I_m = \frac{m}{6}\left(l^2 + \frac{h^2}{4}\right) \approx \frac{ml^2}{6};$$

$$\frac{d\omega}{dt} = \frac{a}{l} = \frac{4\alpha}{\tau^2} \, .$$

$$T_{act} = \frac{2}{3}\frac{ml^2\alpha}{\tau^2} = \frac{\rho A l^3 \alpha}{3\tau^2}$$

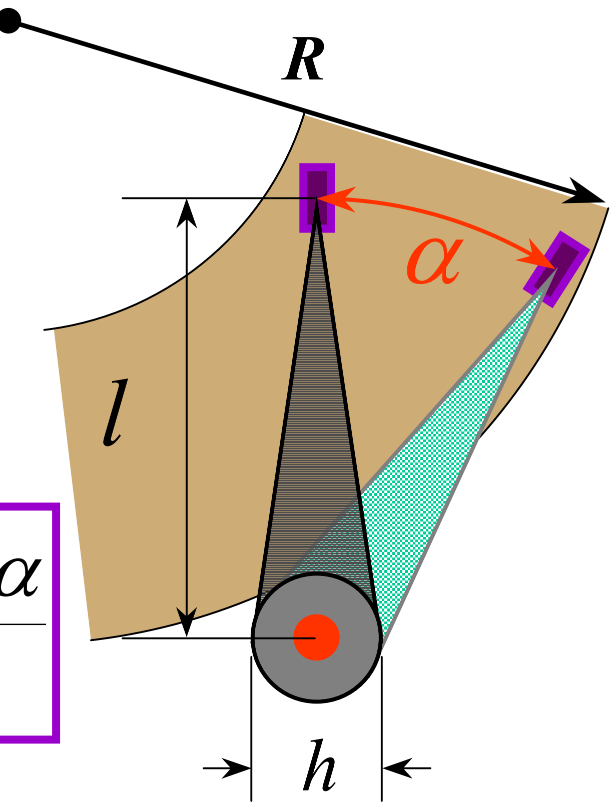

Deflection at the head position

$$\Delta s_{act} = \frac{1}{2}\frac{T_{act}\, l^2}{EI_o} = \frac{1}{6}\frac{\rho\, A\, l^5\, \alpha}{\tau^2 EI_o} = n \times TP$$

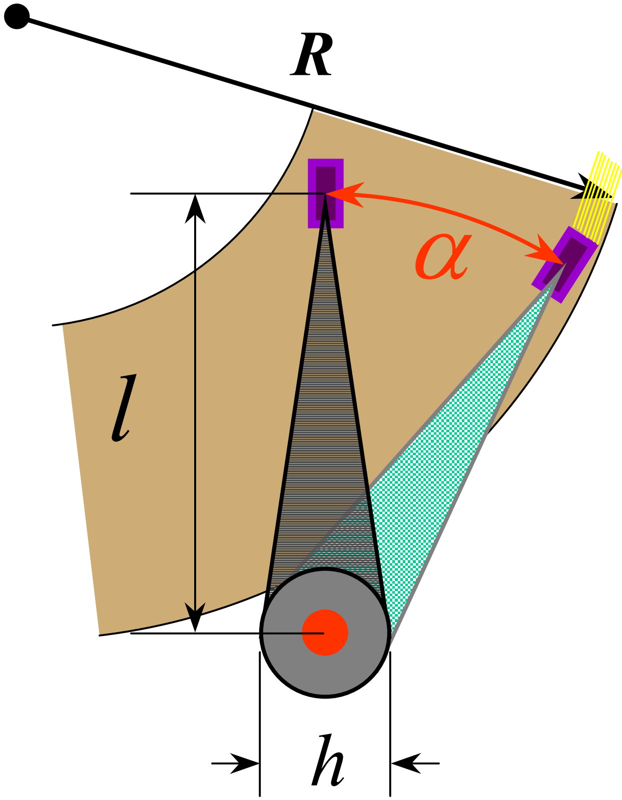

E:	Young's modulus
I_o:	moment of inertia (cross- section)
A:	cross section area
ρ:	density

The flexure deflection by the bending moment of actuation should be compared against the track pitch.

Capacity vs. Areal Density (I)

$$l^2 = \left(\frac{6n\tau^2 EI_o}{\rho A \alpha} \right)^{\frac{2}{5}} \cdot \left(\frac{1}{TD} \right)^{\frac{2}{5}}$$

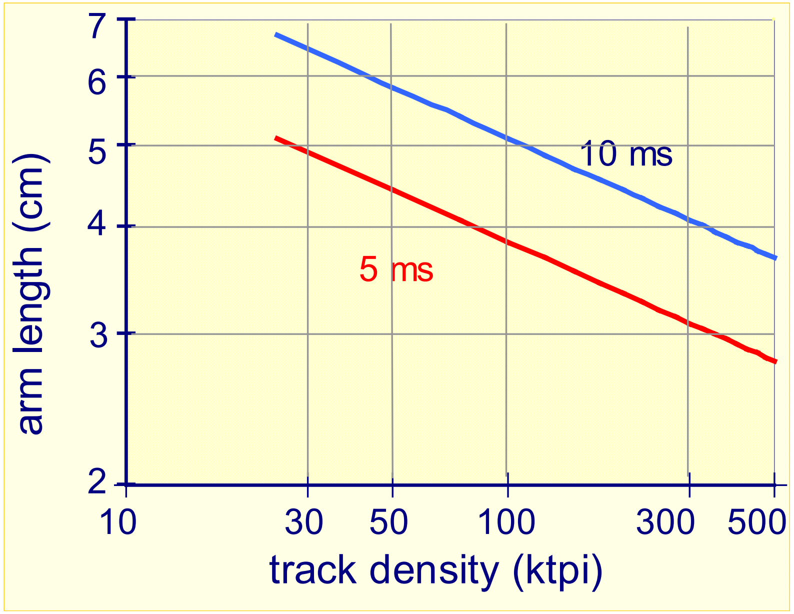

- Disk radius R ~ arm length
- With a constraint on the amplitude of vibration of the arm, the capacity grows slower than linearly with TD:

$$C_{act} = 2\pi \left(\frac{6n\tau^2 EI_o}{\rho A \alpha} \right)^{\frac{2}{5}} \cdot LD \cdot TD^{\frac{3}{5}}$$

per platter

Non-repeatable Perturbations

- A different scaling attains if the drive's environment creates a random, non-repeatable torque of rms value $\breve{T}_{nr}$. For instance, in a RAID
- T_{nr} has a probability distribution, and a 3-σ TMR event would happen for a 3-σ T_{nr} occurrence
- The rms arm deflection is

$$\Delta s_{nr} = \frac{1}{2} \frac{\breve{T}_{nr} l^2}{EI_o}$$

Non-repeatable Arm Deflection

- The non-repeatable arm deflection is equated to the 1-σ TMR,

$$\Delta s_{nr} = \frac{1}{2}\frac{\breve{T}_{nr}\, l^2}{EI_o} \approx O(1\sigma\ TMR) = \frac{1}{30\ TD}$$

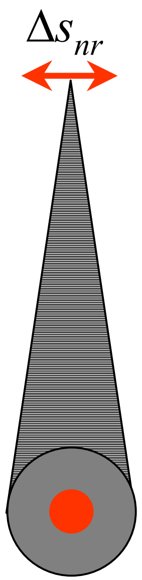

- The arm length is:

$$l^2 = \frac{2E \cdot I_o}{30 \cdot TD \cdot \breve{T}_{nr}}$$

For a constant disk radius, the ambient perturbations must decrease linearly with track density. If $\breve{T}_{nr}$ does not decrease, then the disk must become smaller.

Capacity vs. Areal Density (II)

$$l^2 = \frac{2E \cdot I_o}{30 \cdot TD \cdot \breve{T}_{nr}}$$

- Disk radius R ~ arm length
- With a constraint from irreducible $\check{T}_{nr}$, the capacity is:

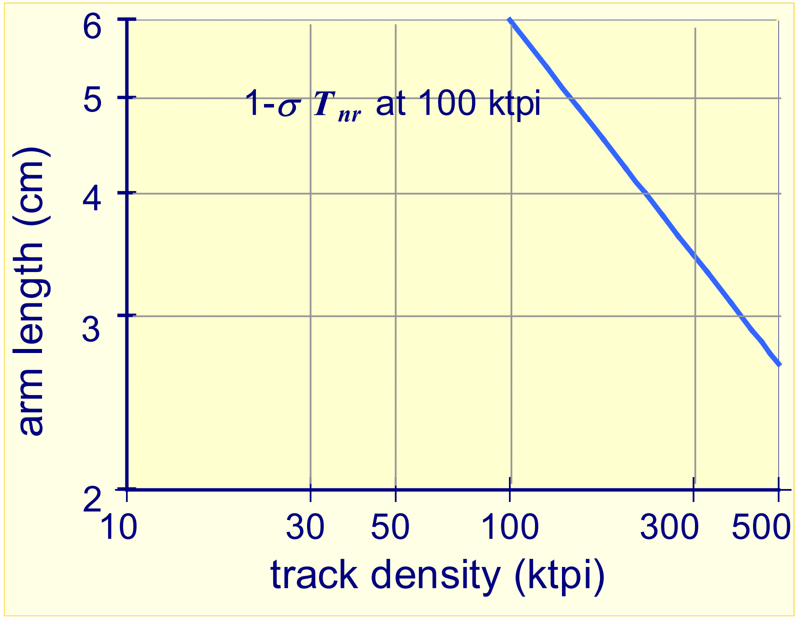

$$C_{nr} = 2\pi \left(\frac{2E \cdot I_o}{30 \cdot TD \cdot \breve{T}_{nr}} \right) \cdot LD \cdot TD = 2\pi \left(\frac{2E \cdot I_o}{30 \cdot \breve{T}_{nr}} \right) \cdot LD$$

Capacity vs. Areal Density (III)

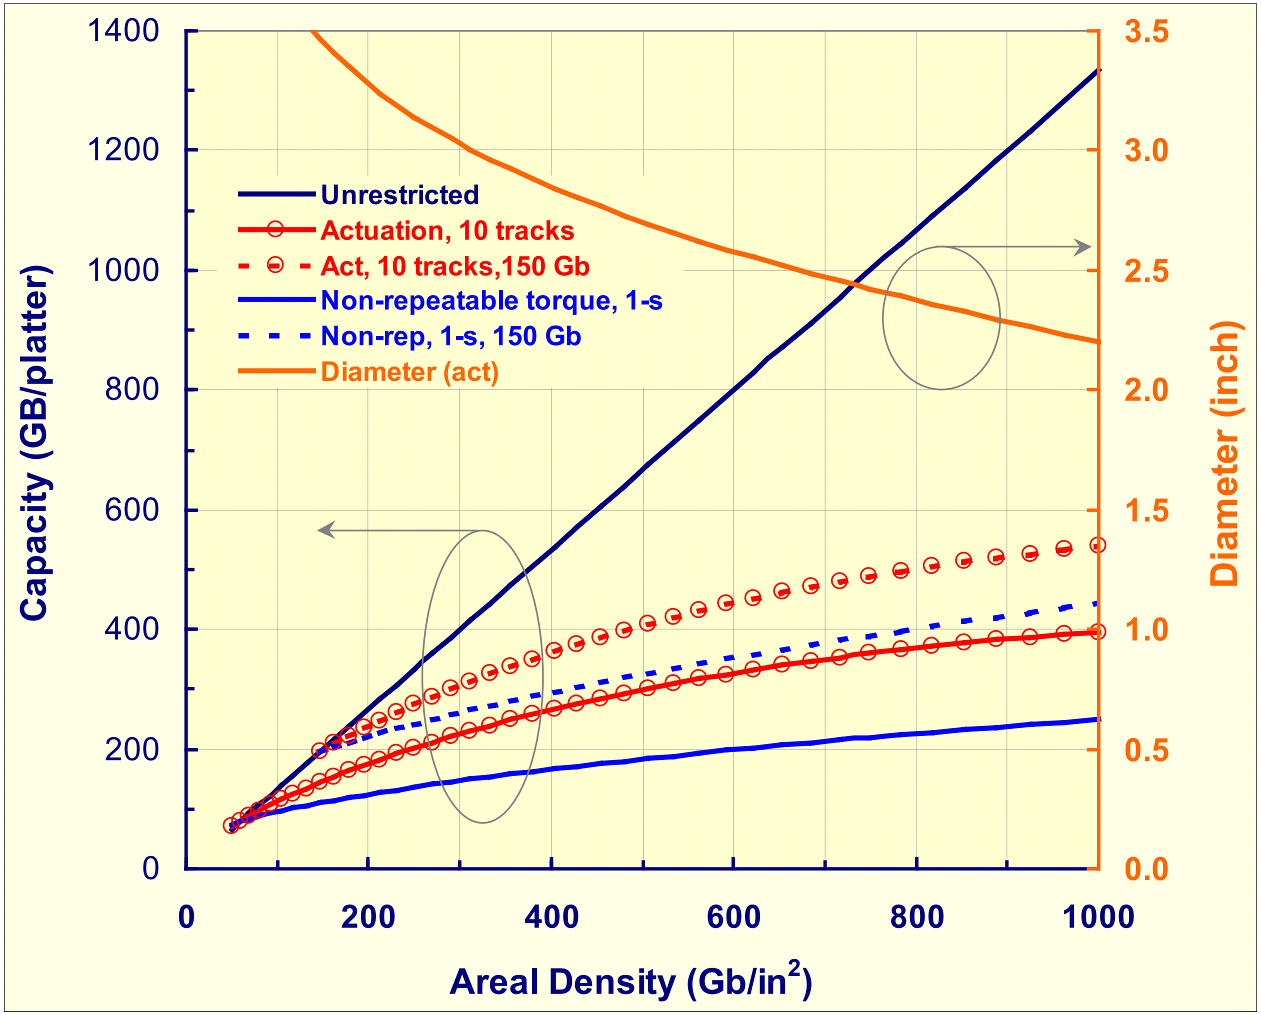

- BAR: 7.2 ➔ 4
- τ : 10 ➔ 5 ms
- 1 Tb/in^2 may not be cost-effective for producing a higher capacity drive with today's low price and ruggedness

Capacity ~ (Areal Density)$^{\alpha}$, 0< α<<1

- The requirement of a rugged, inexpensive device leads to a capacity growth significantly slower than the areal density growth
- The limited SNR expected at ~ 2 Mbpi likewise limits the capacity growth due to an increasing ECC overhead, e.g., 35% at 9.5 dB (rms/rms) in *Wood et al.* analysis
- There is a technological and economical optimum areal density that would enable powerful, useful products below 1 Tb/in^2

The 20 dB, 200 Gb/in^2, \$200 Challenge

Parameters		
AD	200	Gb/in^2
LD	1000	kbpi
TD	200	ktpi
SNRsys	≥ 20	dB
Thermal decay	< 0.5	%/decade
Rotation	≥ 7200	rpm
Disk	65	mm
Capacity	120	GB/platter
Transfer rate	≥ 250	Mbyte/s
Operating T	5 - 55	° C
Operating shock	200	G
No slider-level microactuator		

Drive		
Capacity	480	GB
Platters	4	
Heads	8	
Weight	155	g
Retail price	<<200	$

System		
Capacity	3.84	TB
Drives	8	
Physical size		
Width	9	in
Height	4.25	in
Depth	6	in
Price	<2,400	$

- 0.5 Terabyte in mobile format, high rpm & data rate
- Desirable, rugged, affordable, universally used. Raid 5.

The disk “tape” challenge

- Psychological adherence to tape
- A disk “tape” cartridge is possible and useful
- 200 GB in LTO/DLT cartridge format
- Non-op shock: 10-ft drop, any and all axes
- Hot-swappable
- Price = Tape cartridge/2
- Bandwith = 2 x BW(tape)
- Shelf-life: indefinite

The value ADD of AD Growth

- Optimum areal density is that which delivers:
  - Data Permanency
    - Absolute non-volatility // Fixed content // Regulatory compliance
  - Reliability
  - Ruggedness
  - SNR_{sys}
  - Bandwidth and IOPS
  - Product Longevity
    - Areal Density Plateau: Extended product cycles