

IEEE Magnetics Society
Santa Clara Valley Chapter

Shingled Magnetic Recording and Two-Dimensional Magnetic Recording

Roger Wood

Hitachi GST, San Jose, California

 **Hitachi Global Storage Technologies**

October 19th, 2010



Acknowledgements:

- **INSIC EHDR Program**
 - **for early support of the “TDMR” project**
- **SRC (Storage Research Consortium)**
 - **for support of “Shingled-Writing” Projects**
- **Universities and University colleagues**
 - **for their enthusiasm and hard work**
- **Committee 144* and Magnetism Soc. of Japan**
 - **for providing an excellent forum (PMRC 2010)**
- **Y. Shiroishi & I. Tagawa (Hitachi Ltd., CRL)**
 - **for particular help in preparing material for this talk**

*Committee 144 (Magnetic Recording) of the Japan Society for the Promotion of Science

Shingled-Writing and Two-Dimensional Magnetic Recording

Session Chairs: Roger Wood (Hitachi GST) and Ikuya Tagawa (Hitachi Ltd.)

“Future HDD Technologies and the Prospects for Shingled Recording”

W. Cain, E. Champion, C. Stevens (Western Digital)

“The Potential of Bit Patterned Media in Shingled Recording”

S. Greaves, H. Muraoka, Y. Kanai (Tohoku University, Niigata Institute of Technology)

“Shingle write Recording Assessment with Spin Stand Measurement”

H. Kiyono, O. Nakada, T. Mori, T. Oike (TDK Corporation)

“Minimization of Erase-Band in Shingled PMR with Asymmetrical Writer”

*I. Tagawa, Y. Urakami, M. Maeda, Y. Maruyama, K. Kudo, H. Shiina, M. Mochizuki
(Hitachi Ltd., Hitachi GST)*

“Drive Based Recording Analyses at >800 Gbit/in² Using Shingled Recording”


R. W. Cross, M. Montemorra (Seagate Technology)

“Investigation of Position and Timing Uncertainty of Two-dimensional Magnetic Recording (TDMR) at 4 Terabits per Square Inch”

E. Hwang, R. Negi, B.V.K. Kumar, R. Wood (CMU, Hitachi GST)

“Comparisons of One- and Two-dimensional Detectors on Simulated and Spin-stand Readback Waveforms”

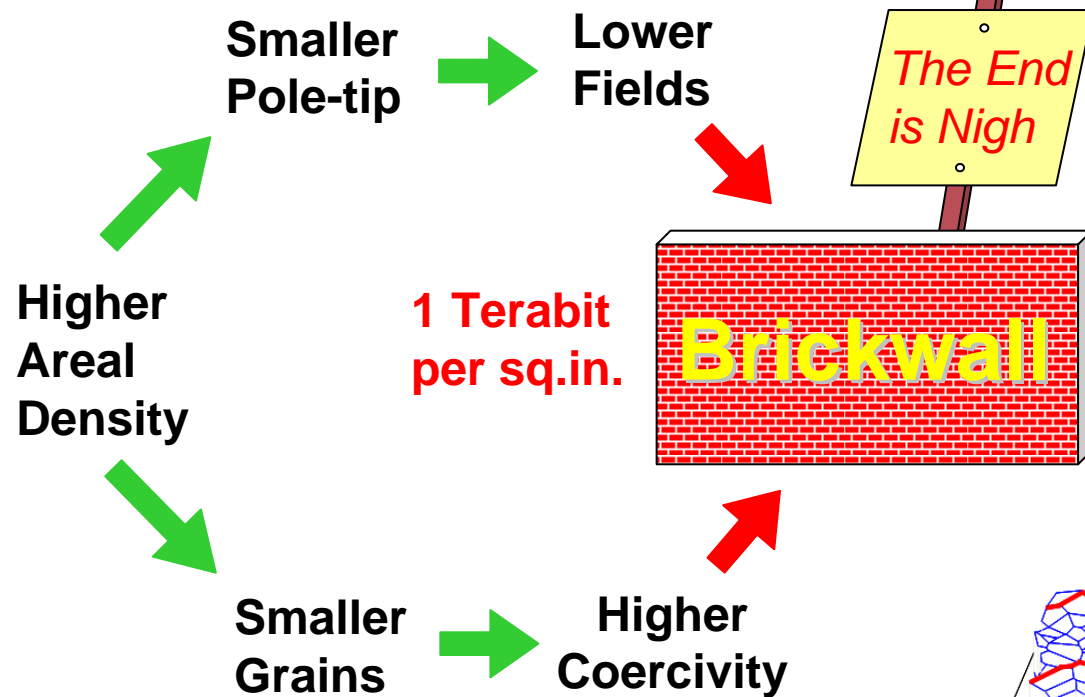
K. CHAN, M. Elidriss, K. Eason, R. Radhakrishnan, K Teo (DSI, Singapore)

- 
- Limits on Magnetic Recording
 - Technology options for 1 Tbit/sq.in & beyond
 - Shingled Magnetic Recording (SMR)
 - Two-Dimensional Magnetic Recording (TDMR)
 - Data architecture and Systems Issues
 - Future Scenarios

“Superparamagnetic Brickwall” – J. Judy

Diminishing returns for magnetic spacing

- already at about $\sim 1/2$ media thickness
- little freedom to reduce media thickness
(must maintain grain volume for thermal stability)

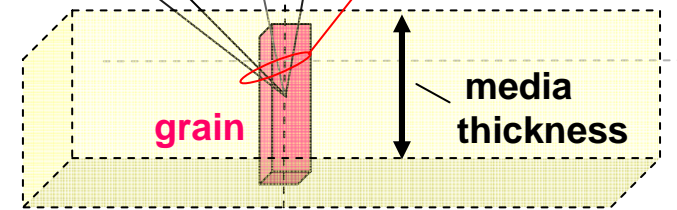


Trapezoidal Pole-Tip

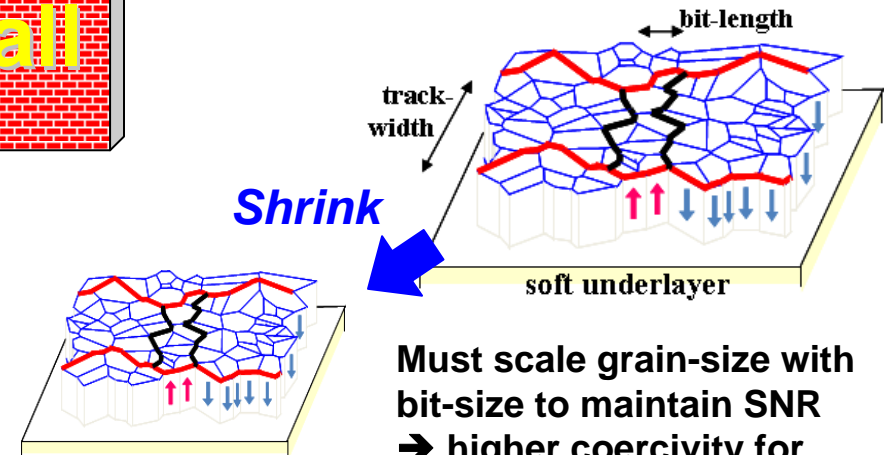
saturated pole-tip

Head field in medium is proportional to Solid-Angle

Solid-Angle shrinks with pole-tip



Shrink



Grain volumes shrink faster than scaling

- halving ‘grain-pitch’ leaves only 19% of core area
(9 → 4.5 nm pitch, assuming 1nm grain-boundary required)

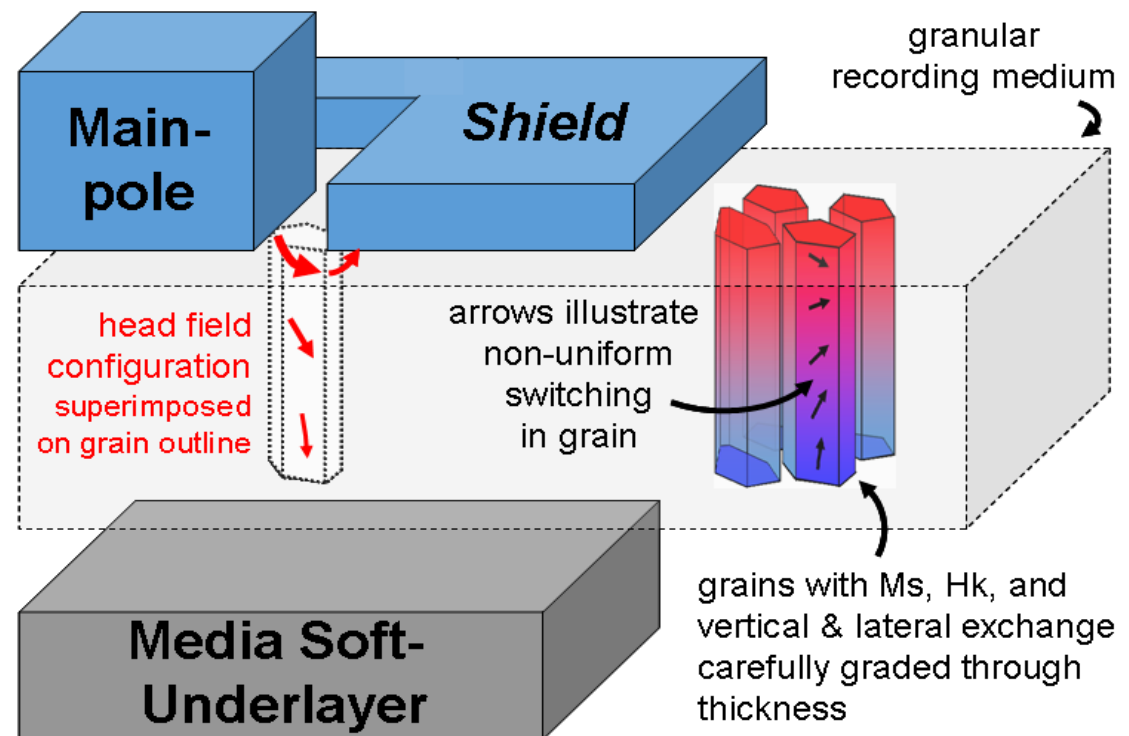
Limits on Grain size (Areal-Density)

- *Today's media:* $H_0 \approx 10$ kOe, $M_s \approx 500$ emu/cc, size $8 \times 8 \times 16$ nm $\rightarrow K_u V / kT = 90$
 - Little opportunity to increase H_0 (writability) or M_s (demag reduces stability)
 - Small reductions in $K_u V$ possible, but energy-barrier already $< K_u V$ due to demag.
 - Thicker media (smaller diameter grains) causes loss of vertical field strength

Only Limited gains available from 'graded' media:

Head field configuration is already close to ideal for switching grain:

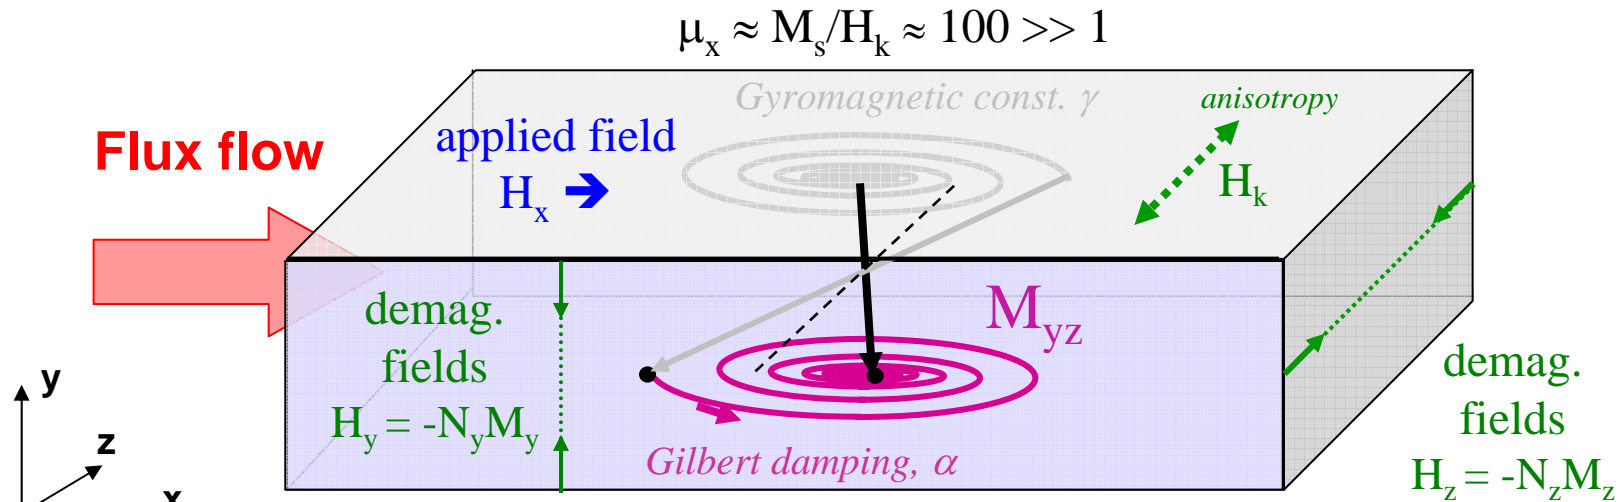
- 'uniform' vertical field to lower energy barrier
- strong in-plane field tweaks top of grain to initiate switching
- High gradients close to head ensure formation of sharp transition



R. Wood, J. MMM 321 (2009) pp. 555–561

Current Grain-sizes are close to limit for Conventional Recording

Flux flow in write-head structures limited by gyromagnetic effects



$$f_r \approx \gamma_e M_s \sqrt{[(N_y - N_z)/\mu_x]} - \text{resonant frequency}$$

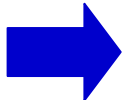
$$\xi \approx \alpha \sqrt{[(N_y - N_z)\mu_x]} - \text{damping factor}$$

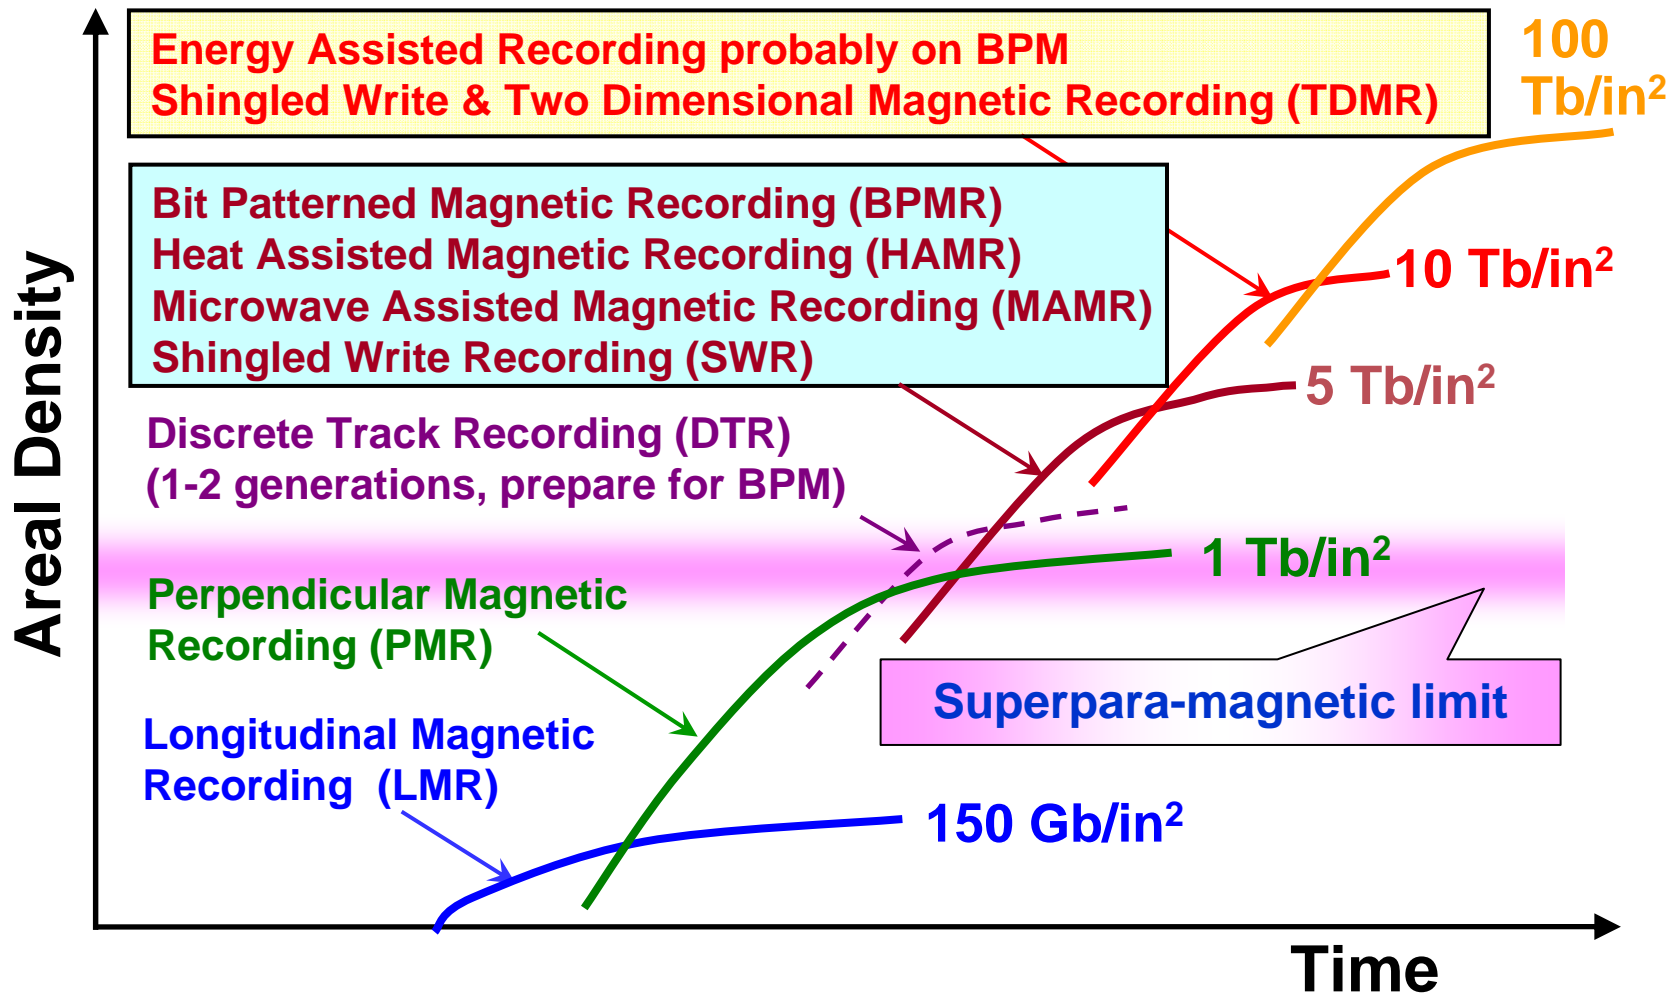
R. Wood et al., IEEE Trans. Magn.,
Vol. 38, pp. 1711-1716, July 2002

- Bandwidth limited by ferromagnetic resonance (and ability to adjust damping)
- Bandwidth depends on permeability, sat. magnetization, shape anisotropy
- Permeability of ~ 100 required if yoke is 100x longer than gap
- Maximum bandwidth ~ 4.8 GHz ($\mu = 100$, $M_s = 2.4$ T, $N_y - N_z = 0.5$, $\gamma = 28$ GHz/T)
- Need to write 3rd harmonic of “all-1’s” \rightarrow max channel data-rate = 3.2 Gb/s

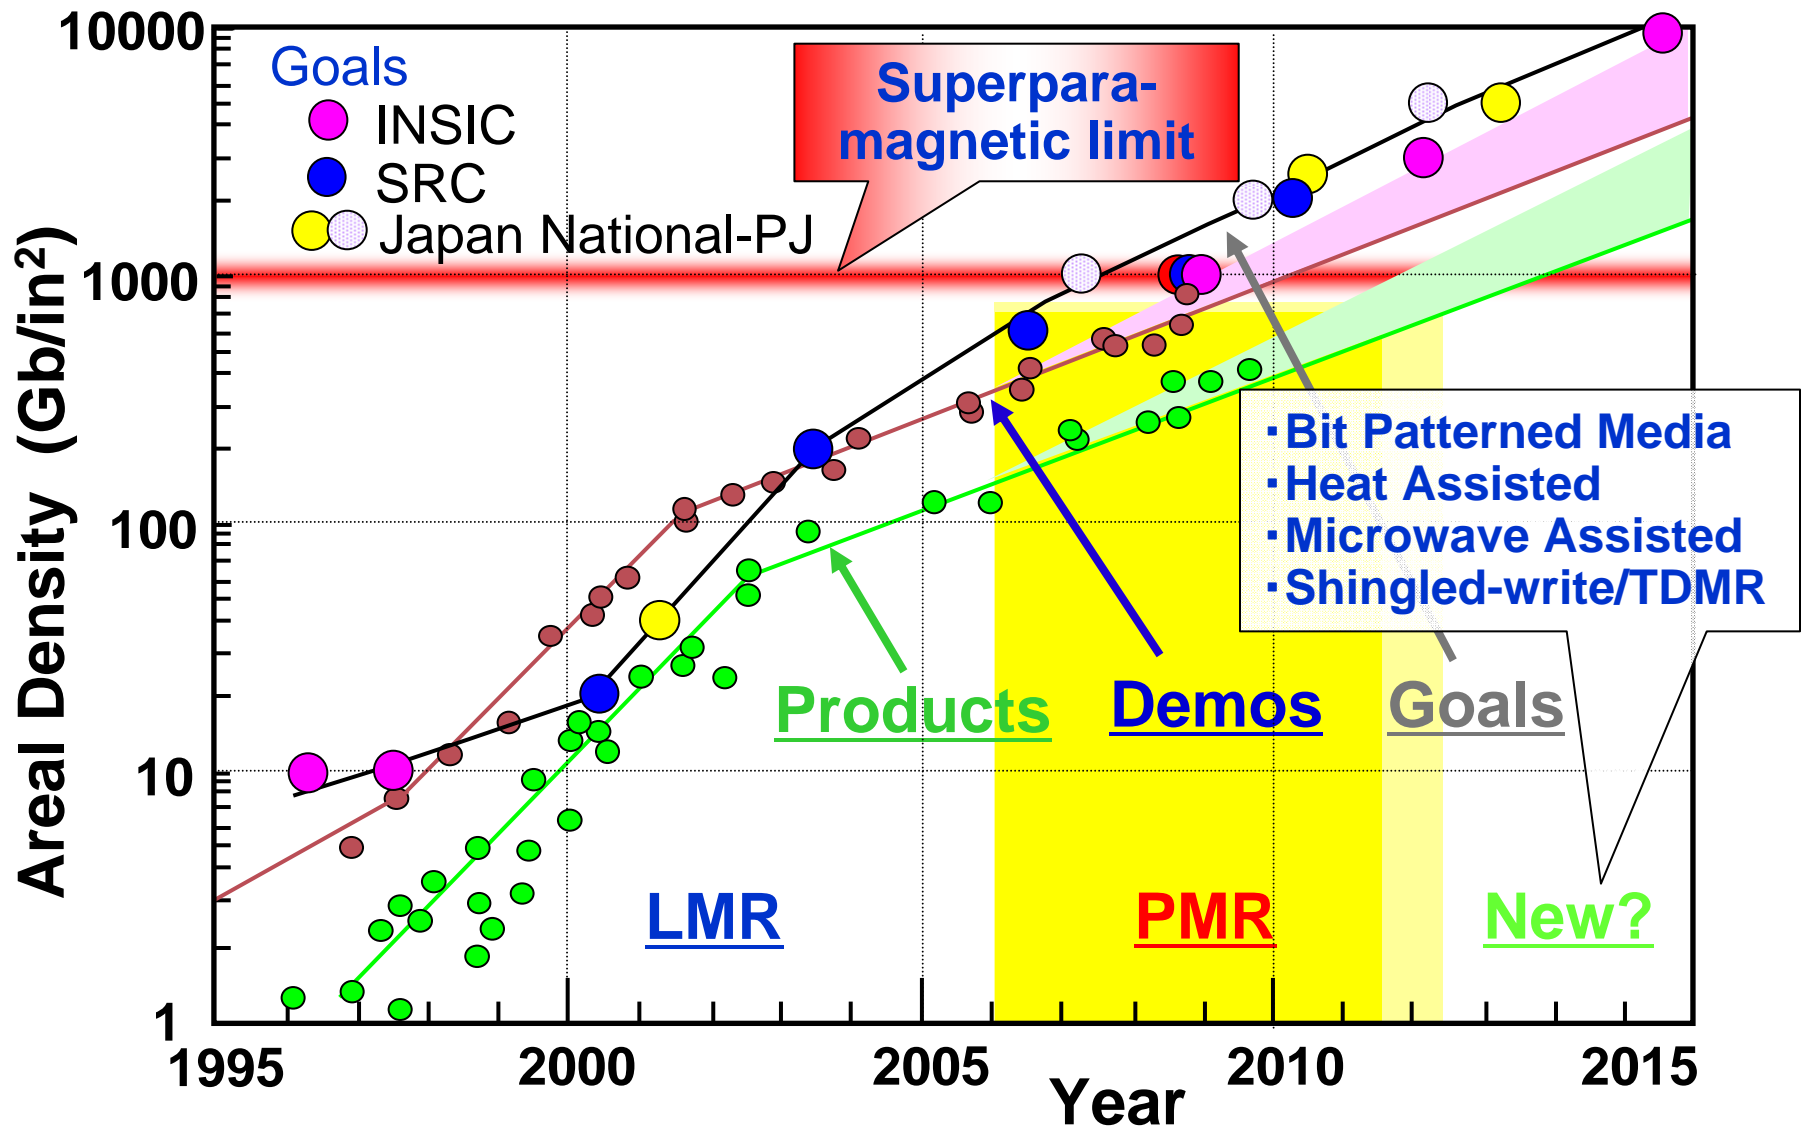
Maximum data-rate limited to approx. 3 Gbit/s for conventional recording

Future Technology Options & Limits for Hard Disk Drives

- Limits on Magnetic Recording
-  ■ Technology options for 1 Tbit/sq.in & beyond
- Shingled Magnetic Recording (SMR)
- Two-Dimensional Magnetic Recording (TDMR)
- Data architecture and Systems Issues
- Future Scenarios



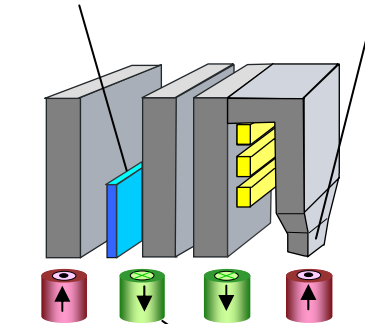
Y. Shiroishi, Intermag 2009, FA-01



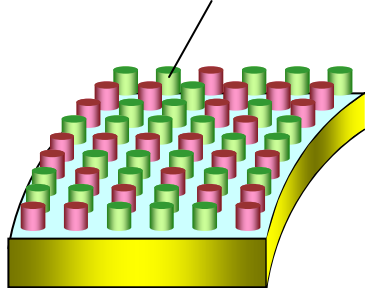
Y. Shiroishi, Intermag 2009, FA-01

Bit Patterned
Magnetic nano-islands w/
exchange coupled grains

Reader Main Pole

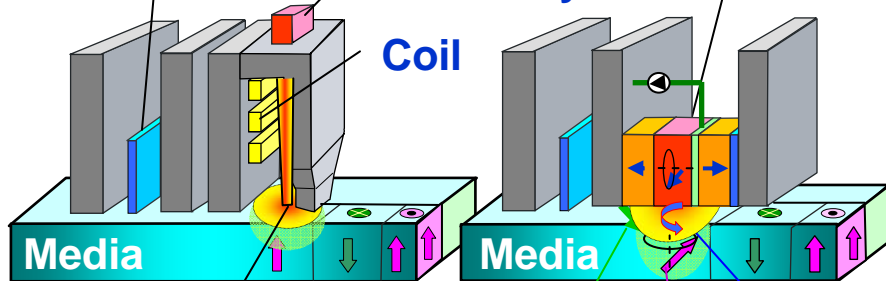


1 bit=1 Island

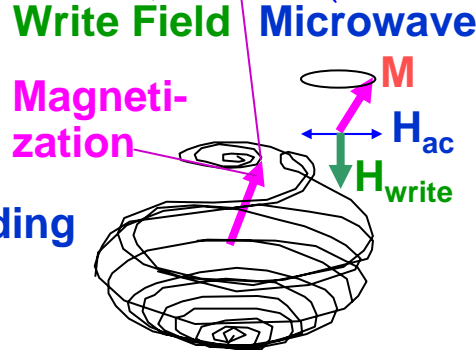
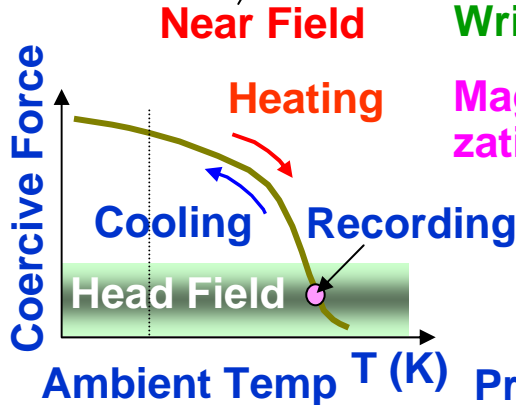


Heat/Microwave Assisted
Energy assisted writing to thermally
stable & hard-to-write media

Reader Laser Field Generating Layer

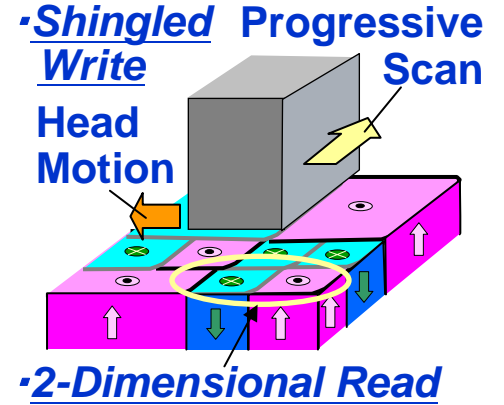
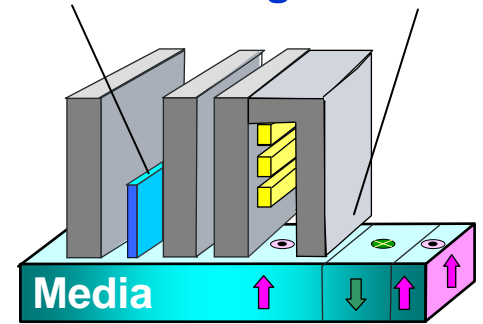


Coercive Force



SWR/TDMR
Shingled write w/ 2-dim
read & signal processing

Reader Shingled Writer



(a) Bit Patterned

(b) Heat Assisted

(c) Microwave Assisted

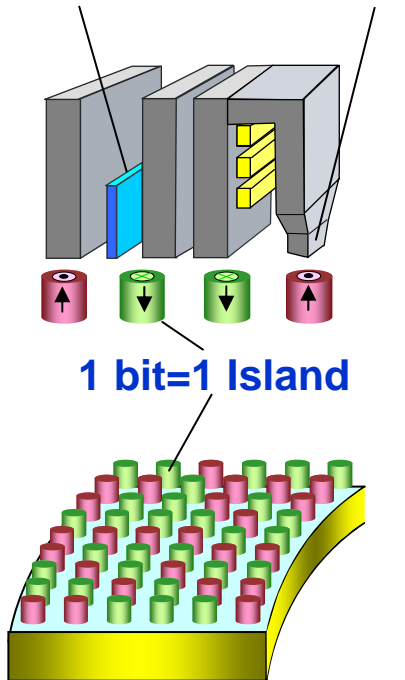
(d) SWR/TDMR

Y. Shiroishi, Intermag 2009, FA-01

Bit Patterned

Magnetic nano-islands w/
exchange coupled grains

Reader Main Pole



(a) Bit Patterned

Key Advantages:

1. Single large 'island' with well-defined position and geometry replaces the several smaller grains that are necessary if grain positions & geometries are unknown
2. Islands are separated by a lithographic process, so much more freedom to choose best material set (i.e. process not required to simultaneously create grain core and grain boundary)

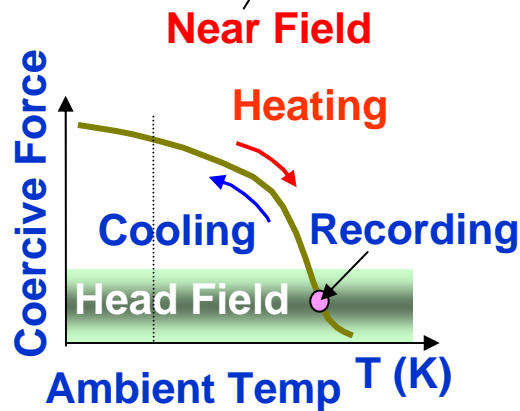
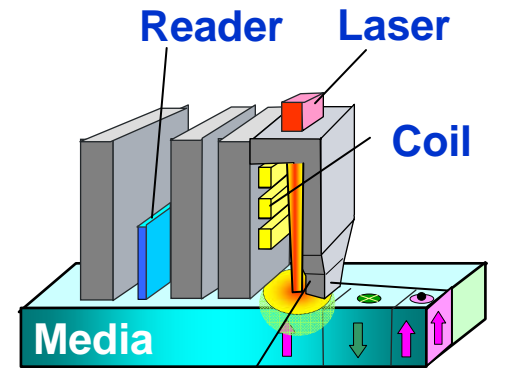
Challenges/disadvantages:

1. Extreme nano-lithography/imprinting requires massive/expensive change to disk mfg.
2. Intensive processing of disk surface may compromise head/disk interface & mag.-spacing
3. Write process must be accurately synchronized
4. 'Natural' 'self-assembled' hexagonal pattern is simplest to fabricate, but requires use of TDMR

Key Advantages:

1. Very high-anisotropy materials enable media with smaller grains
2. High temperature-gradients can write very 'sharp' transitions

Heat/Microwave Assisted
Energy assisted writing to thermally stable & hard-to-write media



(b) Heat Assisted

Challenges/Disadvantages:

1. Development of small-grain recording medium with good thermal and magnetic properties
2. Development of integrated optical and magnetic write-head
3. Elevated temperatures in head, medium, and at head-disk interface accelerate failure mechanisms

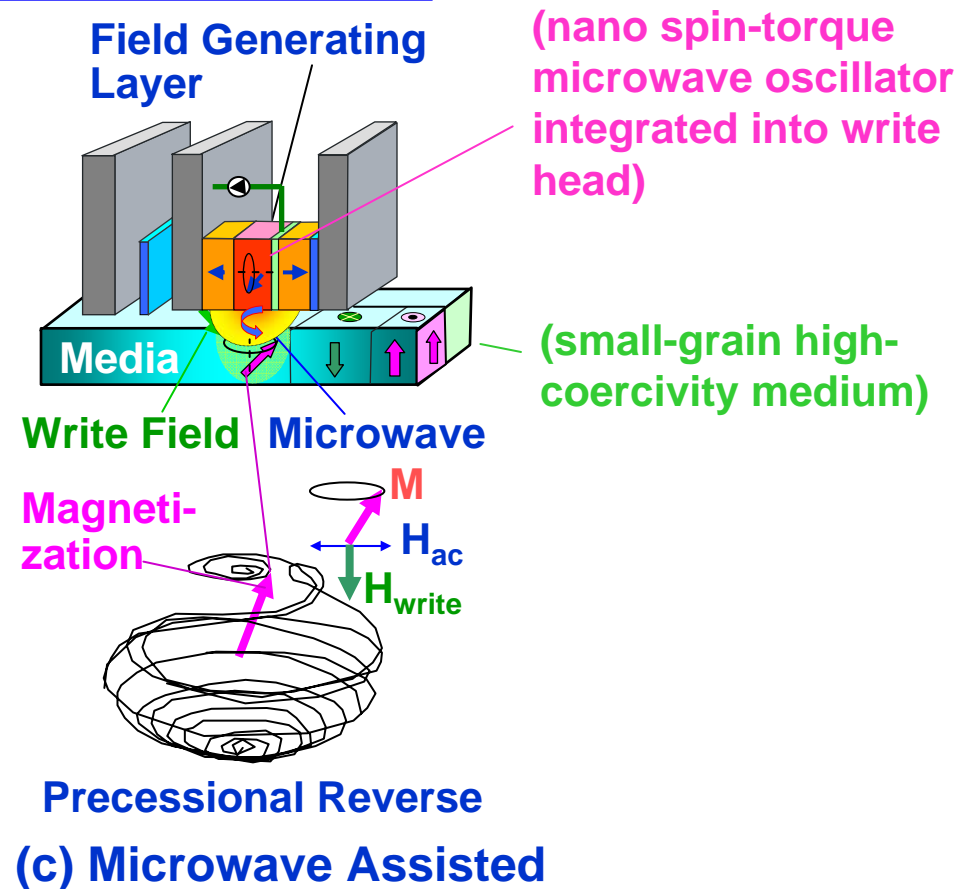
Heat/Microwave Assisted
Energy assisted writing to thermally stable & hard-to-write media

Key Advantages:

1. Best option for compatibility with current head and media manufacturing & HDD data architecture

Challenges/Disadvantages:

1. Development of small-grain recording medium with good microwave and magnetic properties
2. Development of nano-spin-torque oscillator with high power density
3. Experimental and theoretical development lagging

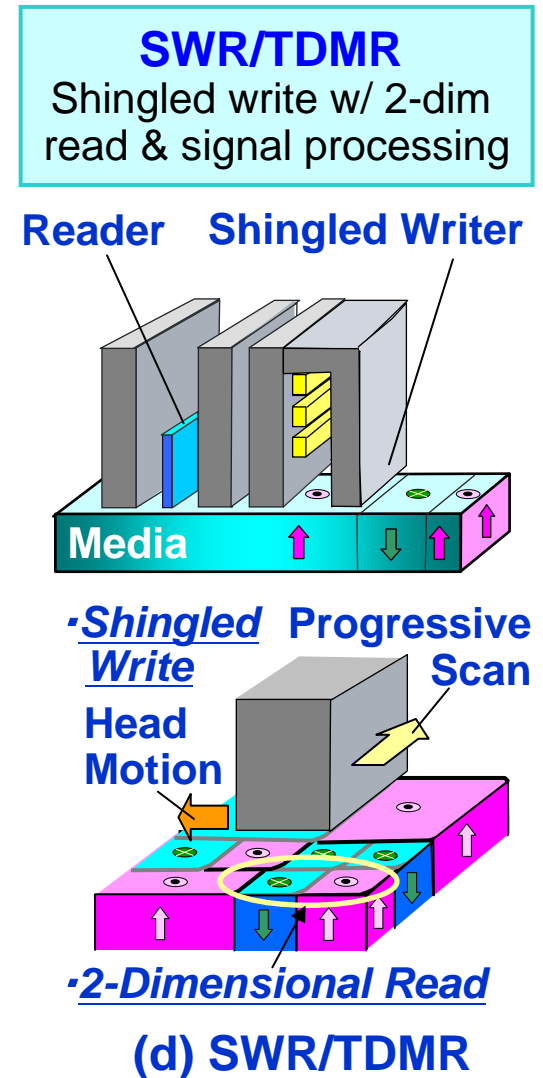


Key Advantage:

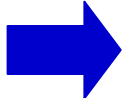
1. **Best option for compatibility with current head and media manufacturing processes**

Challenges/Disadvantages:

1. **Internal HDD data architecture must be changed or operating system modified. (“Shingling” places constraints on data flow onto disk that are incompatible with current HDD usage.)**
2. **Two-Dimensional Readback implies either several revs of latency, or a read head with three or more immediately adjacent sensor elements.**



Future Technology Options & Limits for Hard Disk Drives

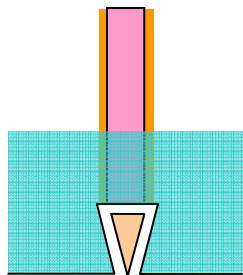
- Limits on Magnetic Recording
- Technology options for 1 Tbit/sq.in & beyond
-  ■ Shingled Magnetic Recording (SMR)
- Two-Dimensional Magnetic Recording (TDMR)
- Data architecture and Systems Issues
- Future Scenarios

Concept

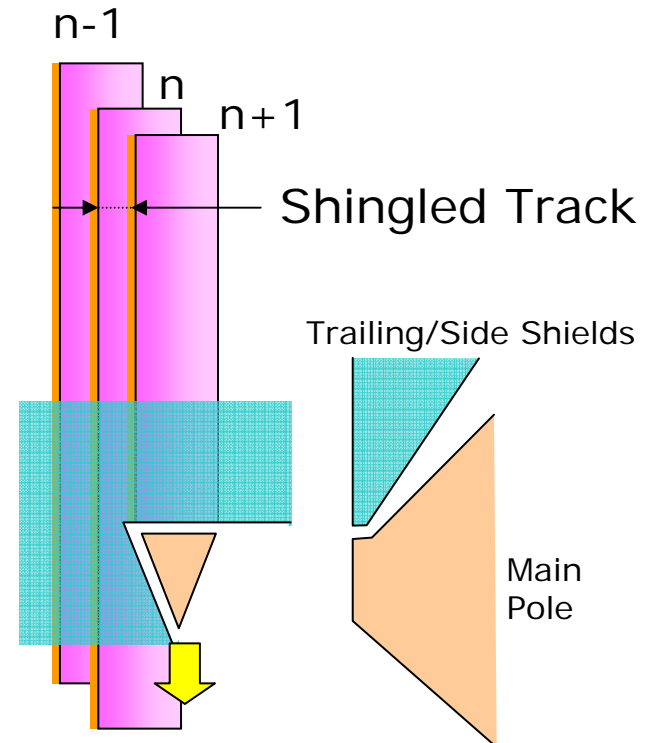
- Heavily overlapped writing at corner edge of wider head.
- Remaining narrow track reading.
- Sequential write, random read.

Extendability

- Shingle-write with normal random read access → Shingle-write with 2D readback (TDMR).



Conventional WAS writer
(PhysWW = ~25nm required for 2Tb/in²)



Corner Writer for Shingle
(PhysWW = ~70nm)

I. Tagawa, Intermag 2009, FA-02

Advantages

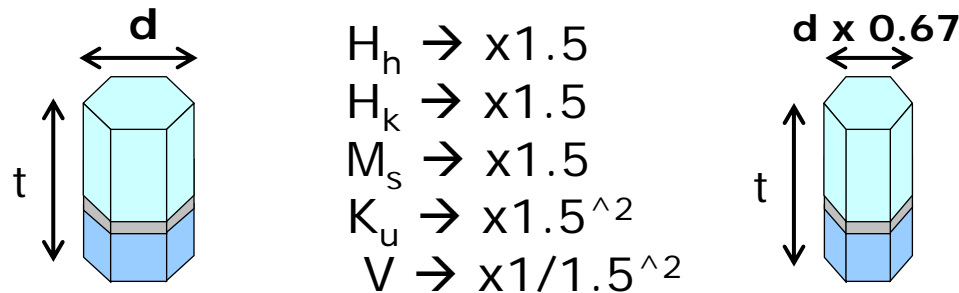
- Much stronger write field due to larger pole.
- No adjacent track erasure due to multiple write (ATE) brings further stronger field.
 - Stronger field brings improvement on linear density.
- Sharp corner-edge field brings narrower erase band.
 - These enables us to increase track density (TPI).
 - TPI gain expected even on conventional head at higher skew angle.
- Track pitch independent of head magnetic write width (MWW).
 - Much relaxed tolerance on MWW, while tight MWW screening required on discrete track & patterned media recording.

Challenges

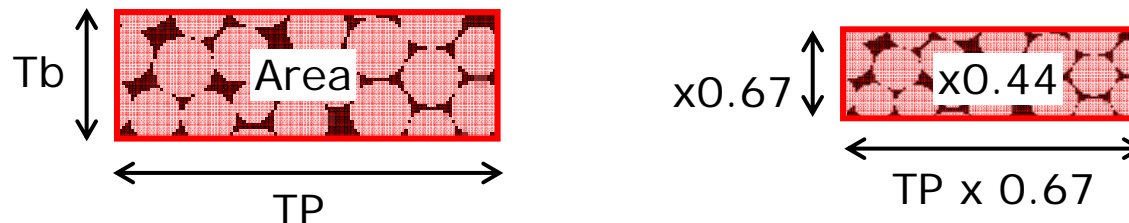
- New format architecture for random access emulation.
 - How to avoid performance loss (sustained data rate) due to Read-modify-write (de-fragmentation).
- Head wafer process development.
 - One-sided WAS structure should be optimized.

Gain prediction due to smaller grain

- 1.5 times larger write field H_h enables 33% reduction of grain diameter
 - ✓ x1.5 larger H_k possible without losing write-ability.
 - ✓ One assume to increase both H_k and M_s simultaneously.
 - ✓ $x1/1.5^2$ volume reduction with the same $K_u V/kT$.
 - ✓ Grain diameter reduced to $\sqrt{1/1.5^2}=0.67 \rightarrow 33\%$ smaller

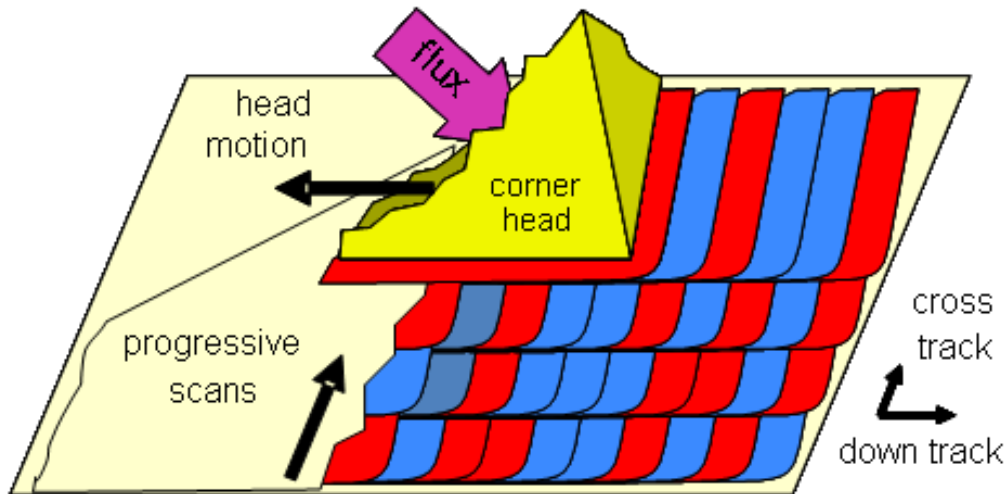


- 33% smaller grain brings more than twice (x2.25) areal density.
 - ✓ Assuming the same # of grains in a bit, occupied area reduced to 0.44
 - ✓ Areal density gain potential; $1/0.44=2.25 \rightarrow 125\%$ gain



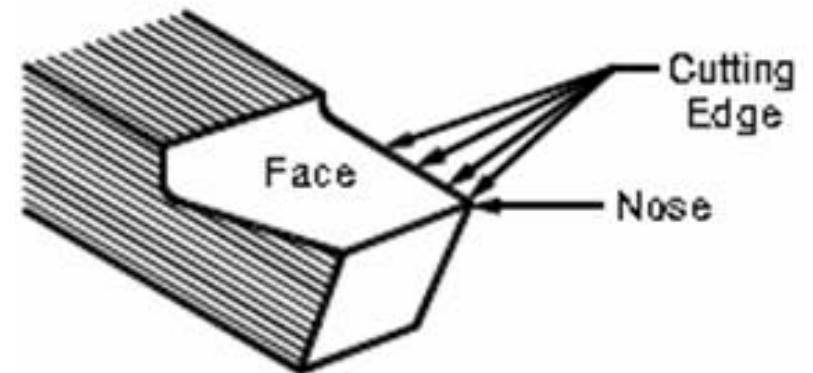
Intermag 2009, FA-02

I. Tagawa, Intermag 2009, FA-02

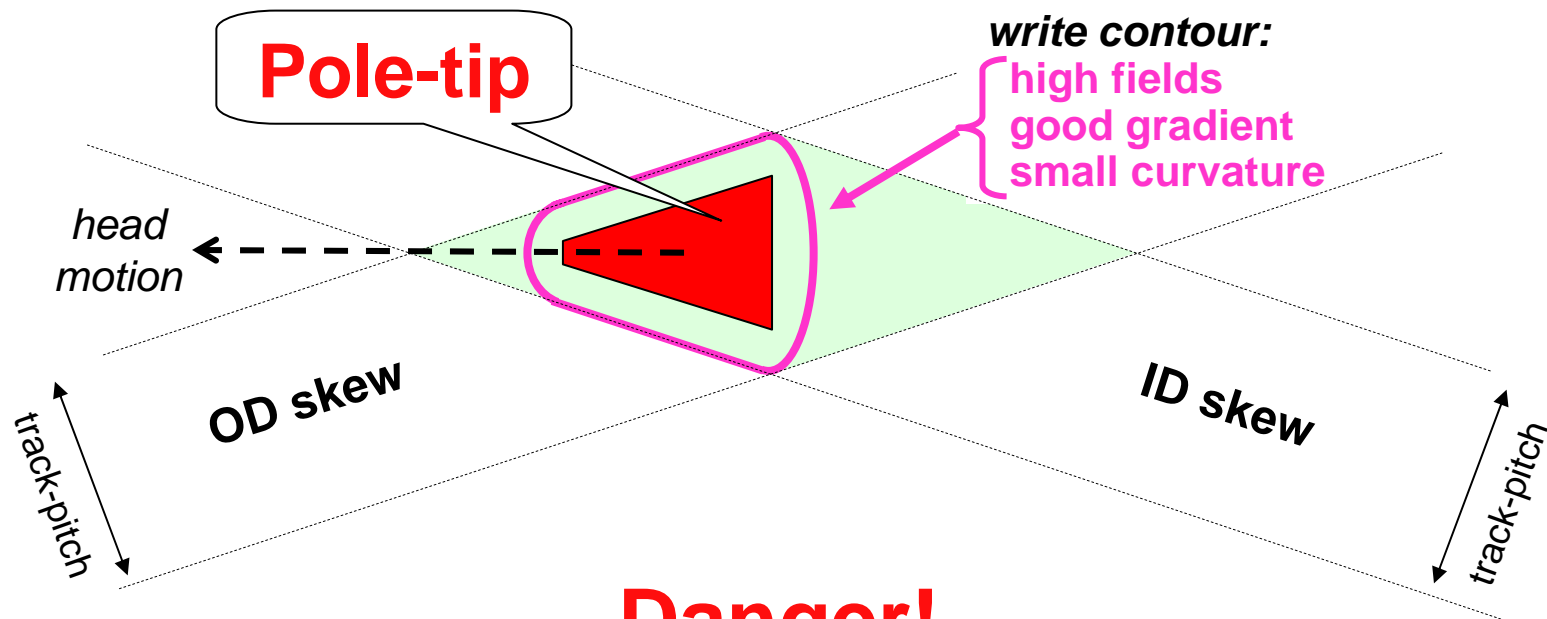


Shingled Writing Head

Cutting tool for lathe or milling machine



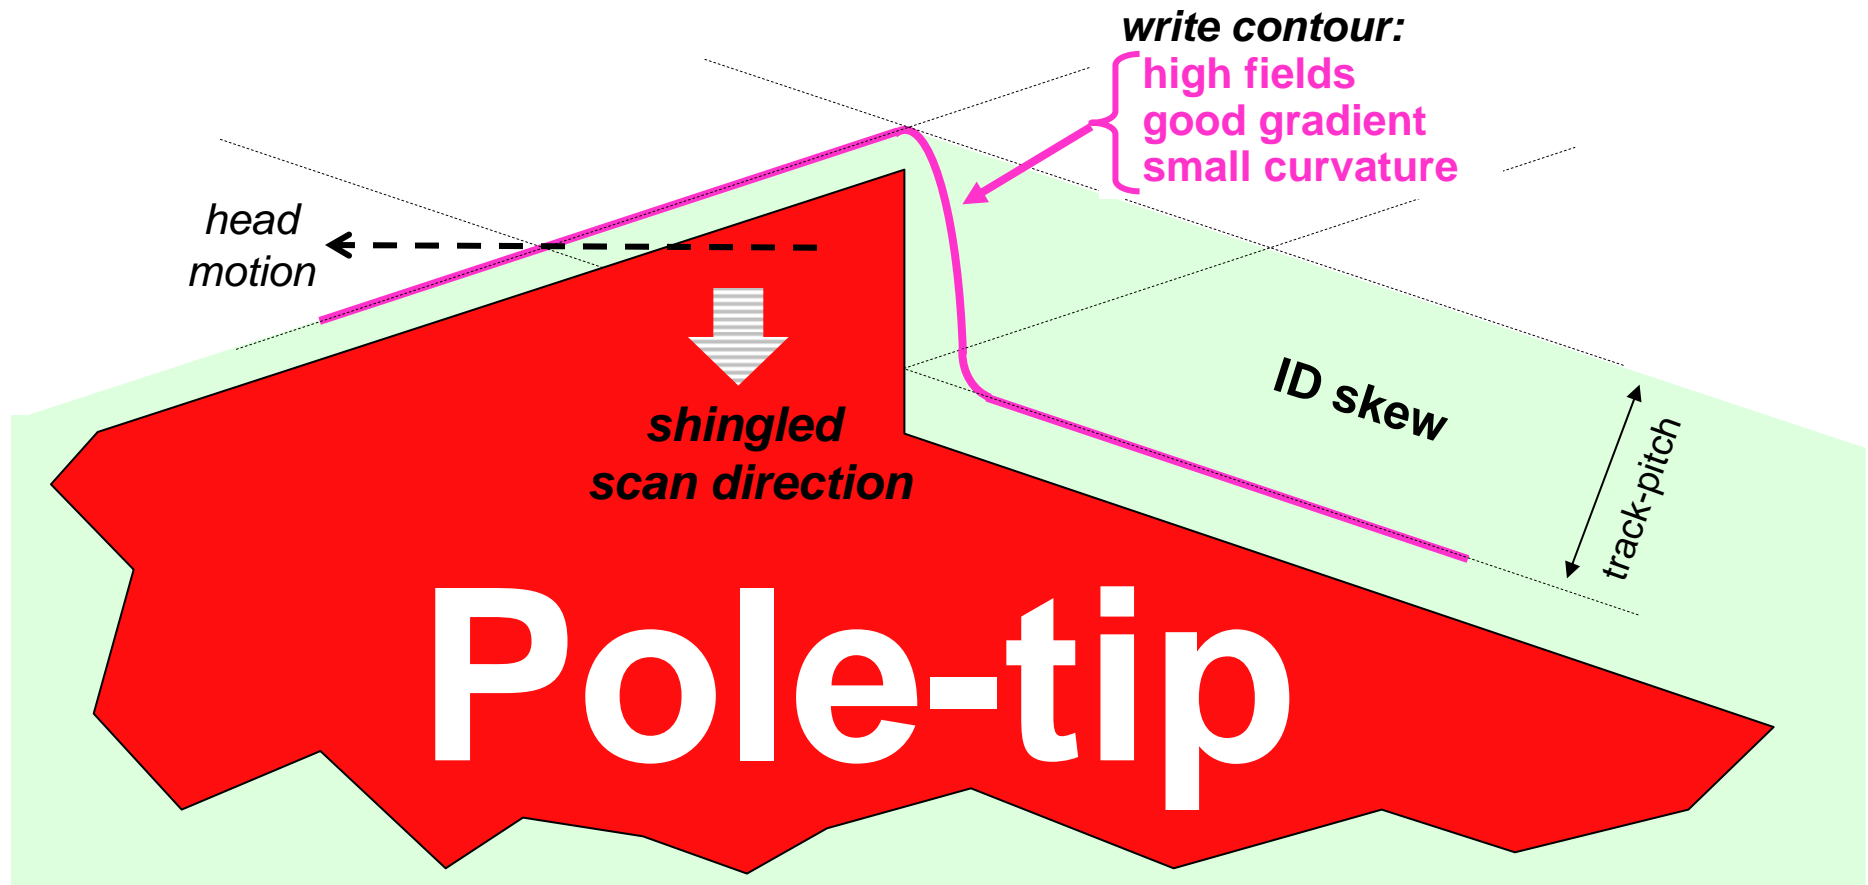
- Size and shape of pole-tip is *extremely* constrained
 - especially given need to operate over wide range of skew angles



Danger!
**No significant fields allowed outside green area
(risk of adjacent track erasure!)**

huge design freedom available: large pole-tip → high fields


no ATI problem, but write-field must not extend up here



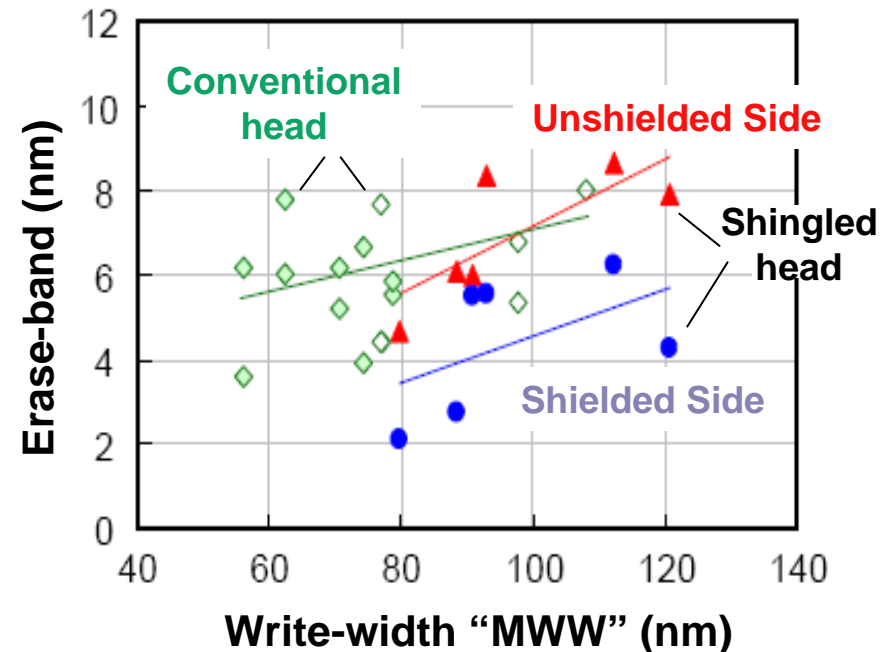
Shingled Head: design freedoms!

- **No need to shield both sides of head**
(shields suck flux from pole-tip and reduce fields)
- **Tighter Side-gap improves side-gradients**
 - higher side-gradients give smaller 'erase' band between tracks
(erase-band has no signal & lot of noise → reduced TPI capability)

ABS shape observed with SEM



Shingled head fabricated by Hitachi



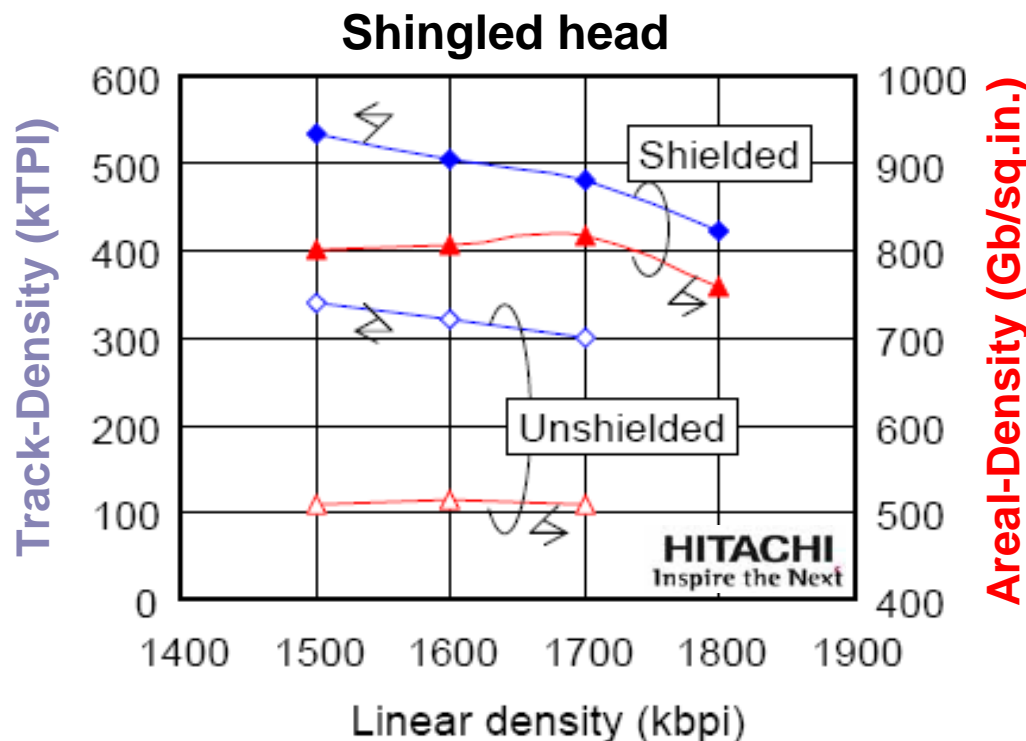
I. Tagawa et al., PMRC 2010

Shielded-side SMR-TPI is ~40% higher than unshielded side

- TPI failure-point defined at bit error-rate of 10^{-3}
- SMR Track-Pitch = One-side-squeeze Track-Pitch-Failure point

> 800 Gb/in² feasibility confirmed at 1700 kTPI

- TPI drops with density, but Areal-density almost constant at 800 Gbit/in²



**Spin-stand
measurement:**

**SMR head
with 1-side
shield**

**Conventional
recording
medium**

I. Tagawa et al., PMRC 2010

Many expected advantages and a concern

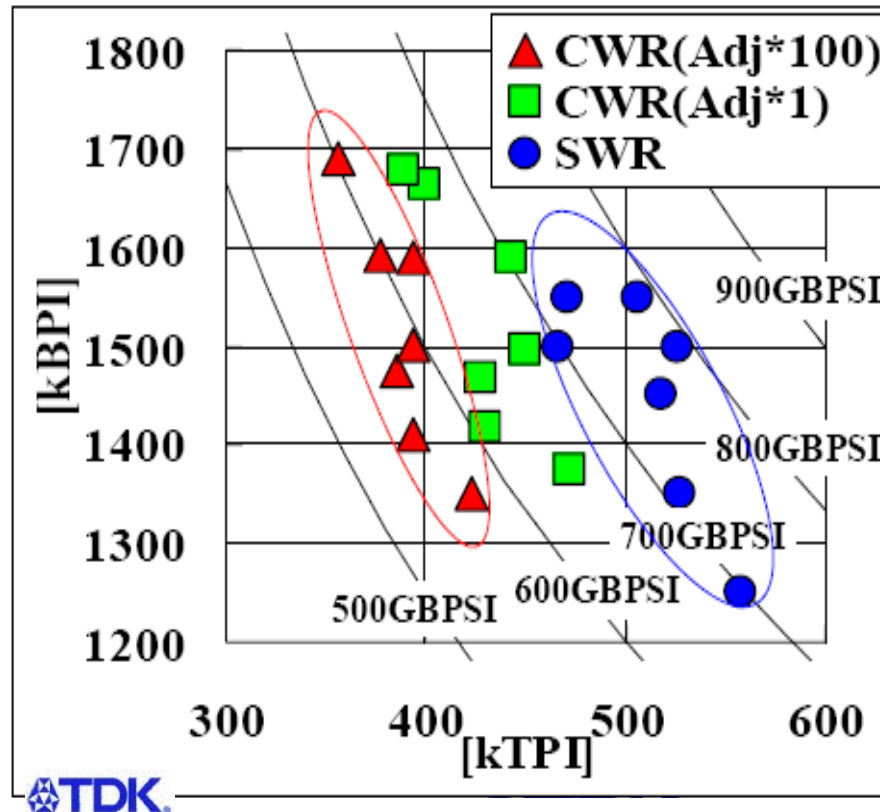
Advantages

- **Higher write-ability to breakthrough the famous tri-lemma**
- **Wider data track due to single side squeeze**
- **Less skew sensitivity**
- **Usage of continuous media**
- **Less pole width related yield losses**
- **Less adjacent track erasure concern**
- **Less pole erasure concern**

Concern

- **Slow HDD response**

- TPI was improved by 31% by Shingled-Writing, BPI dropped by 4% and in total the Areal-Density was improved by 26% from Conventional Recording with 100 adjacent track writes.
- Approximately half of the TPI gain came from the reduction of number adjacent writes from 100 to 1, and rest of the gain came from reduced track-width.



Spin-stand evaluation:

5400 rpm, 2.5" MD condition

Error-rate threshold = $10^{-2.5}$

Write-width: MWW+ 2EB = 62 nm

Read-width: MRW_μ = 38 nm

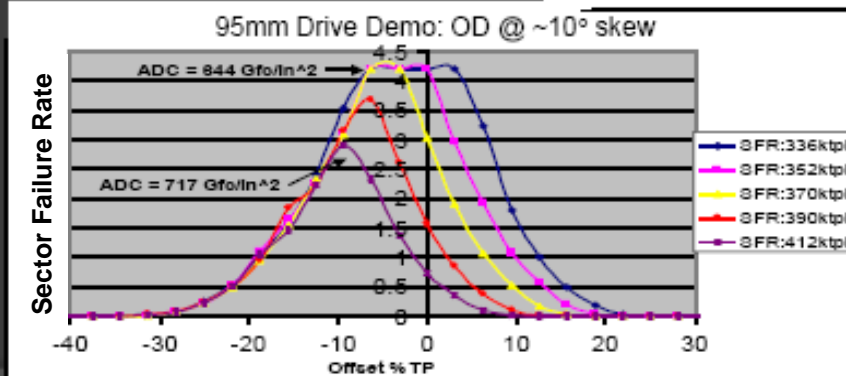
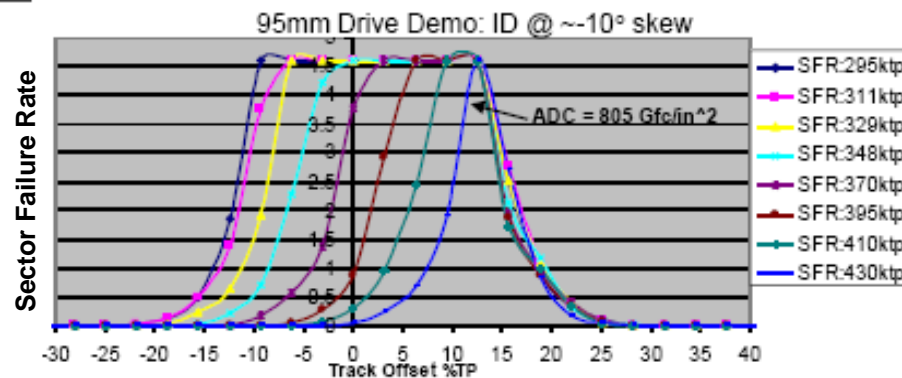
SWR = Shingled-Write
Recording
CWR = Conventional Write
Recording

H. Kyono et al., PMRC 2010

TMR Limits of SMR: 95mm example

- NRRO, especially windage effects, significantly degrade achievable areal density

Note: 95mm drive data shown to illustrate TMR effects (TMR = track-misregistration)



Conventional head:

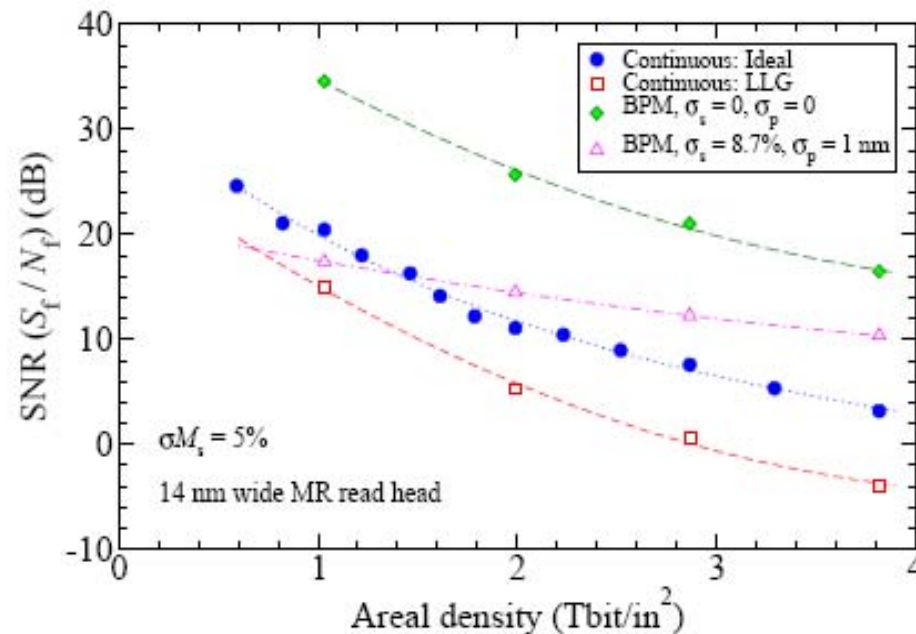
Physical Writewidth	85nm
Pole wall angle	12 °
Trailing pole gap	30nm
Side shield gap	65nm
Physical readwidth	38nm
Reader shield to shield	23nm

PMRC 2010

© Seagate Technology

B. Cross, M. Montemorra, PMRC 2010

SNR from Micromagnetic Simulation



Comparison of SNR for continuous and patterned media

Continuous media: BAR = 1:1

BPM: Dot diameter = $0.5 \times$ track pitch

- SNR of written tracks was calculated for continuous and patterned media.
- Patterned media had the highest SNR for areal densities exceeding 0.7 Tbit/in^2 .

S. Greaves, et al., PMRC 2010

Future Technology Options & Limits for Hard Disk Drives

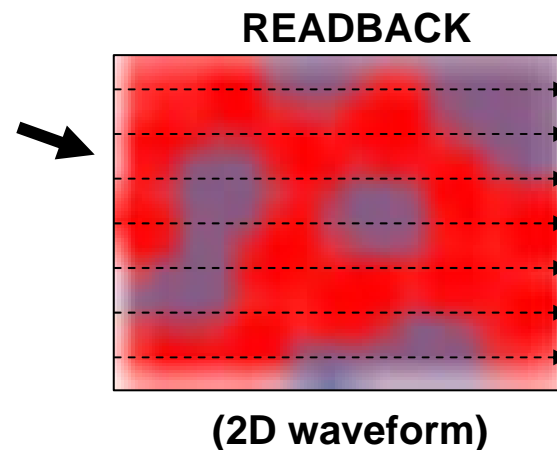
- Limits on Magnetic Recording
- Technology options for 1 Tbit/sq.in & beyond
- Shingled Magnetic Recording (SMR)
- ➔ ■ Two-Dimensional Magnetic Recording (TDMR)
- Data architecture and Systems Issues
- Future Scenarios

Two-Dimensional Magnetic Recording (TDMR)

= Shingled Write + 2D Readback

- In conventional systems, data is recovered from a single waveform gathered along the center of the wanted data-track
 - Any inter-track interference (ITI) is inevitably destructive
- In 2D-readback, a complete 'picture' is built up from multiple tracks
 - ITI is no longer destructive. ITI contains information about the data that powerful detectors can extract, just as is done with ISI

2D 'image' is built up either by multiple-passes of a single head or by a single pass of a multi-track reader



Accurate timing & position of bit-cells must be established

Powerful 2D detection & decoding algorithms recover the original user data

'2D-readback' opens a huge array of new possibilities:

- **Powerful 2D-equalization, detection, & decoding that treats inter-track interference just like inter-symbol interference**
- **Read head-width can be comparable to or larger than the track-pitch, control of read-width can be relaxed**
- **'Electronic fine-tracking' that relaxes the need for ultra-accurate mechanical track-following (*still need accurate mechanical track-following during write*)**

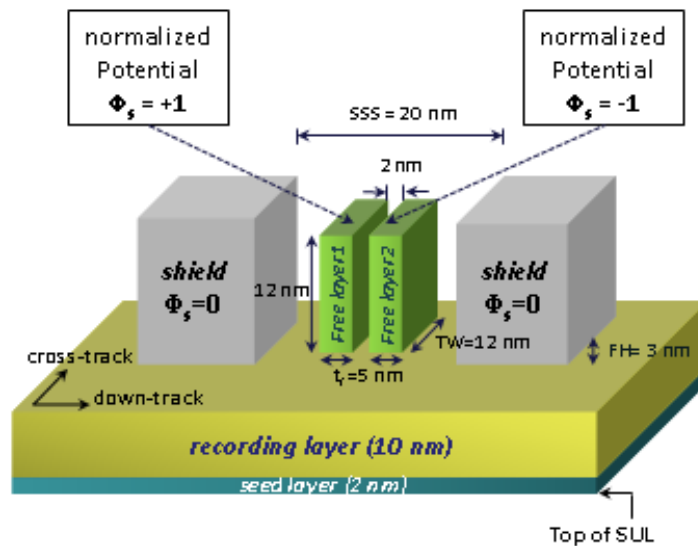
'2D readback' requires either an array-head or consecutive passes of a single head:

- **Building a 2-D image by consecutive incremental passes of a single head will require extra silicon memory and involve an increase in latency by at least one and perhaps several revs.**
- **An array head that addresses immediately adjacent tracks is very difficult to build. Using widely spaced heads staggered along a shallow slope is easier, but implies a linear actuator**

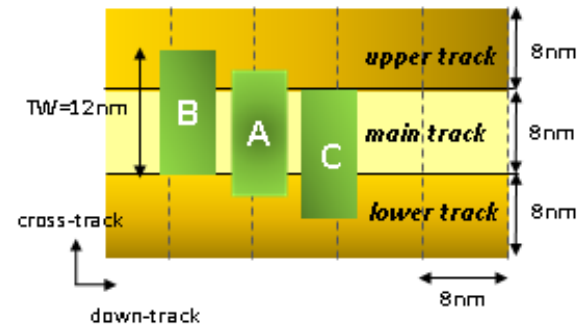
3-element array of differential-sensors staggered across track

Geometry of Dual Differential Head

o Schematic design of a dual differential head above a recording layer



o Planar view of an array of three dual differential heads on continuous media



For composite media,
$$V_{1\text{bit}} = \int_{\text{media}} \Phi_{\text{head}} \cdot \nabla M$$

The head magnetic potentials are obtained by means of subtraction of those on the middle from on the top of the recording media. ($\Phi_{\text{top}} - \Phi_{\text{middle}}$)



UNIVERSITY OF MINNESOTA

E. Cho, Y. Dong, R.H. Victora, INSIC Annual Meeting 2010

2D readback example: MFM Images

- Initial study conducted by looking at a series of high-resolution MFM images from 2007 (250 Gb/in² components)

Shingle to approach 1Tb/sq. in.

511-bit PRBS data
Shingle-write

1) 250 kTPI x 1270 kBPI
= 318 Gb/sq.in.

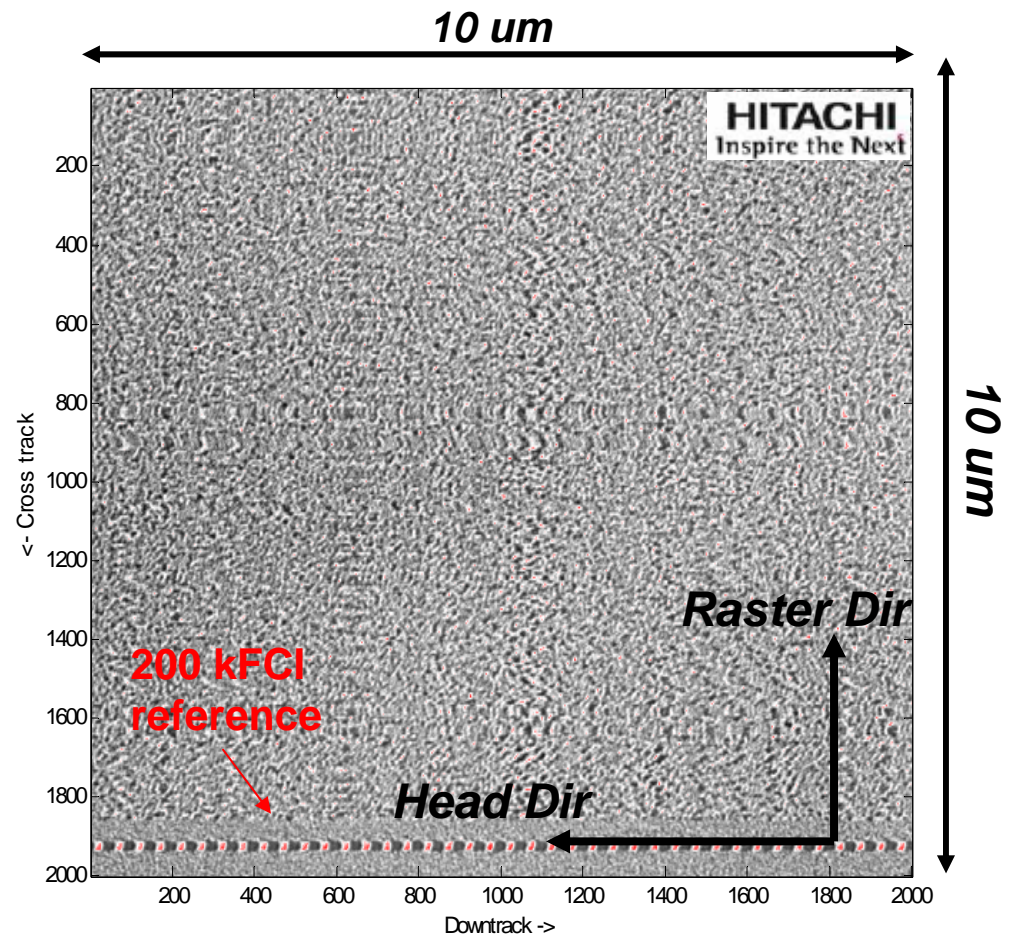
2) 500 kTPI x 1270 kBPI
= 635 Gb/sq.in.

3) 500 kTPI x 1483 kBPI
= 742 Gb/sq.in.

4) 700 kTPI x 1483 kBPI
= 1039 Gb/sq.in.

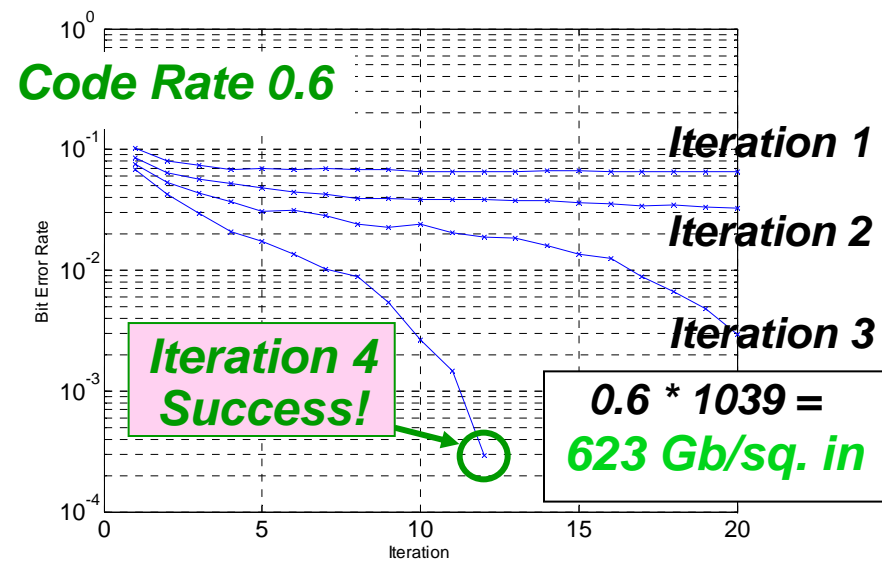
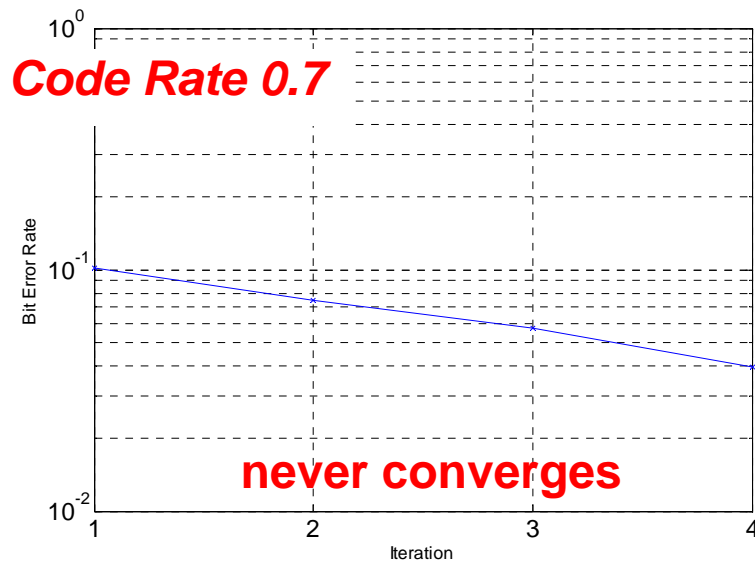
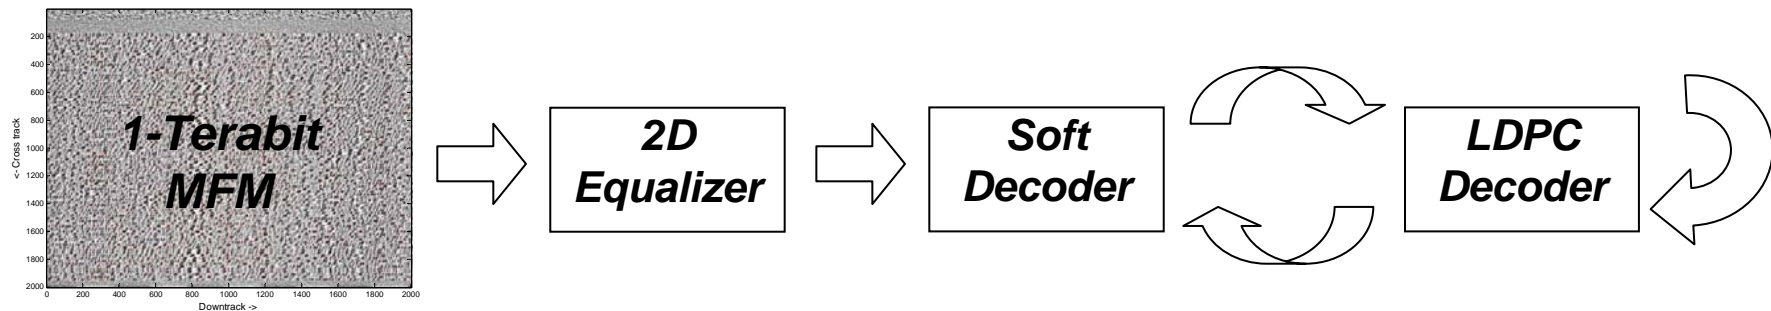


Each image is 10 μm x 10 μm
2000 x 2000 pixels



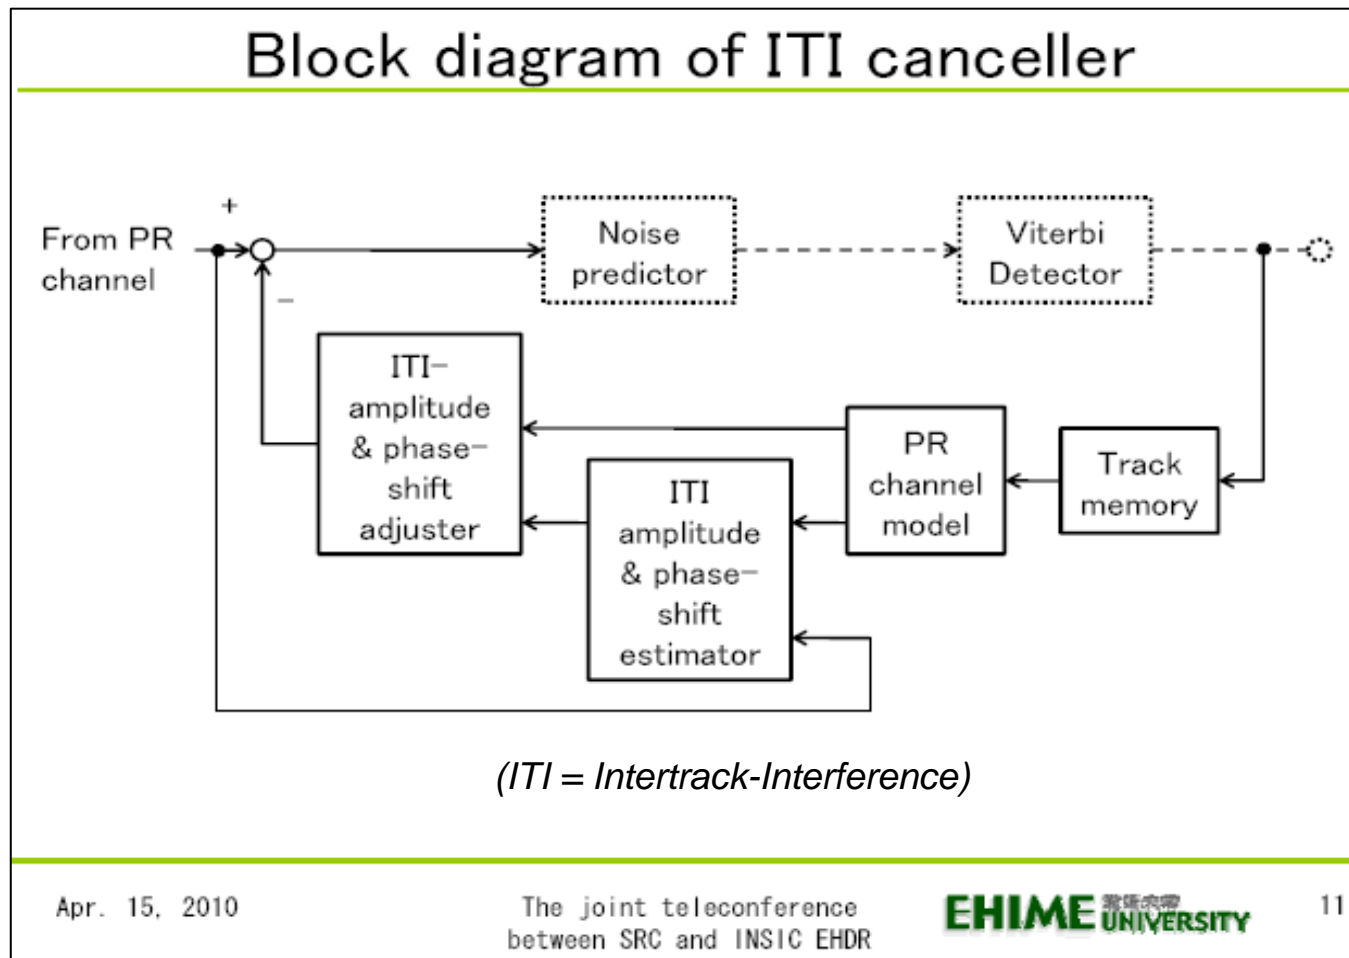
2D Signal Processing to recover data

- Is it possible to recover any data from that 1 Tb/sq. in MFM?
 - Yes, using an LDPC iterative decoding scheme.
 - Succeeded at code rate 0.6 → 623 Gb/sq.in.



F. Lim et al, Intermag 2010

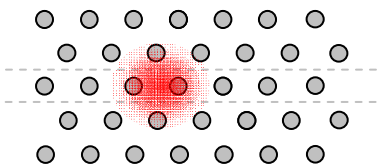
- Proposed for sequential readback of tightly Shingled tracks
 - Data from prior track is held in memory until it can be used to subtract the interference it causes on the subsequent track



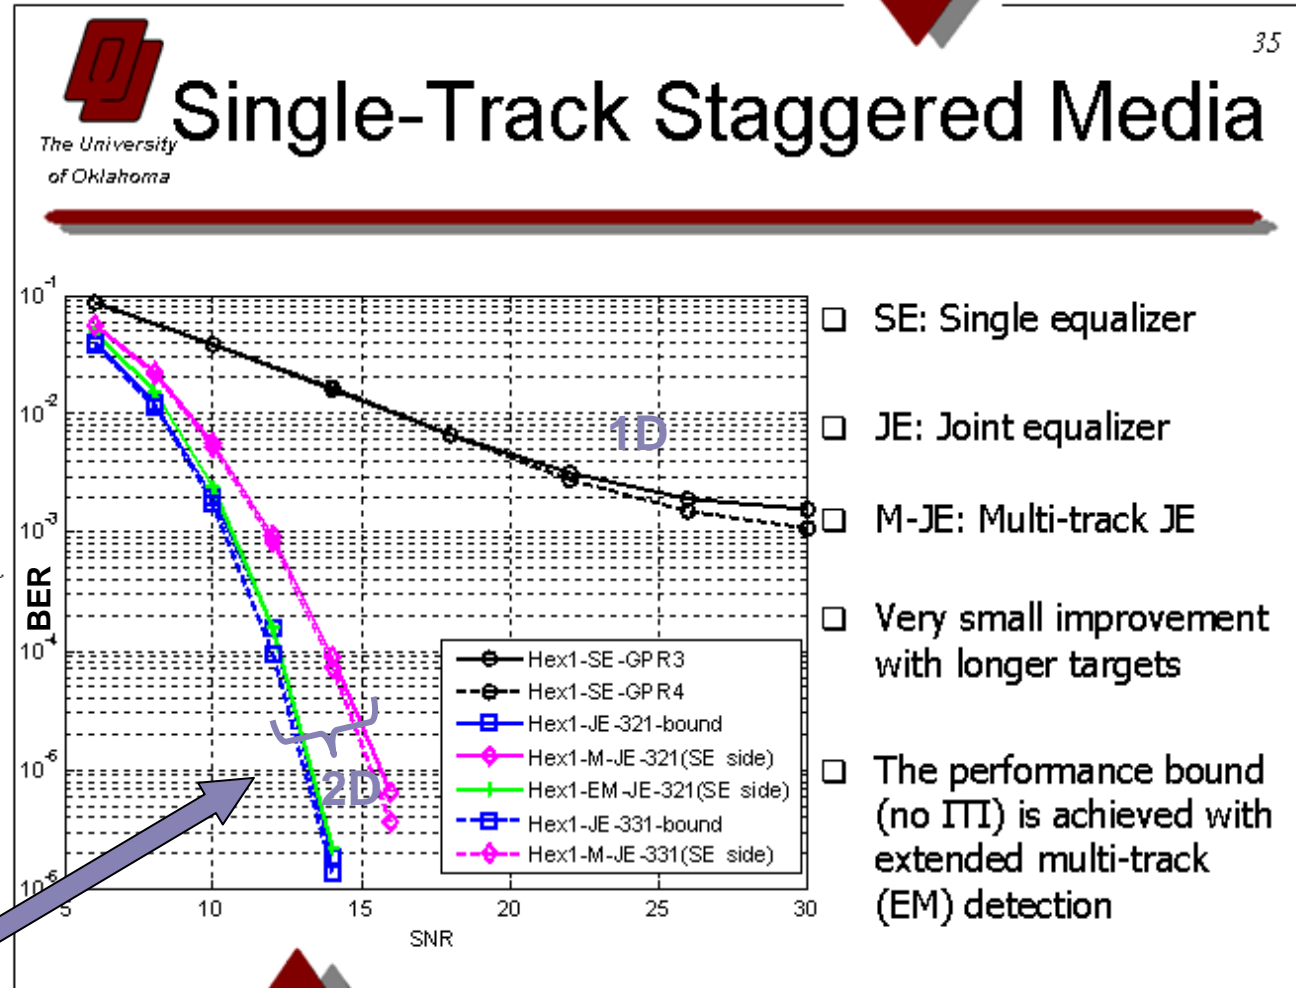
Y. Okamoto, K. Ozaki, Ehime University

Bit Patterned Media on a hexagonal array ("Staggered Media") is easiest to fabricate but has bit-aspect-ratio < 1

For practical read head designs, there will be a lot of side-reading (~30%)



The high level of Inter-Track Interference (ITI) will necessitate 2D-equalization to get good error-rates



8/4/2010

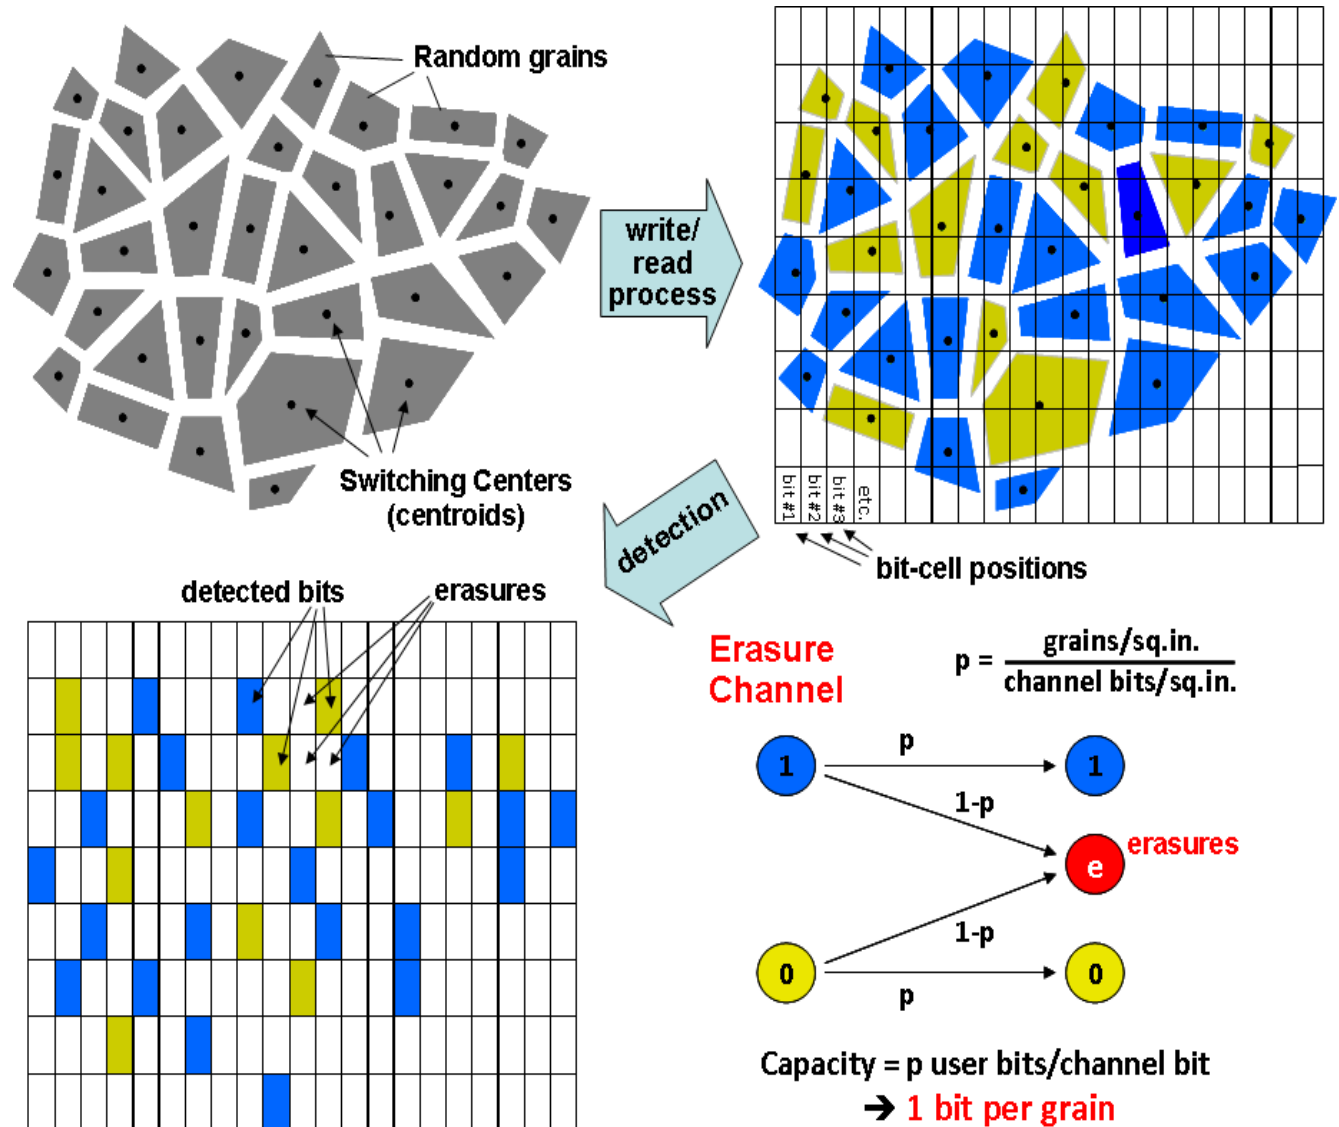
CSPLab

JR Cruz, University of Oklahoma

TDMR Central Argument: “one bit-per-grain”

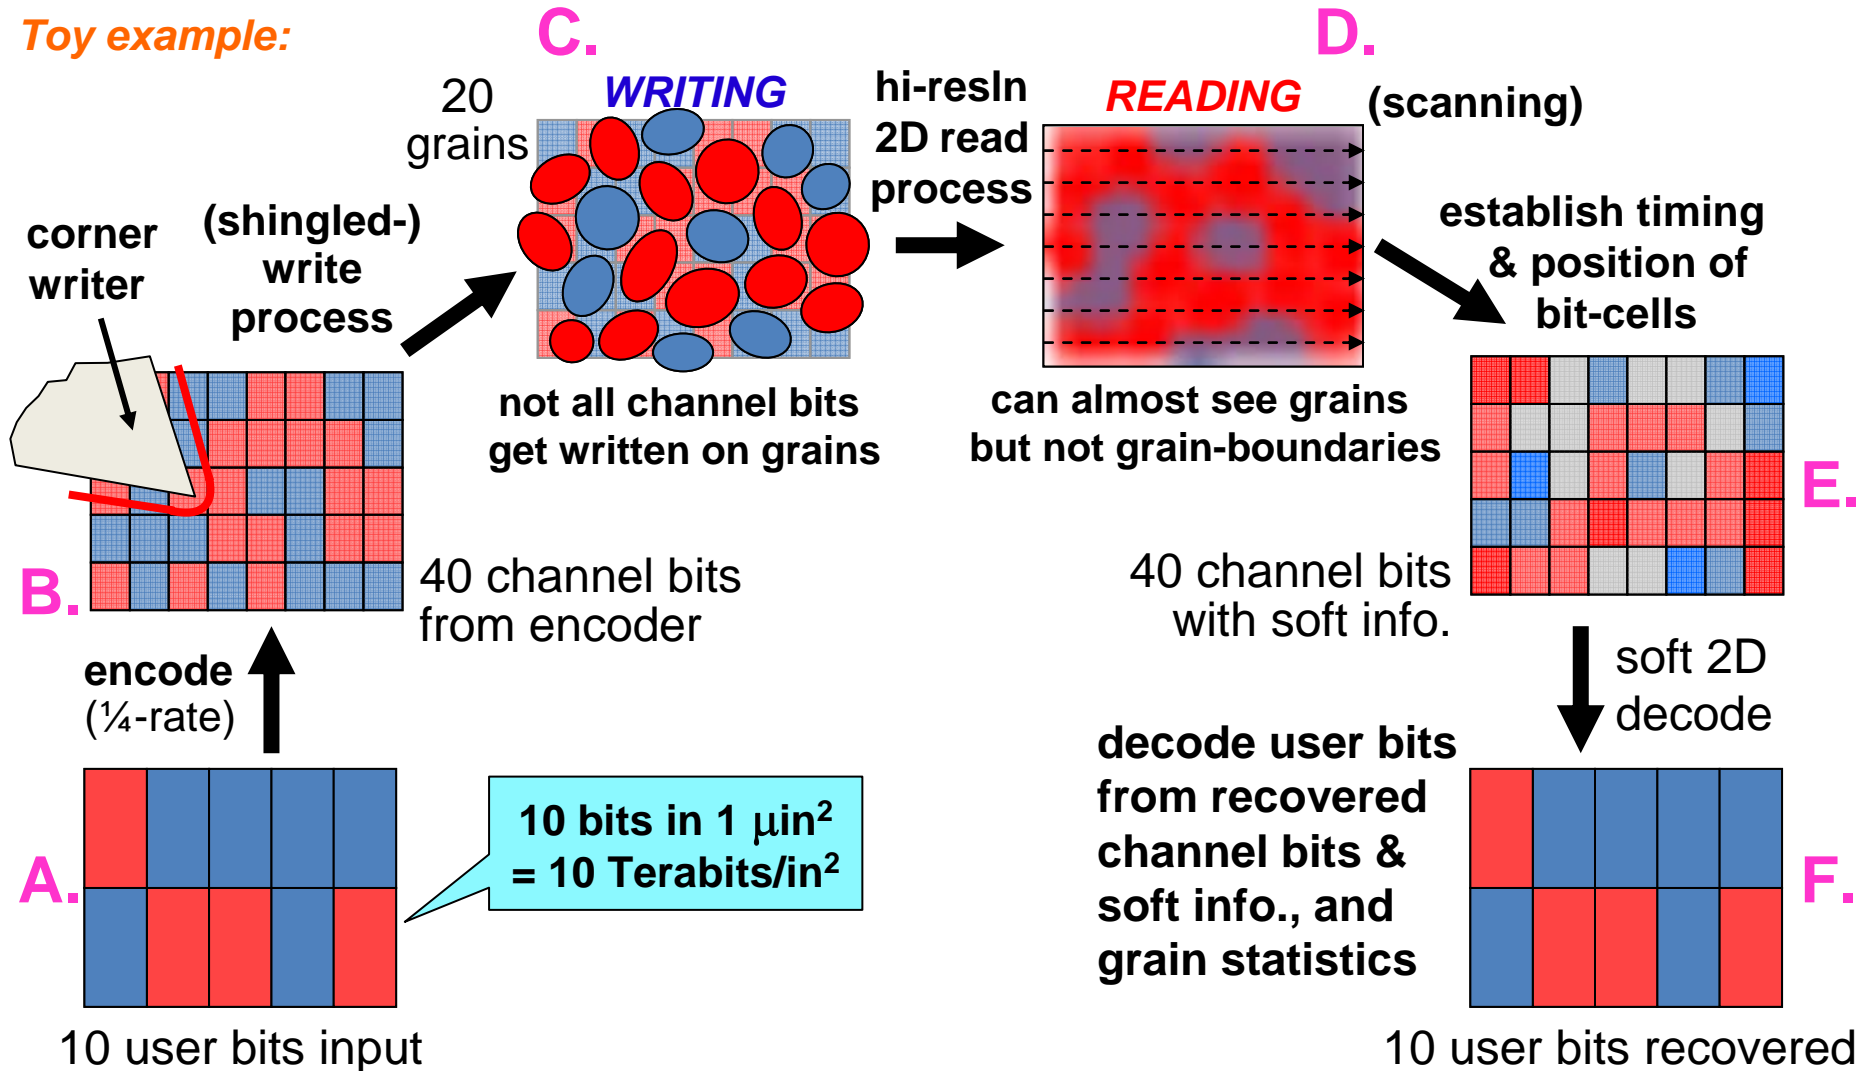
Even with random grains and no knowledge of grains during writing, it is still theoretically possible to approach 1 user bit per grain

[During readback & detection, one must be able to exactly discern the grain outlines and have complete knowledge of how the write-process works]



Concept: Two-Dimensional Magnetic Recording (TDMR)

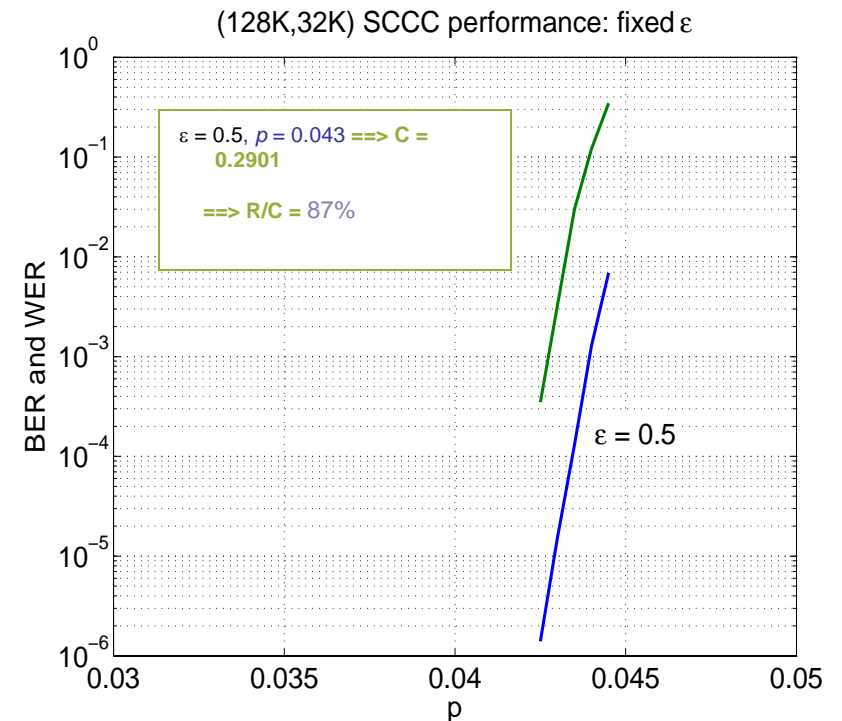
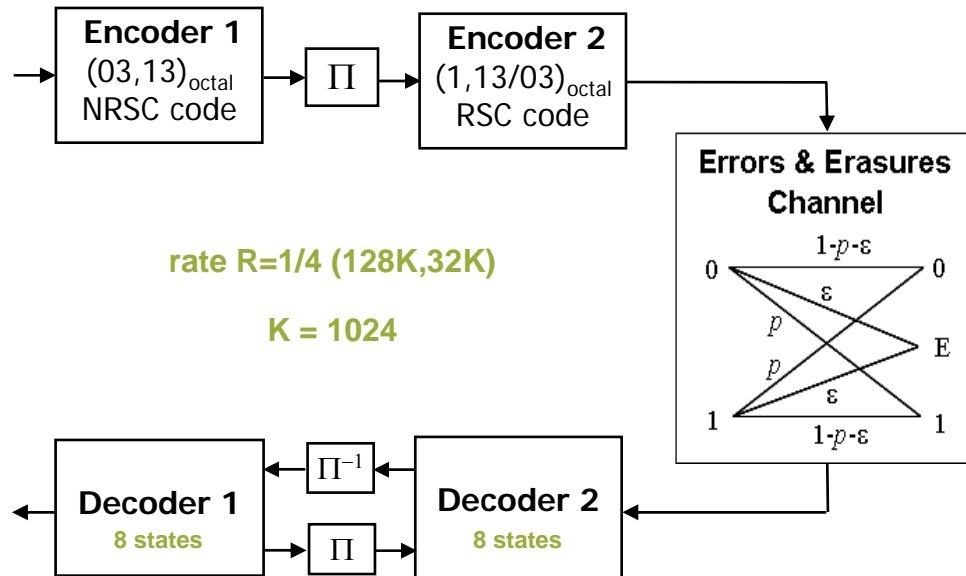
Toy example:



TMRC 2008

Coding/Detection on Erasures & Errors Channel

Serially Concatenated Convolutional Codes



Practical codes and detectors are reaching > 85% of theoretical information capacity for errors and erasures channel


ϵ = erasure rate,
 p = bit error-rate
 R = code rate = $1/4$,
 C = Shannon Capacity

W. Ryan, U. Arizona

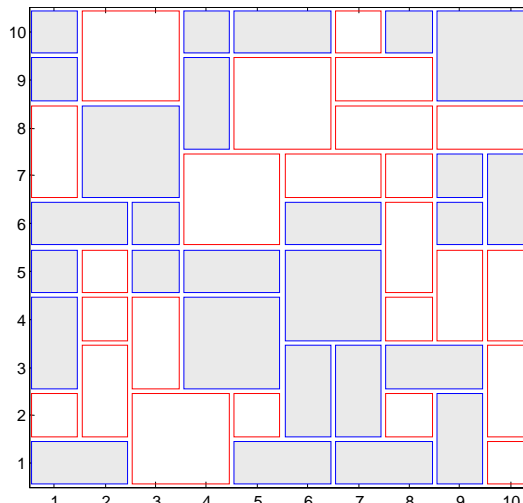
Shannon Capacity of 2D grain-limited channel?

- Assume cannot distinguish grain boundaries where no polarity change
- Create simplest possible 2D model with random grain sizes & shapes

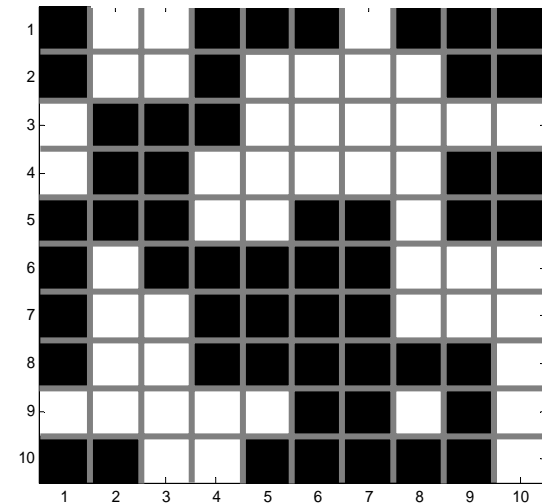
Map of input channel bits



Grain Outlines & Polarities



Map of output bits



Grains are 1x1, 1x2, 2x1, or 2x2 channel bits in size.

Mean grain area is 2 bits, sigma is 50% of mean grain

Write-process: raster scan from bottom to top, left to right

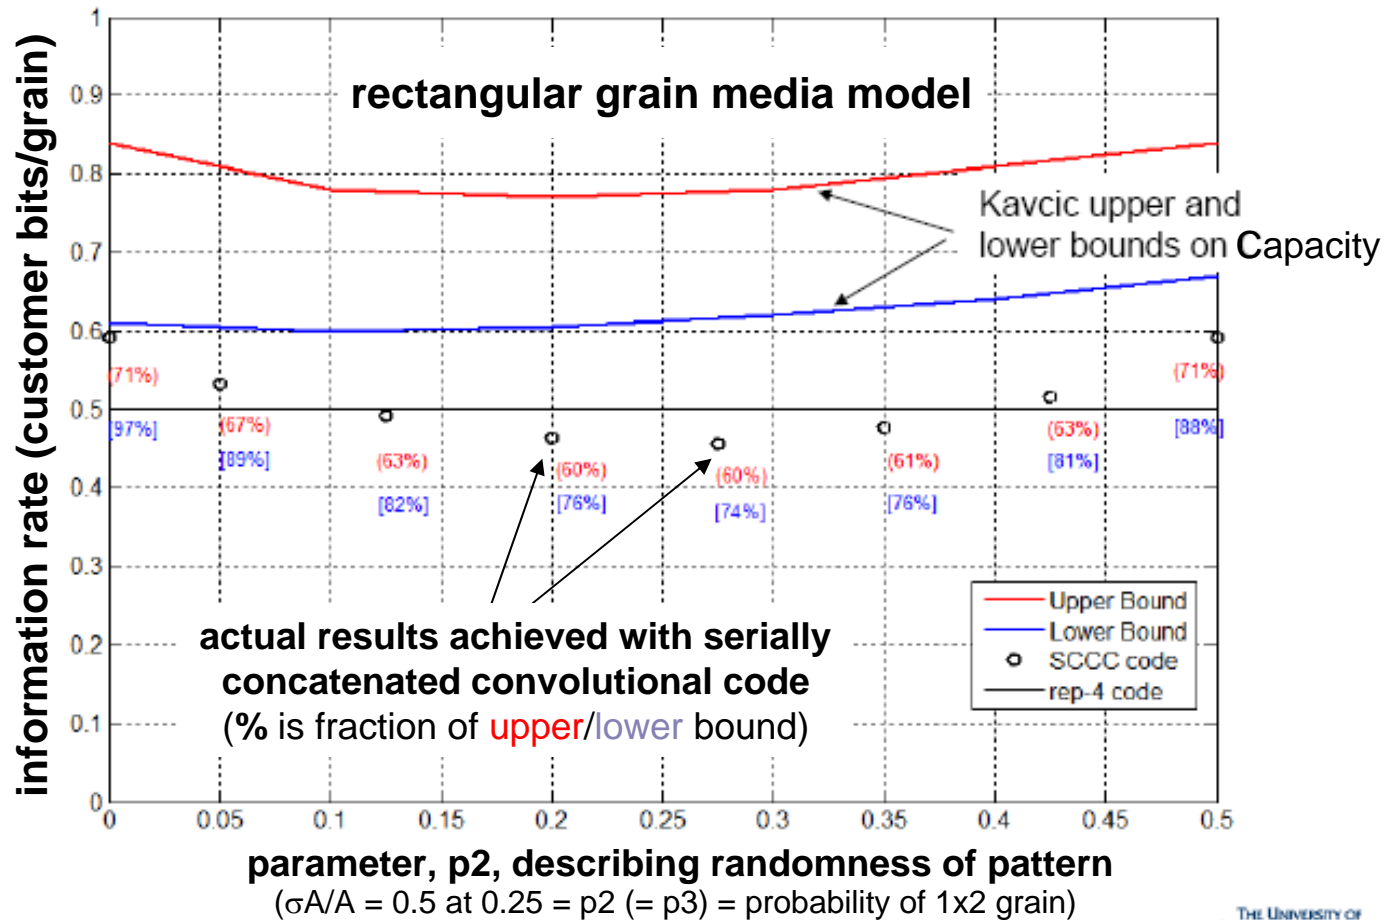
grain takes polarity of last channel bit seen (top right corner)

*establish bit-cell
positions*

For random input data, the error-rate is 0.25 which, for independent bits would give an information rate of only 0.38 bits per grain

But what is the true Shannon capacity?

Simulation Results



April 14, 2010



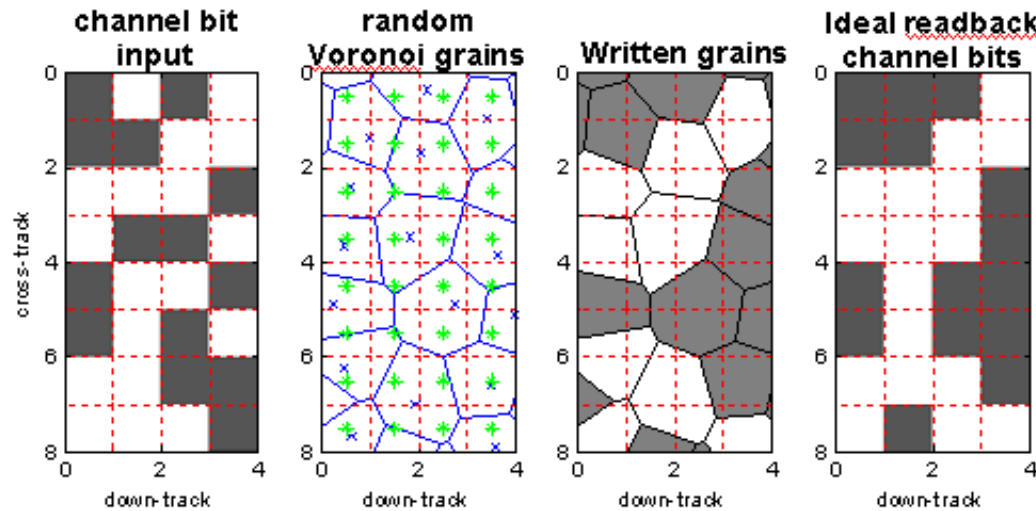
L. Pan, W. Ryan, Joint SRC/INSIC telecon, Apr. 15th, 2010

Random Voronoi Medium: Information Rates

Voronoi Media Model: Information Rates



Write-Process:
grain takes polarity of bit-cell into which its centroid falls



Only able to see polarity
Unable to distinguish grains or grain-boundaries

Illustrative example of an ideal readback TDMR channel

Lower bounds on Symmetric Information Rate (info. bits per grain) in ideal readback TDMR of half grain per channel bit model ($\kappa = 0, 1, \text{ and } 2$)

Voronoi Grains	$\{p(x_{i,j} y_{i,1}^{i,N})\}$	$\{p(x_{i,j} y_{i-1,1}^{i+1,N})\}$
$\kappa=0$ ($\sigma_A=0.488$)	0.401	0.461
$\kappa=1$ ($\sigma_A=0.352$)	0.429	0.488
$\kappa=2$ ($\sigma_A=0.249$)	0.449	0.511

Information rate in Voronoi media model can also exceed 1/2-bit per grain

“YOU WANT to BUILD WHAT?!”

INSIC Annual Meeting 2009, Symposium on Future Technologies

■ 10 TByte 2.5”x 12.5mm high, 5400 rpm

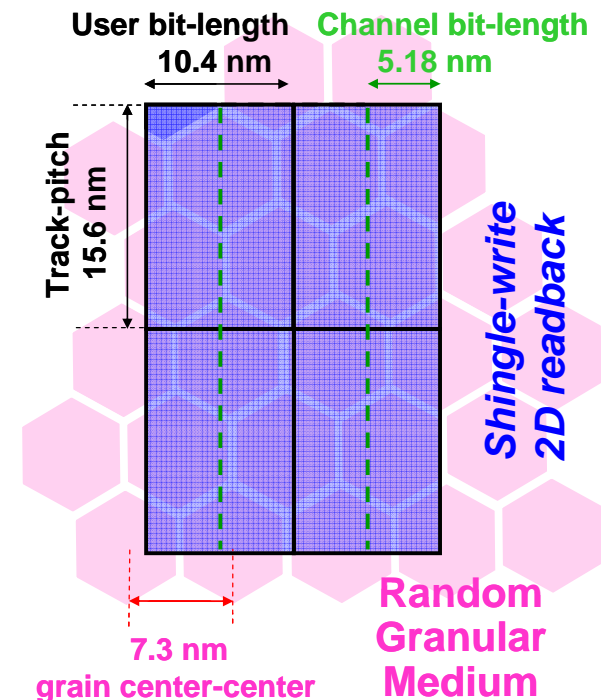
Shingled Writing & Two-Dimensional Readback

■ 4 Terabit/in² (~10x today's densities)

- 2449 kBPI (~1.5x today's BPI)
- 1633 kTPI (~6x today's TPI)
- bit-aspect ratio = 1.5 (vs. 6 today)
- ~2 Gbit/s (similar to today)
- Mag.-spacing 3 to 5 nm (vs. ~10 nm)
- Grain-size 7.3 nm (8-10 nm today)
- 3.5 Grains per bit (vs. ~20 today)
- Track-pitch = 16nm (Shingled Write)
- Read-width ~16 nm (2D Readback)
- Read-latency ~3-revs (~33 ms)



10 TB
3-disks



Medium for TDMR at 4 Terabits/sq.in

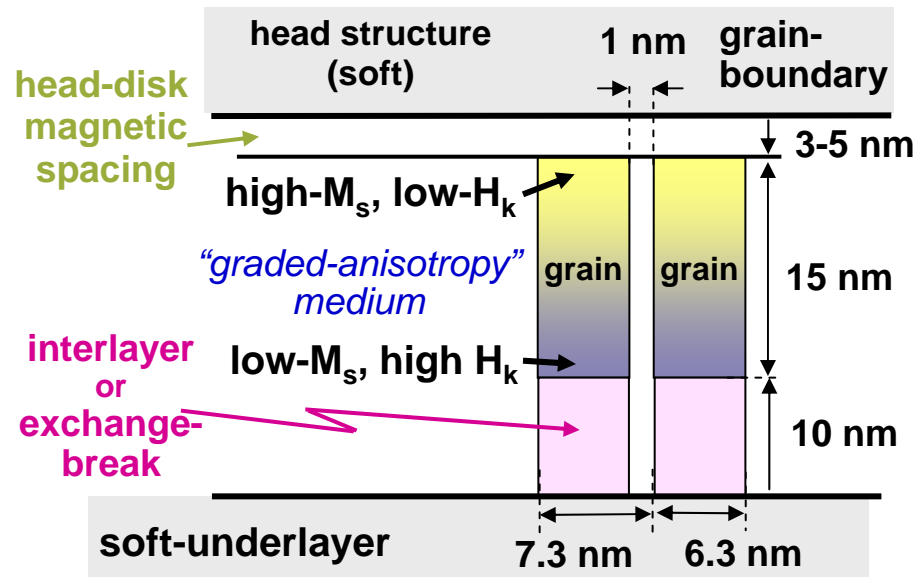
- **Need continuous improvement in conventional media:**

- some reduction in grain-size ($14 \text{ Teragrains/in}^2 = 7\text{-}8 \text{ nm}$ pitch, but need continued tightening of distributions: grain-size & shape, switching-field, exchange-field

- **Need to ensure that medium is 'writable' by the limited head field**

- engineer medium to have low anisotropy & high moment at top of the grains
(where switching process can be readily nucleated by strong head-fields & gradient)
- Bottom of grains needs to have very high anisotropy to ensure thermal stability

Min stability factor
 $K_u V / kT \approx 60$
 together with
 $H_k = 17.9 \text{ kOe}$ &
 $M_s = 445 \text{ emu/cc}$,
 dictates a grain-
 height or **medium
 thickness of 15 nm**

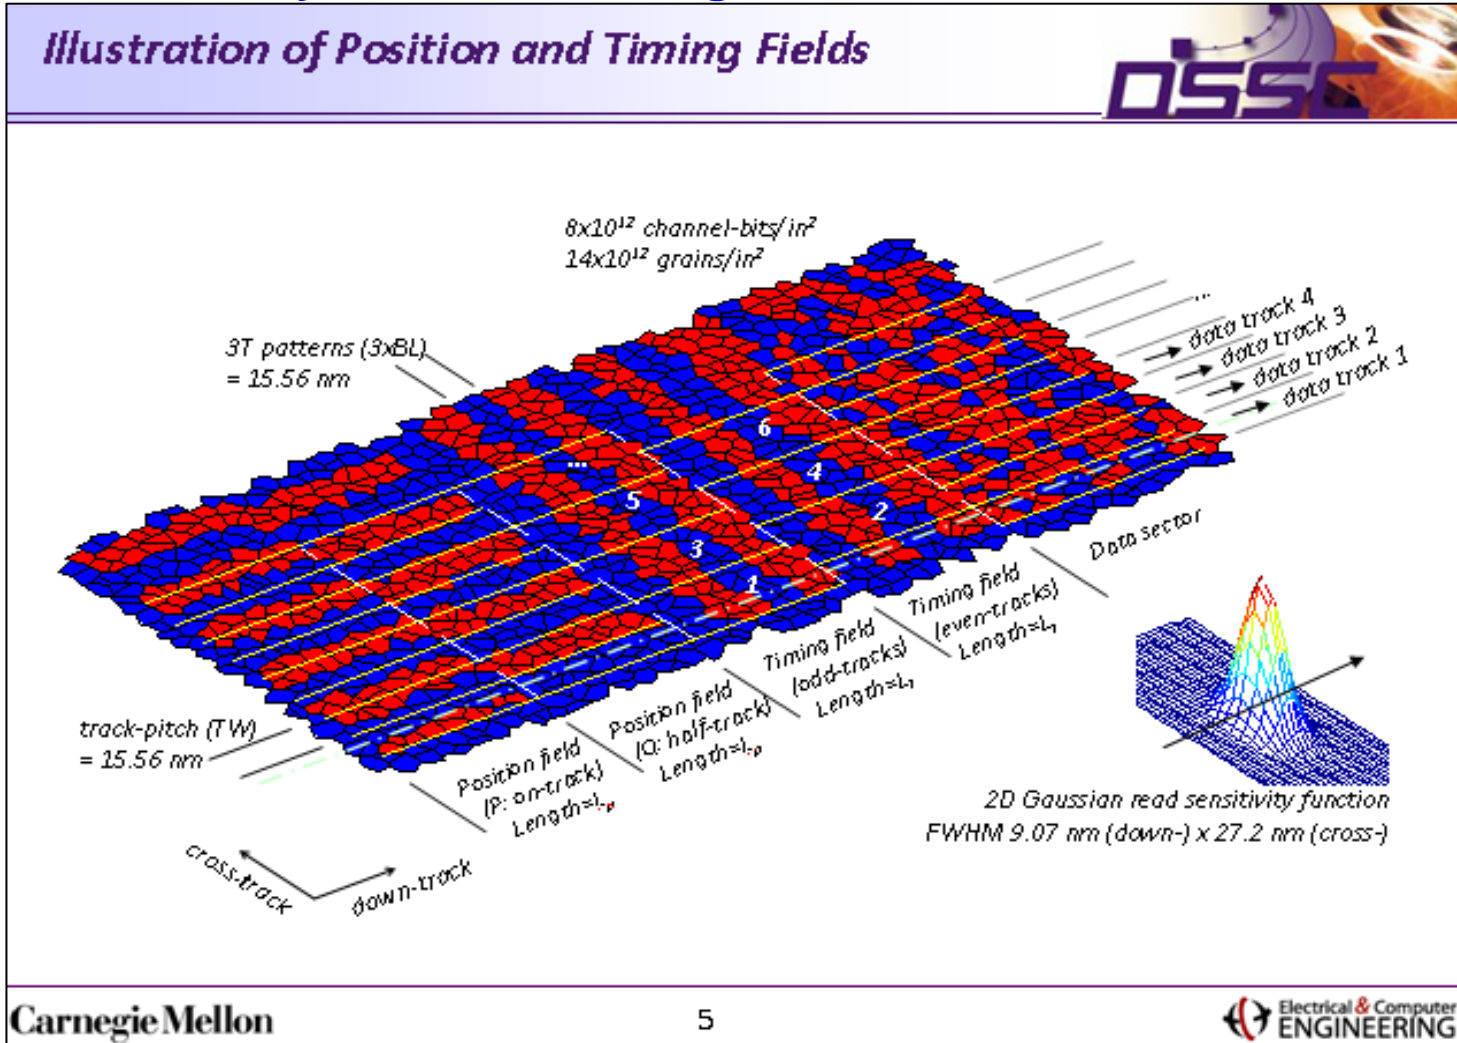


Graded-Anisotropy Medium

Assume axial switching field can be reduced to **12.5kOe** vs. 17.9 kOe required for a 'Stoner-Wohlfarth' medium of similar moment & stability (i.e. $\xi = 1.4$)

Need ~14 Teragrains/sq.in. to support 4 Tbits/sq.in.

Can Timing & Position be recovered with sufficient accuracy without costing unreasonable overhead?



E. Hwang et al., PMRC 2010

4 Tb/sq.in. - including timing & position recovery

TDMR Recording Channel Model:
 - Random Voronoi / Rectangular bit
 - 2D Gaussian Reader Sensitivity

8 Tbit/in² channel bit density

(100 % overhead for coding, timing, positioning, etc.)

15.6 nm x 5.2 nm channel bit

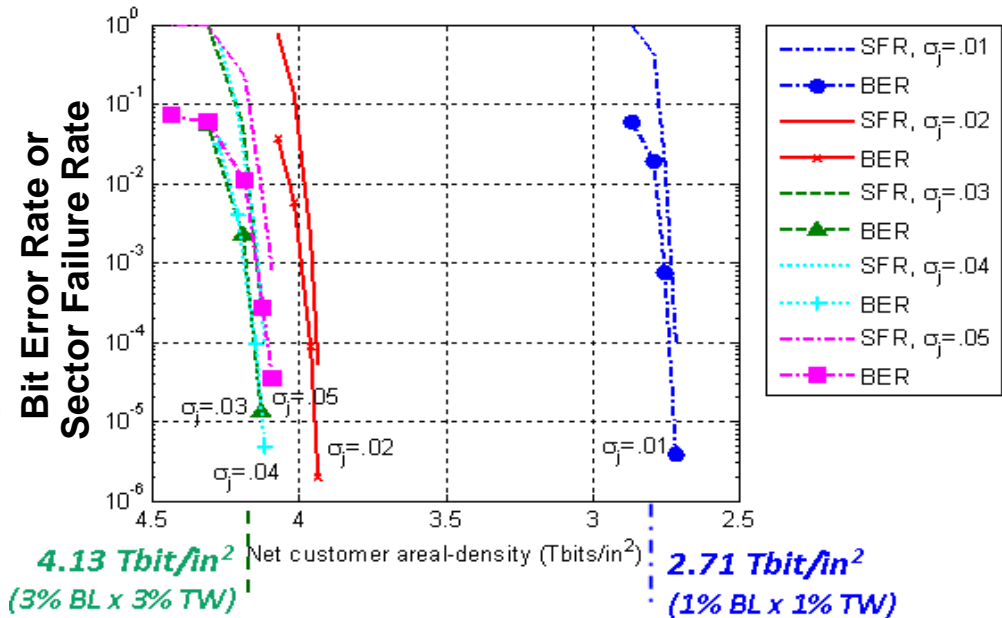
(3:1 BAR, 1632 KTPI x 4900 KBPI)

1.75 grains per channel bit

(14 Tgrains/in², D=6.79 nm, sigma-area=50%)

Target: 4 Tbit/in² customer density

Accuracy (Timing+Posn.)	1%	2%	3%	4%	5%
position/timing overhead, η_h	0.316	0.079	0.035	0.019	0.013
write error rate			7.913×10^{-2}		
raw BER	2.281×10^{-1}	2.283×10^{-1}	2.285×10^{-1}	2.289×10^{-1}	2.295×10^{-1}
equalized BER	1.032×10^{-1}	1.044×10^{-1}	1.063×10^{-1}	1.090×10^{-1}	1.121×10^{-1}
coding overhead, η_c	0.453	0.457	0.461	0.473	0.481
→ corrected BER	3.834×10^{-6}	1.948×10^{-6}	1.314×10^{-5}	4.766×10^{-6}	3.434×10^{-5}
net customer areal-density	2.714 Tbit/in ²	3.931 Tbit/in ²	4.131 Tbit/in²	4.118 Tbit/in ²	4.093 Tbit/in ²



System achieves > 4Tbit/sq.in. on 14 Tgrains/sq.in. media (0.29 customer bits/grain) Including overhead for error-control and timing and position recovery!

E. Hwang et al., PMRC 2010

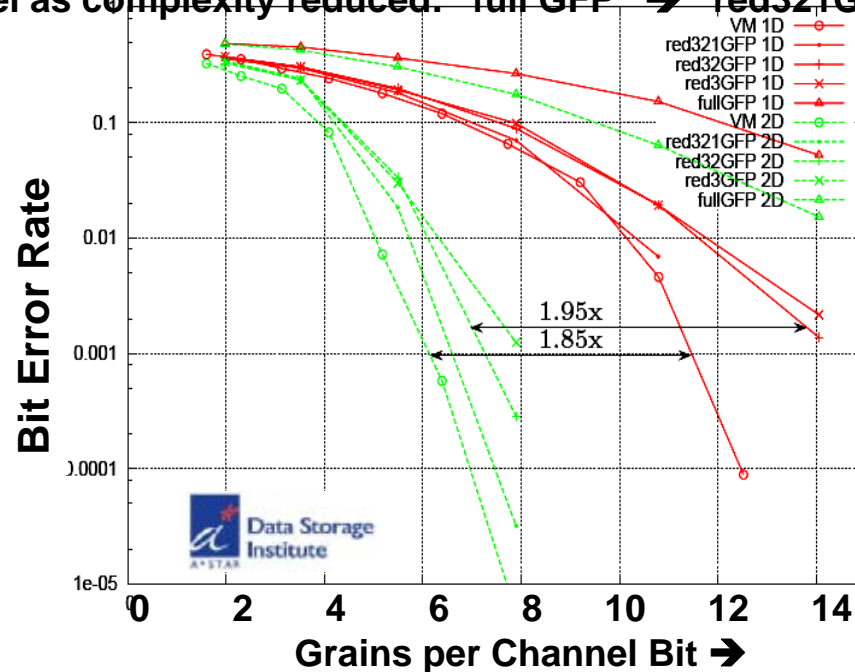
More Realistic Write Process and Bit-Aspect Ratios

- “GFP” model has medium with grain-pitch = 6.5 nm and $\sigma A/A = 17\%$
- Readback is with a 2D Gaussian sensitivity function with $T50 = 24 \times 8$ nm
- VM is simple Voronoi model (grains written by centroid in rectangular bit)
(GFP model becomes similar to VM model as complexity reduced: “full GFP” → “red321GFP”)

Raw detector bit error-rates from the “GFP” model (no parity coding)

Red curves denote 1D-BCJR,
Green curves denote a 2D-BCJR

BAR is held constant at 2.7 while bit-area is varied along horizontal axis



- ~2x Areal-Density increase from 2D vs. 1D signal-processing
- Significant degradation from increasingly realistic write process

Future Technology Options & Limits for Hard Disk Drives


- Limits on Magnetic Recording
- Technology options for 1 Tbit/sq.in & beyond
- Shingled Magnetic Recording (SMR)
- Two-Dimensional Magnetic Recording (TDMR)
- ➔ ■ Data architecture and Systems Issues
- Future Scenarios

The 26th IEEE Conference on Mass Storage Systems and Technologies (MSST2010): Research Track

Hyatt Regency at Lake Tahoe, Incline Village, Nevada, May 3-7th, 2010

Industry must work closely with customers to understand & address changes in performance characteristics

Growing Interest in Storage Architectures for Shingled recording



Friday:

9:00-10:30

Energy Efficiency and QoS

Leveraging Disk Drive Acoustic Modes for Power Management.
D. Chen, G. Goldberg, R. Kahn, R. Kat, K. Meth

Energy and Thermal Aware Buffer Cache Replacement Algorithm.
J. Yue, Y. Zhu, Z. Cai, L. Lin

Mahanaxar: Quality of Service Guarantees in High-Bandwidth, Real-Time Streaming Data Storage.
D. Bigelow, S. Brandt, J. Bent, H. Chen

10:30-11:00

Break

11:00-12:30

Disk and Memory Recording Technology

A Content-Aware Block Placement Algorithm for Reducing PRAM Storage Bit Writes.
B. Wongchaowart, M. Iskander, S. Cho

Design Issues for a Shingled Write Disk System.
A. Amer, D. Long, E. Miller, J. Paris, T. Schwarz

Indirection Systems for Shingled-Recording Disk Drives.
Y. Cassuto, M. Sanvido, C. Guyot, D. Hall, Z. Bandic

12:30-14:00

Lunch

SCU, UCSC



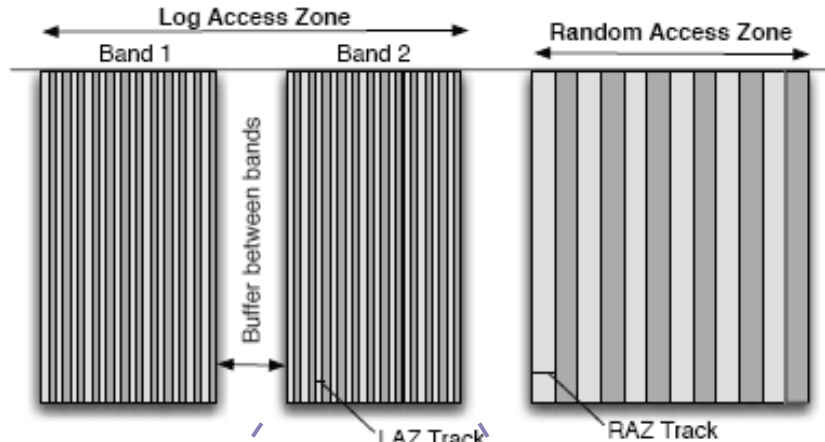
Hitachi GST



Layout

Log Access Zones
(tightly-shingled)

track-pitch
~ = $\frac{1}{2}$ write-width

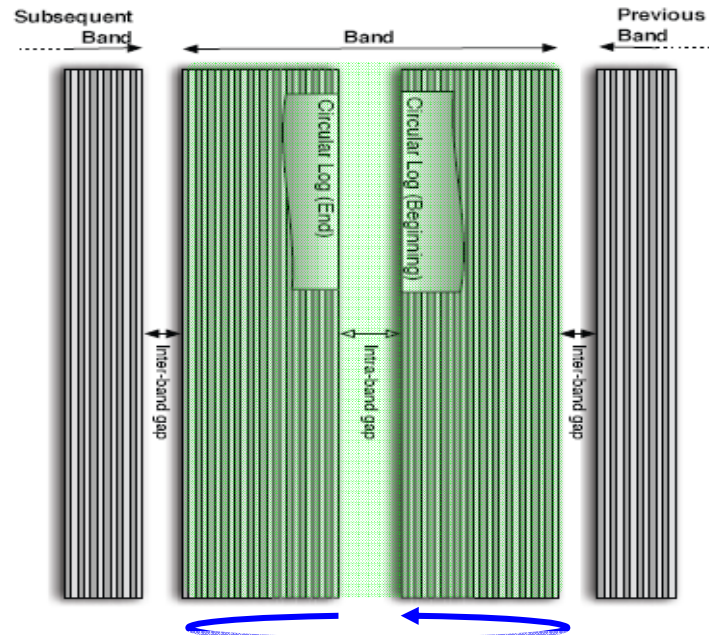


Random access zone
(non-shingled)

track-pitch
~ = write-width

(fast access/settle times)

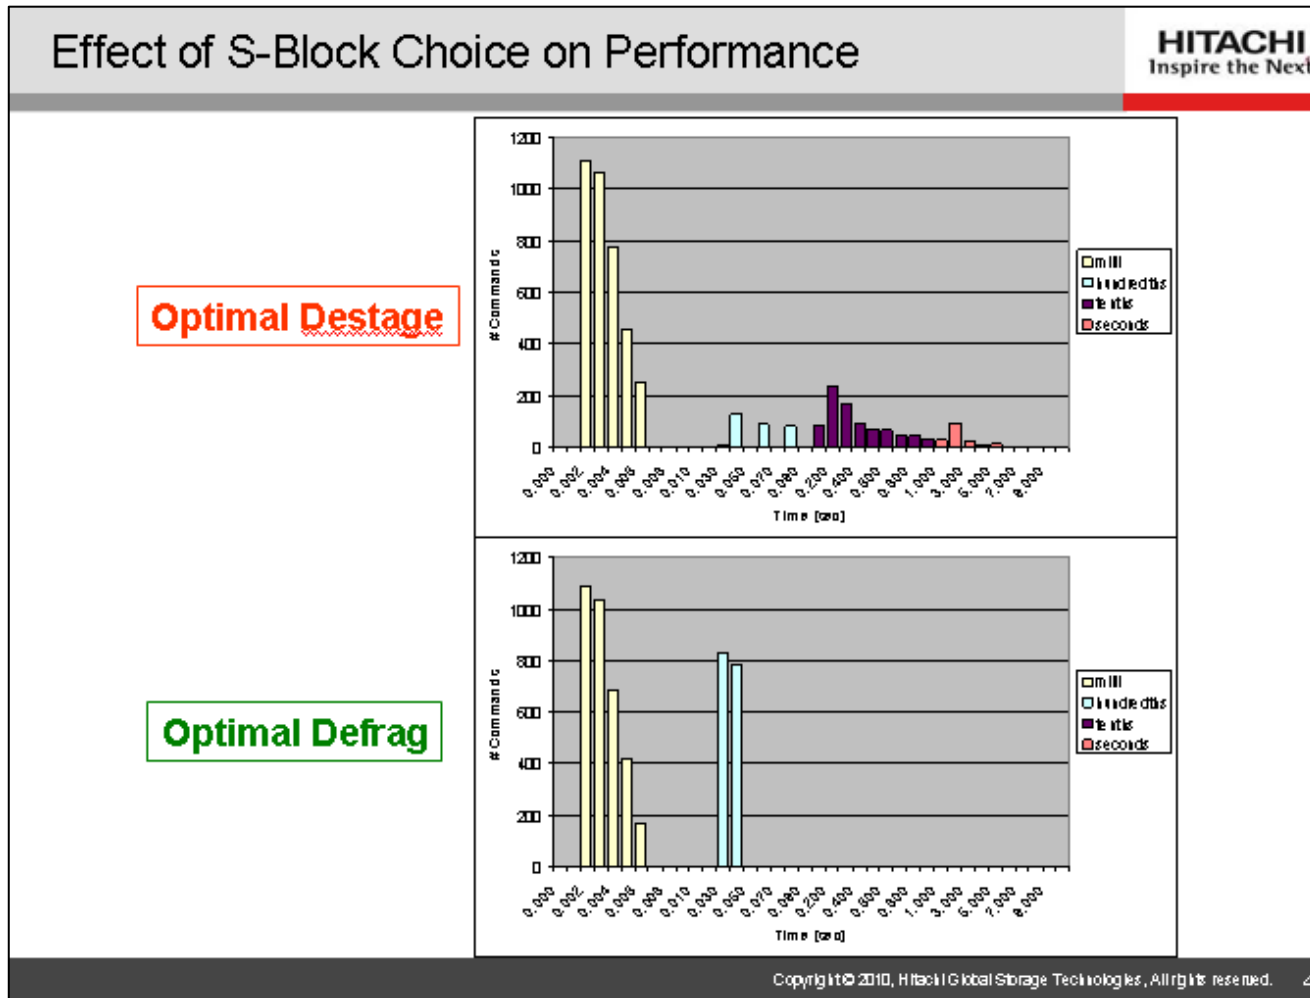
Shingled Zones are
used as circular
log-structured files



Most of data is
stored at much
higher TPI in these
circular log files

A. Amer et al. "Design Issues for a Shingled Write Disk System" MSST 2010

Tune for best *average* performance or best *worst-case* performance?



Y. Cassuto et al. "Indirection Systems for Shingled-Recording Disk Drives" MSST 2010

Garth Gibson, InterMag, April 2009, Paper FA-06

Summary

Shingled-written disk is N bands of sequentially written sectors, each of order GB

Disk can still offer normal commands, write speed using “translation layer” embedded code

Take Flash SSD FTL as starting point

Flash-inspired TRIM command helps

TDMR reading a bigger problem

3-5 revs per small read hard to hide

This could reduce market acceptance

Carnegie Mellon
Parallel Data Laboratory

www.pdl.cmu.edu


14

Garth Gibson, 05/07/2009

(We really need that multi-element reader!)

Future Technology Options & Limits for Hard Disk Drives

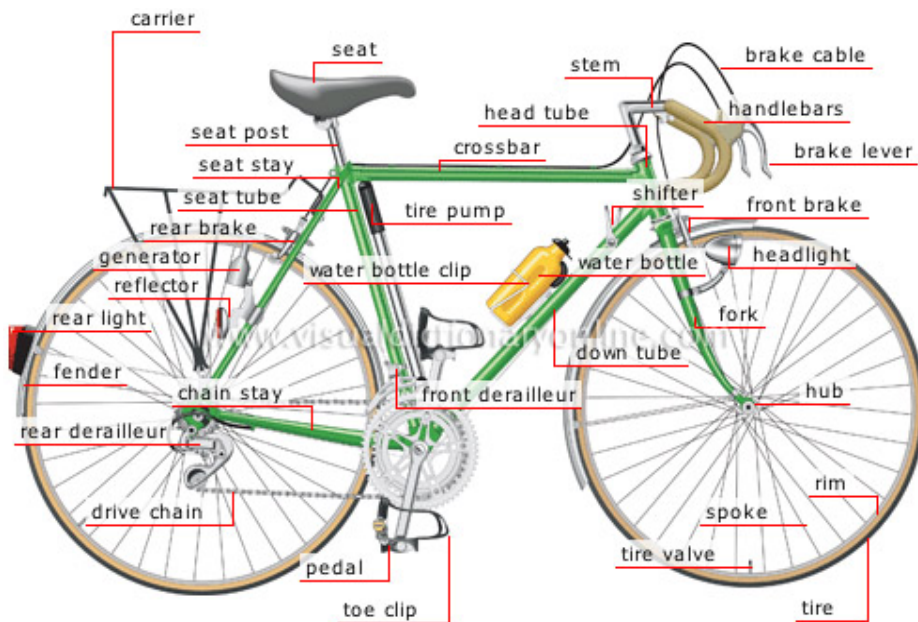
- Limits on Magnetic Recording
- Technology options for 1 Tbit/sq.in & beyond
- Shingled Magnetic Recording (SMR)
- Two-Dimensional Magnetic Recording (TDMR)
- Data architecture and Systems Issues



- Future Scenarios

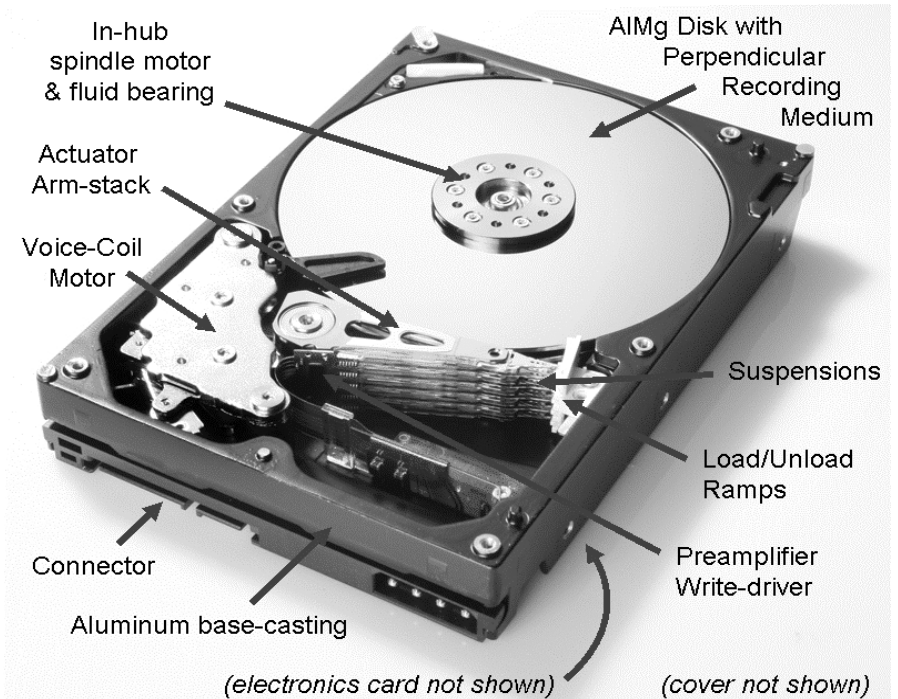
“Perfect” Inventions

Bicycle



- two large similar-diameter wheels
- front-wheel pivots with handlebars
- rear wheel driven through sprocket & chain from rotating pedals
- operated in seated position

Hard Disk Drive

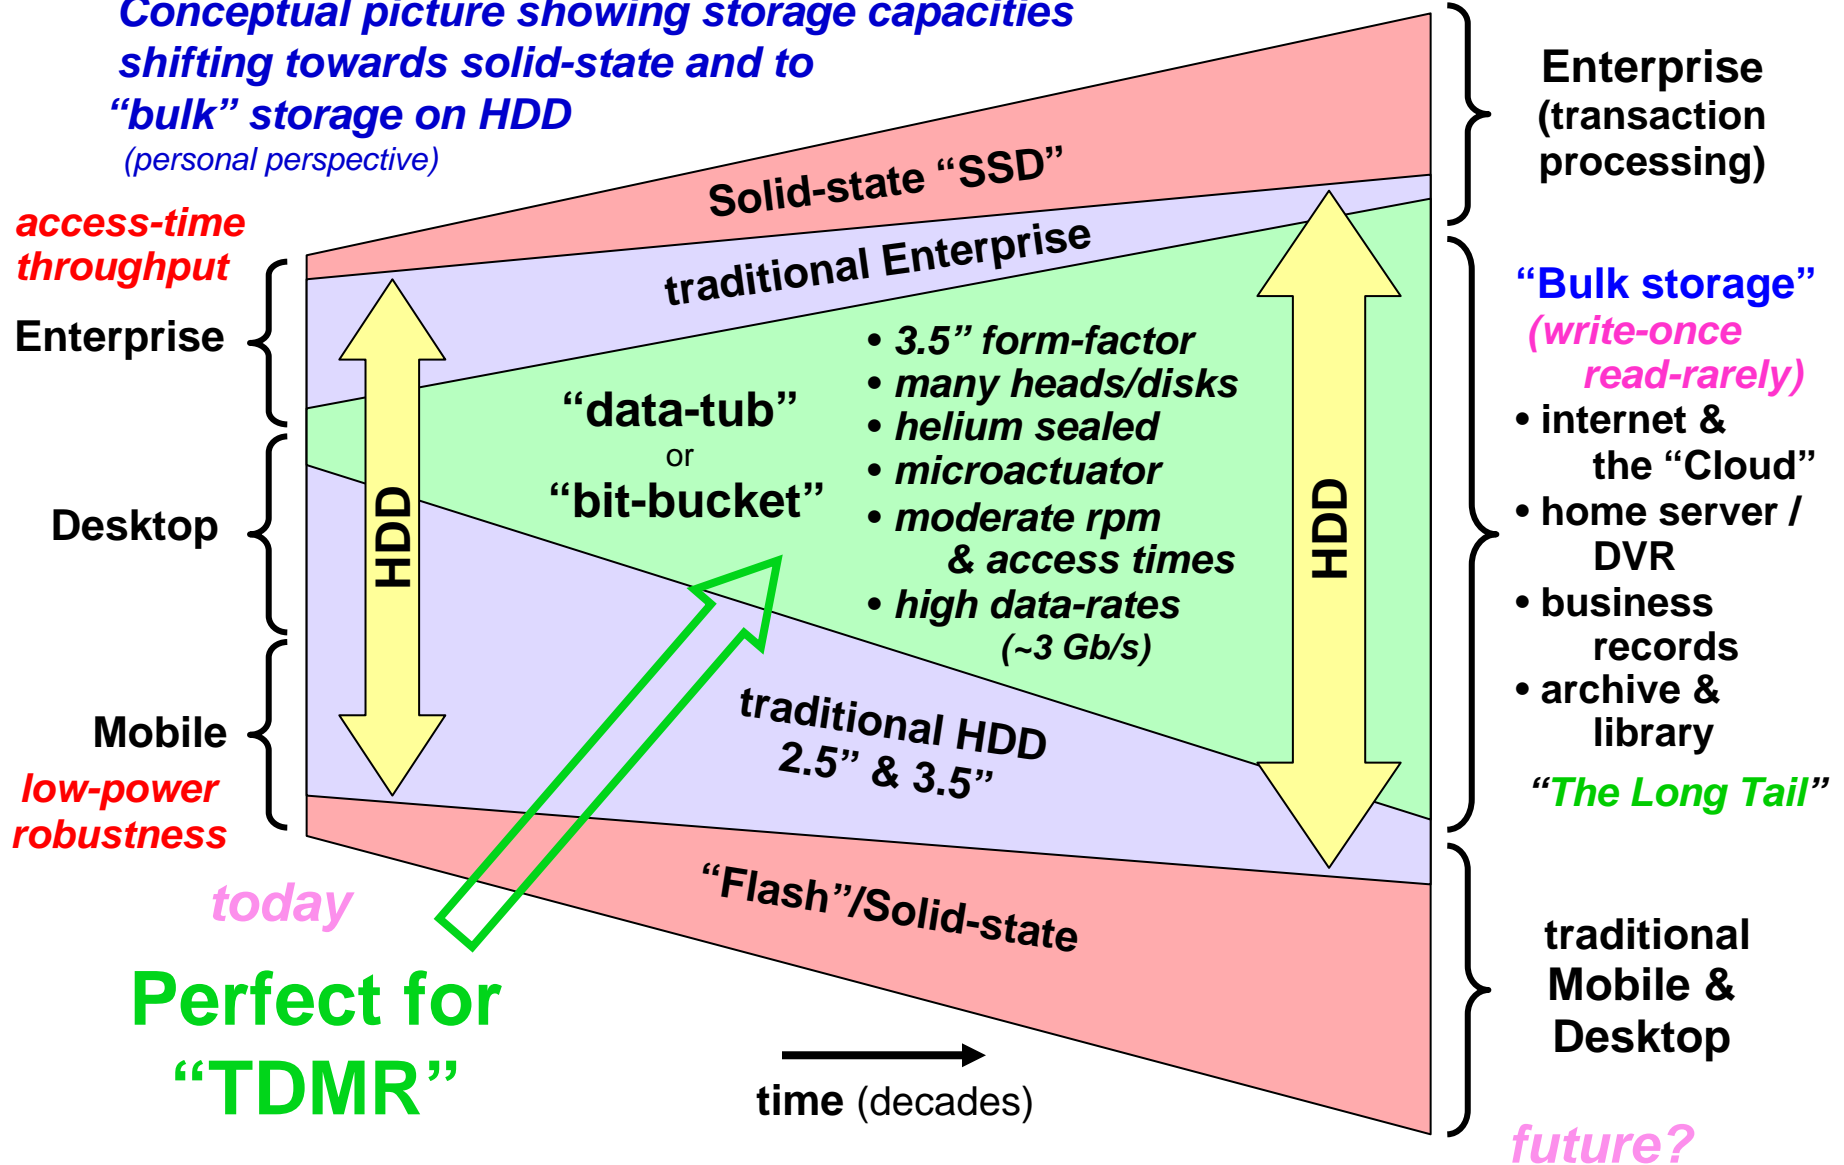


- fluid-bearing spindle with multiple disks
- rotary actuator carrying multiple heads
- slider with self-generated air-bearing
- thermal μ -actuator for magnetic spacing
- perpendicular recording mode

• **Hard Disk drives will be here for many decades to come**

Evolving Markets

Conceptual picture showing storage capacities shifting towards solid-state and to “bulk” storage on HDD
(personal perspective)



■ Limits on Magnetic Recording

- Projections of 1 Tbit/sq.in. & ~3 Gbit/s still seem well founded

■ Technology options for 1 Tbit/sq.in & beyond

- BPMR - lower technology risk but greater manufacturing challenges & cost
- HAMR - higher risks for technology & reliability but more compatible with current HDD architecture & manufacturing processes
- SMR/TDMR - low risk but major customer acceptance issues (especially for TDMR)

■ Shingled Magnetic Recording (SMR)

- Gaining acceptance as at least an interim solution (early-implementation of BPMR and HAMR looks increasingly challenging)

■ Two-Dimensional Magnetic Recording (TDMR)

- Still much work to do to understand & minimize large gap between Shannon capacity promising >0.5 bits per grain and realistic write-processes giving ~ 0.1 bits/grain

■ Data architecture and Systems Issues

- Attention now being drawn to this topic - work starting in both academia & industry

■ Future Scenarios

- HDD will be eased out of traditional markets, but the “bulk” data storage market will become immense and be ruled by HDD data tubs (using TDMR, of course!)



Call for nominations for Invited Speakers & “Symposiums”

- Propose Symposium on Shingled- and 2D Magnetic Recording ?
 - ~6 papers covering key areas:
 - *Shingled Writing*: heads, media, measurements, system integration
 - *2D readback/TDMR*: multi-readers, ITI mitigation, 2D-detection
 - *Data handling*: Architecture, performance, interface, customer impact

Closes on Oct 28th: let me know, please, if you have ideas or suggestions for topics or speakers (roger.wood@hgst.com)

back-up

Tuesday, October 19th, 2010

Western Digital, 1710 Automation Parkway, San Jose, CA

[Directions and Map](#)

Cookies, Conversation & Pizza too at 7:00 P.M.

Presentation at 7:30 P.M.

Shingled Magnetic Recording and Two-Dimensional Magnetic Recording

Roger Wood

Hitachi Global Storage Technologies

Abstract

Magnetic recording as we know it today is approaching its theoretical limit of about 1 Terabit/in². Annual increases in HDD capacity will come to a halt unless a significant new technology can be introduced quickly. There are several new technology options, but all are fraught with serious practical difficulties. The presentation will briefly review the options including Bit-Patterned Magnetic Recording (BPMP), Heat-Assisted Magnetic Recording (HAMR), and Microwave-Assisted Magnetic Recording (MAMR), plus the new concepts of Shingled Magnetic Recording (SMR) and Two-Dimensional Magnetic Recording (TDMR). The focus of the talk will be on the latter concepts, SMR and TDMR. These approaches rely more on evolutionary developments of heads and media, though requiring major changes in data-handling and signal-processing.

The presentation will review recent progress on SMR and TDMR. Numerous papers have been published both on shingled-writing and on 2D signal-processing. Research on shingled-writing has explored novel head designs and real spin-stand and HDD-level measurements are now being reported. Asymmetric "Corner" heads with narrow side-gaps have been fabricated and shown to have superior performance producing narrower less-noisy erase-bands between tracks. Compared to conventional recording, spin-stand measurements of shingled tracks invariably reach a much tighter track-pitch before failure. HDD-level measurements have shown large gains in areal-density in the presence of realistic servo track-misregistration. At very high track-densities there is considerable inter-track interference that will require 2D signal-processing techniques to overcome. Much more work has been done in this area including further confirmation of the earlier claims about the inherent Shannon capacity of random granular media as well as answering basic concerns about the feasibility of timing and position recovery with such extreme levels of media noise. The Achilles heel of the technology is the need to change the way data is handled in the drive and the consequences this may have for the customer. Two recent papers have started to address this issue and an IDEMA standards committee has been established to co-ordinate industry and customer interests.

Biography

Roger Wood



Roger Wood hails originally from the UK and holds degrees from London University and the University of British Columbia. He is currently with Hitachi GST in San Jose. Dr. Wood has a long history in the Magnetic Recording industry starting at Ampex Corporation in 1979 then moving to IBM in 1986. In 1996, he spent a year at the Data Storage Institute in Singapore. In 2003, the IBM HDD operation became part of Hitachi GST and subsequently, in 2003-4, he was fortunate to enjoy an assignment in Odawara, Japan. At Ampex, Dr. Wood was the inspiration and driving force behind the introduction of the first PRML channel. More recently, at Hitachi GST, he led the advanced development effort for perpendicular recording and was delighted by the string of successful products that resulted. Dr. Wood is perhaps best known for dire predictions about conventional magnetic recording being limited to an areal density of about 1 Terabit/in² – a prediction now uncomfortably close. Dr. Wood's interests include magnetism, signal-processing, and mechanical-dynamics. He is the author of over 70 journal papers and 12 patents and is often invited to speak at conferences and technical meetings. Dr. Wood is an IEEE Fellow and was the recipient of the Magnetics Society Achievement Award for 2009.