magnetic information storage

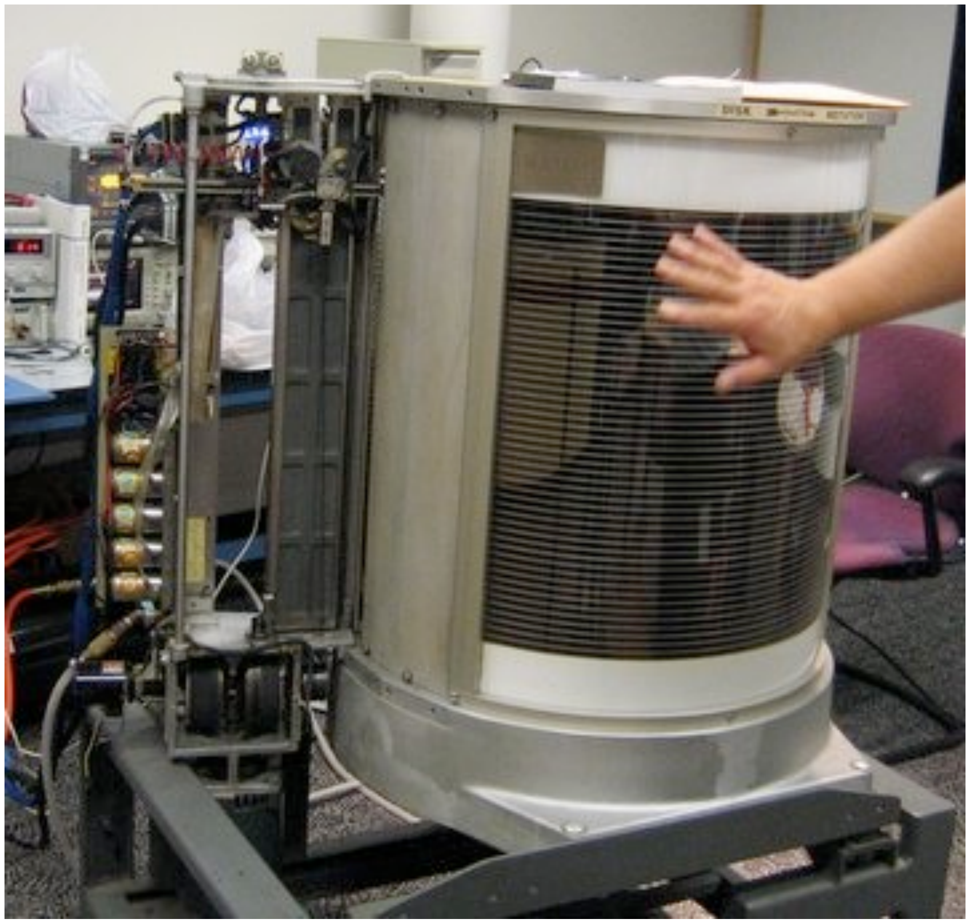

IBM 350 RAMAC, the first hard disk

it stored about 4.4Mb

wikipedia.org - “RAMAC”

what do I mean by magnetic recording?

what do I mean by magnetic recording?

what do I mean by magnetic recording?

hard disks.

mostly hard disks.

why do we use hard disks?

what is their role in a computer?

benefits?

disadvantages?

and the real reason .. $$$

basic PC architecture

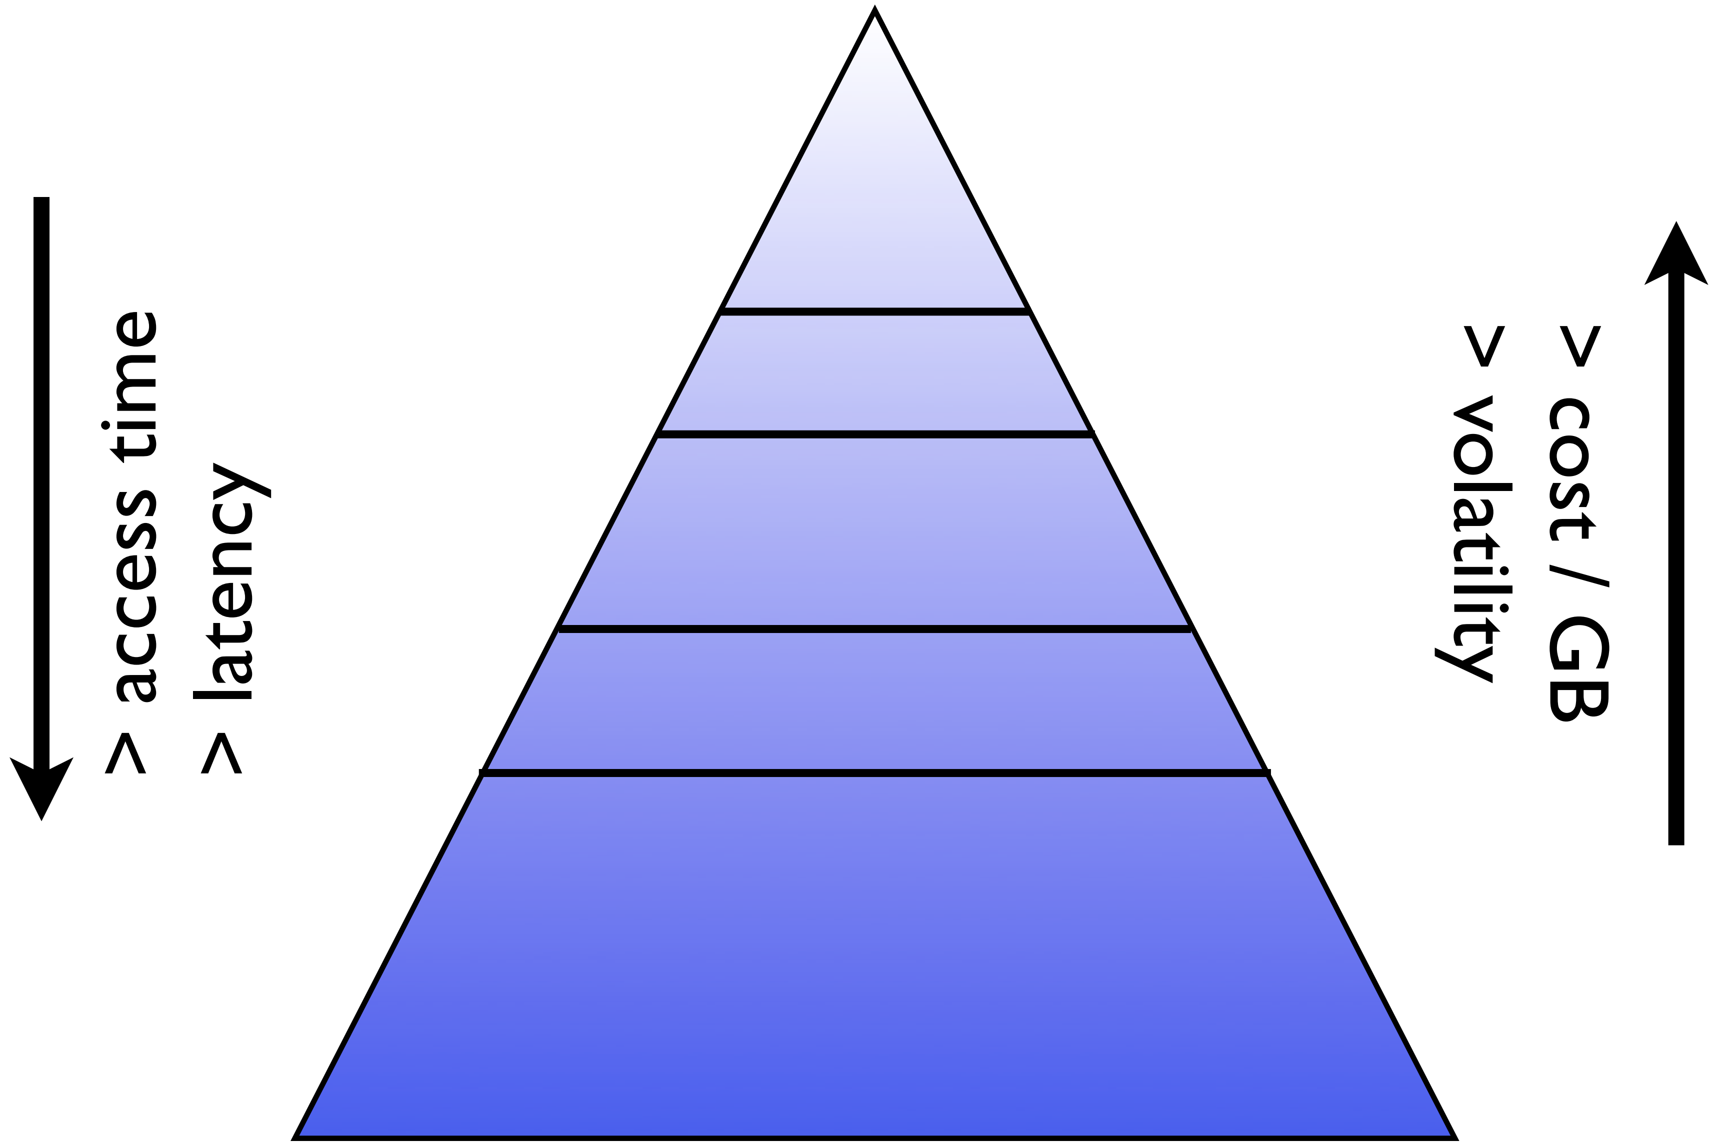

basic PC architecture

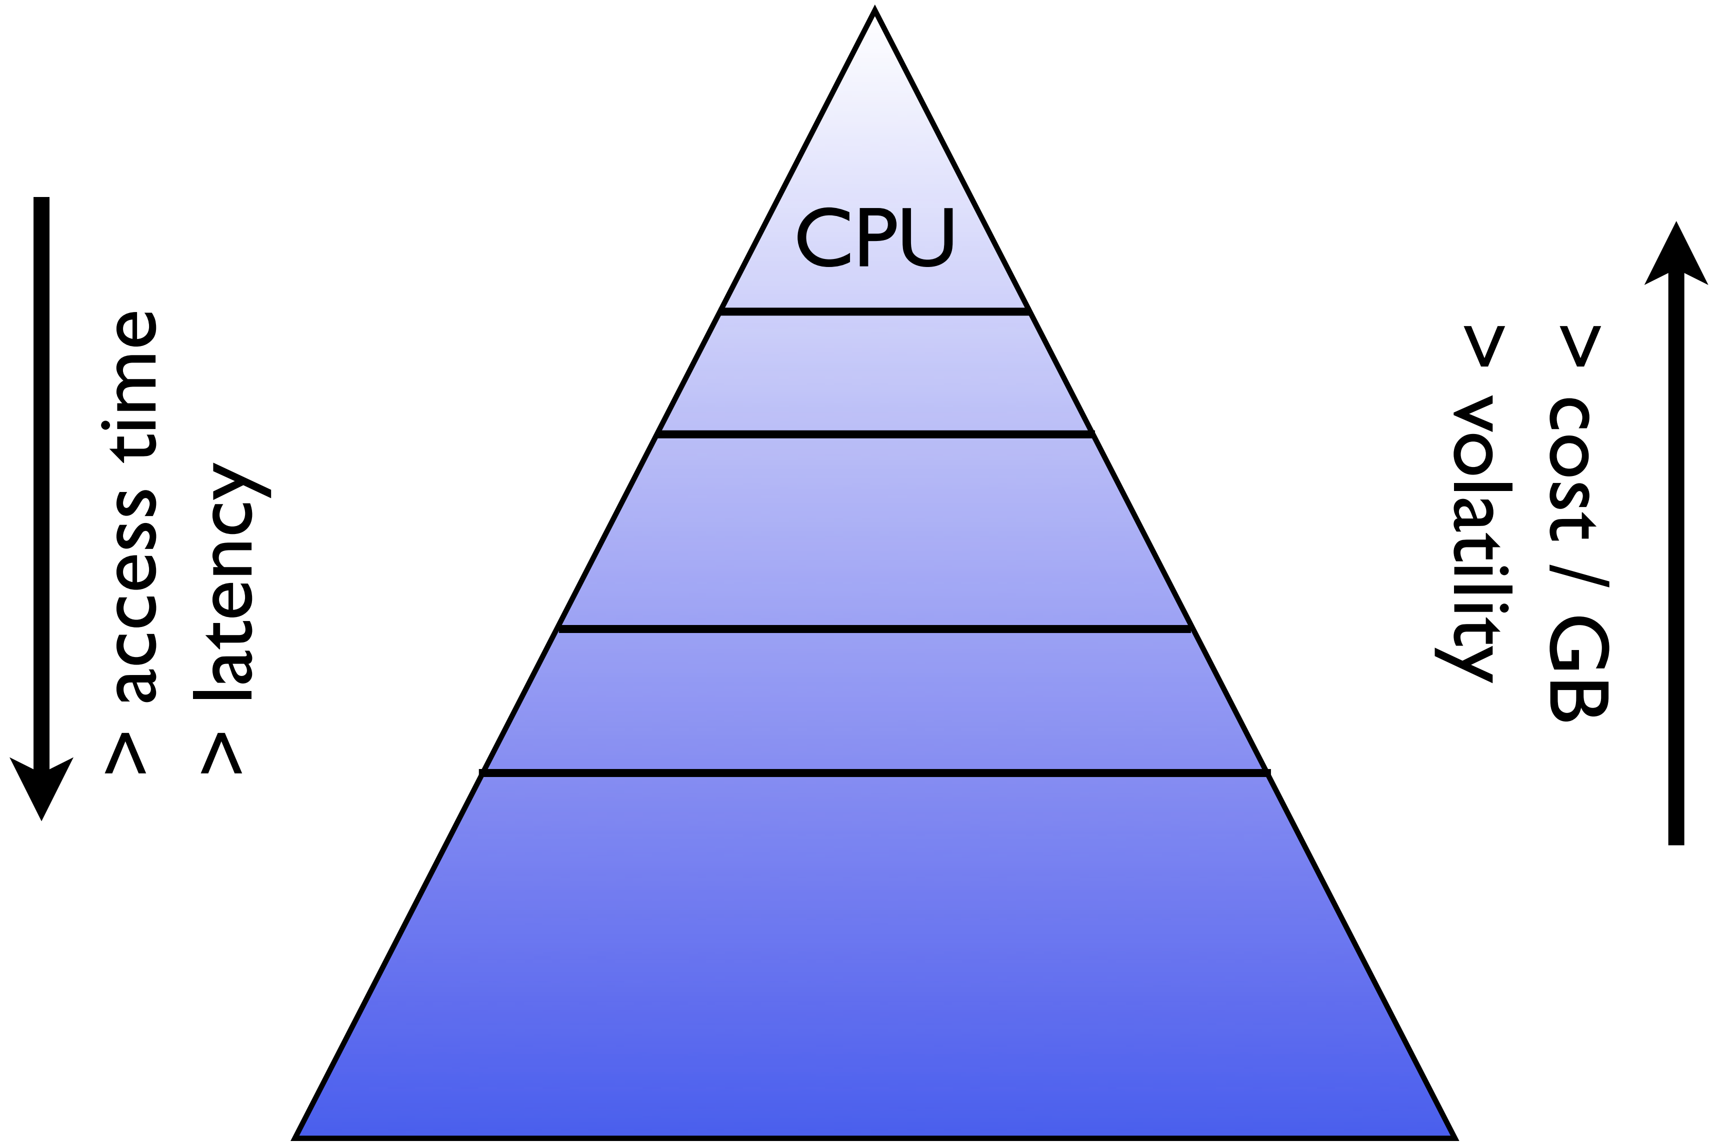

basic PC architecture

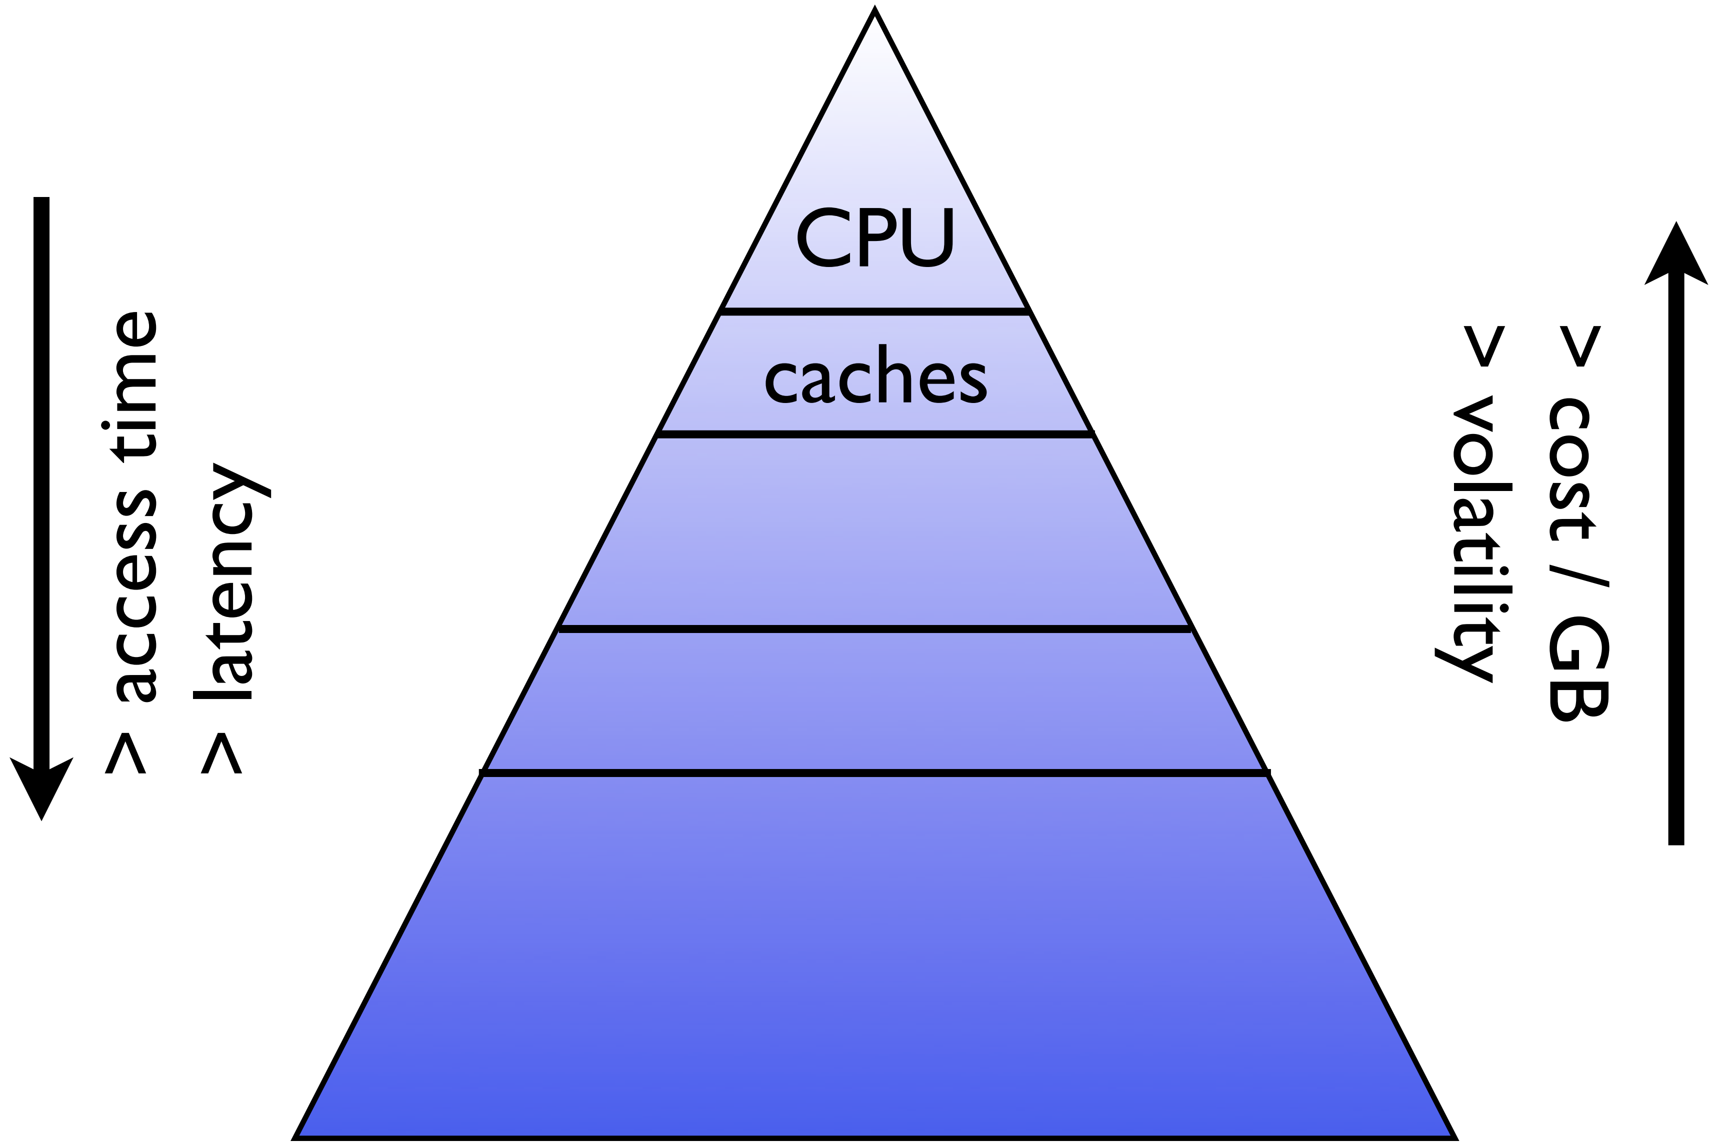

basic PC architecture

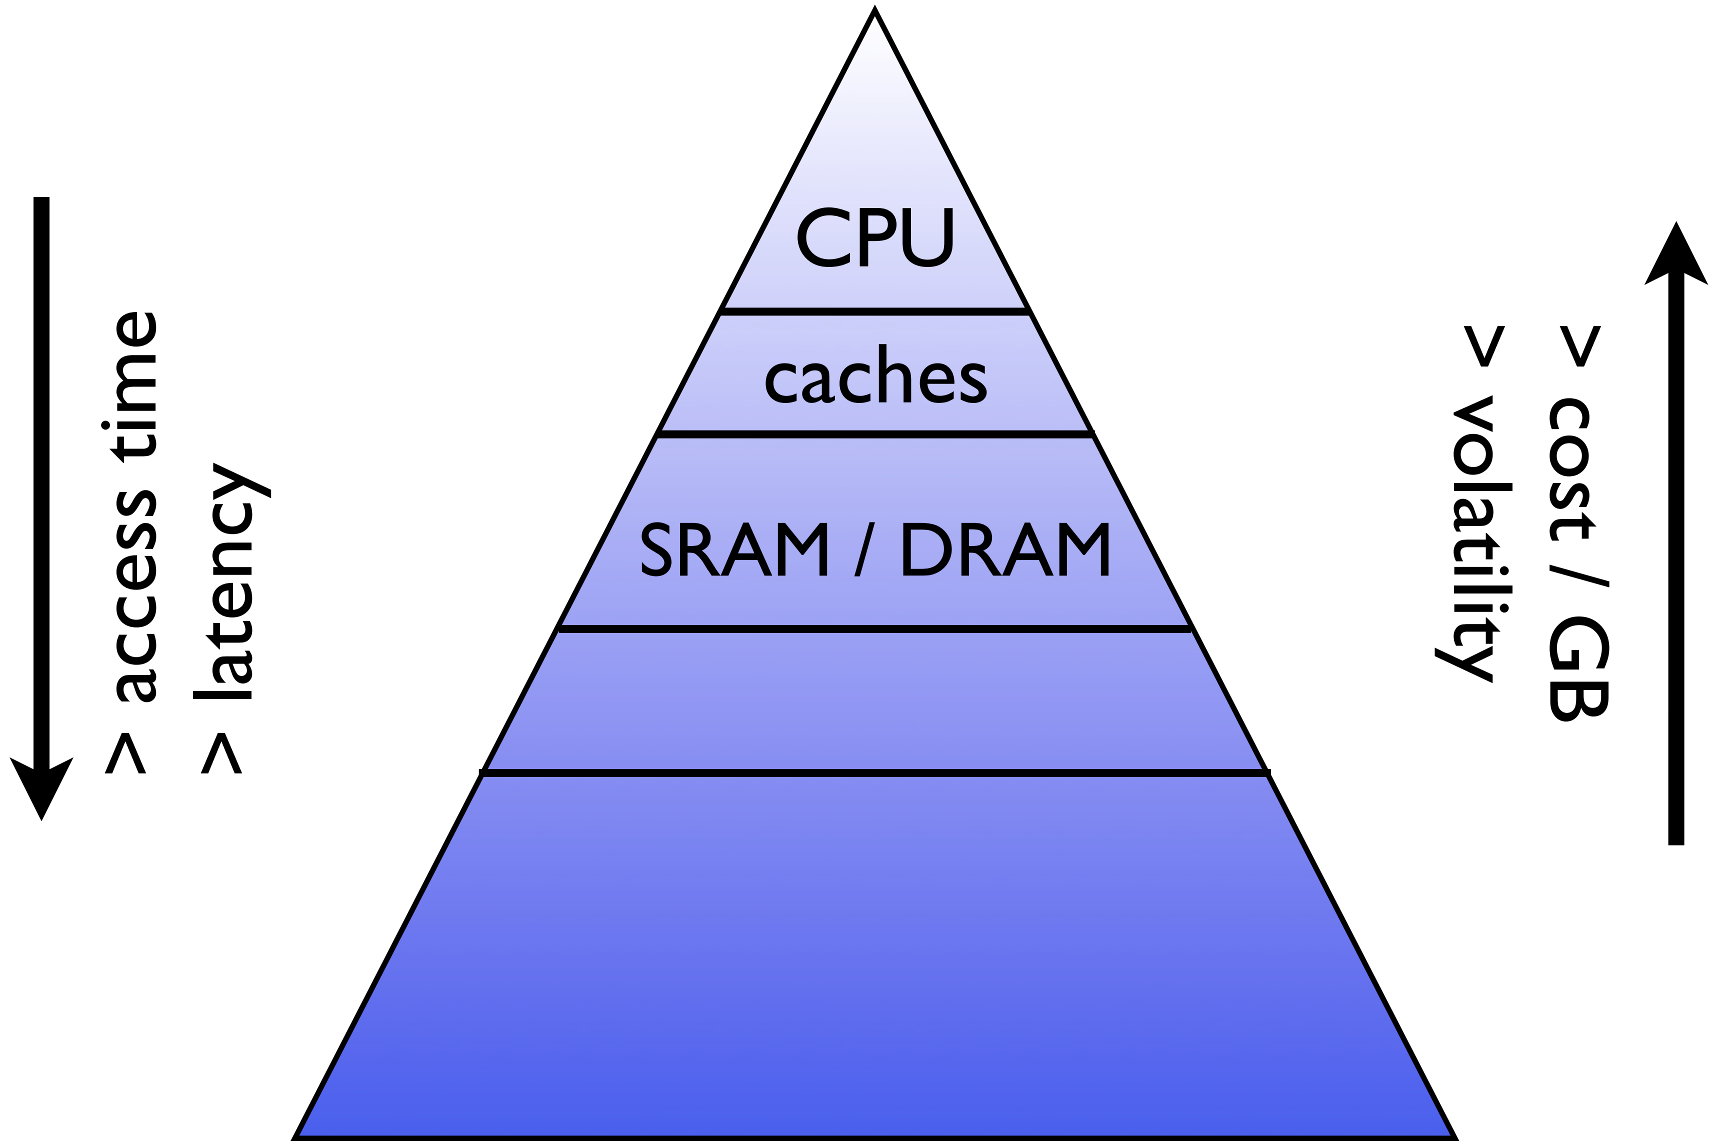

basic PC architecture

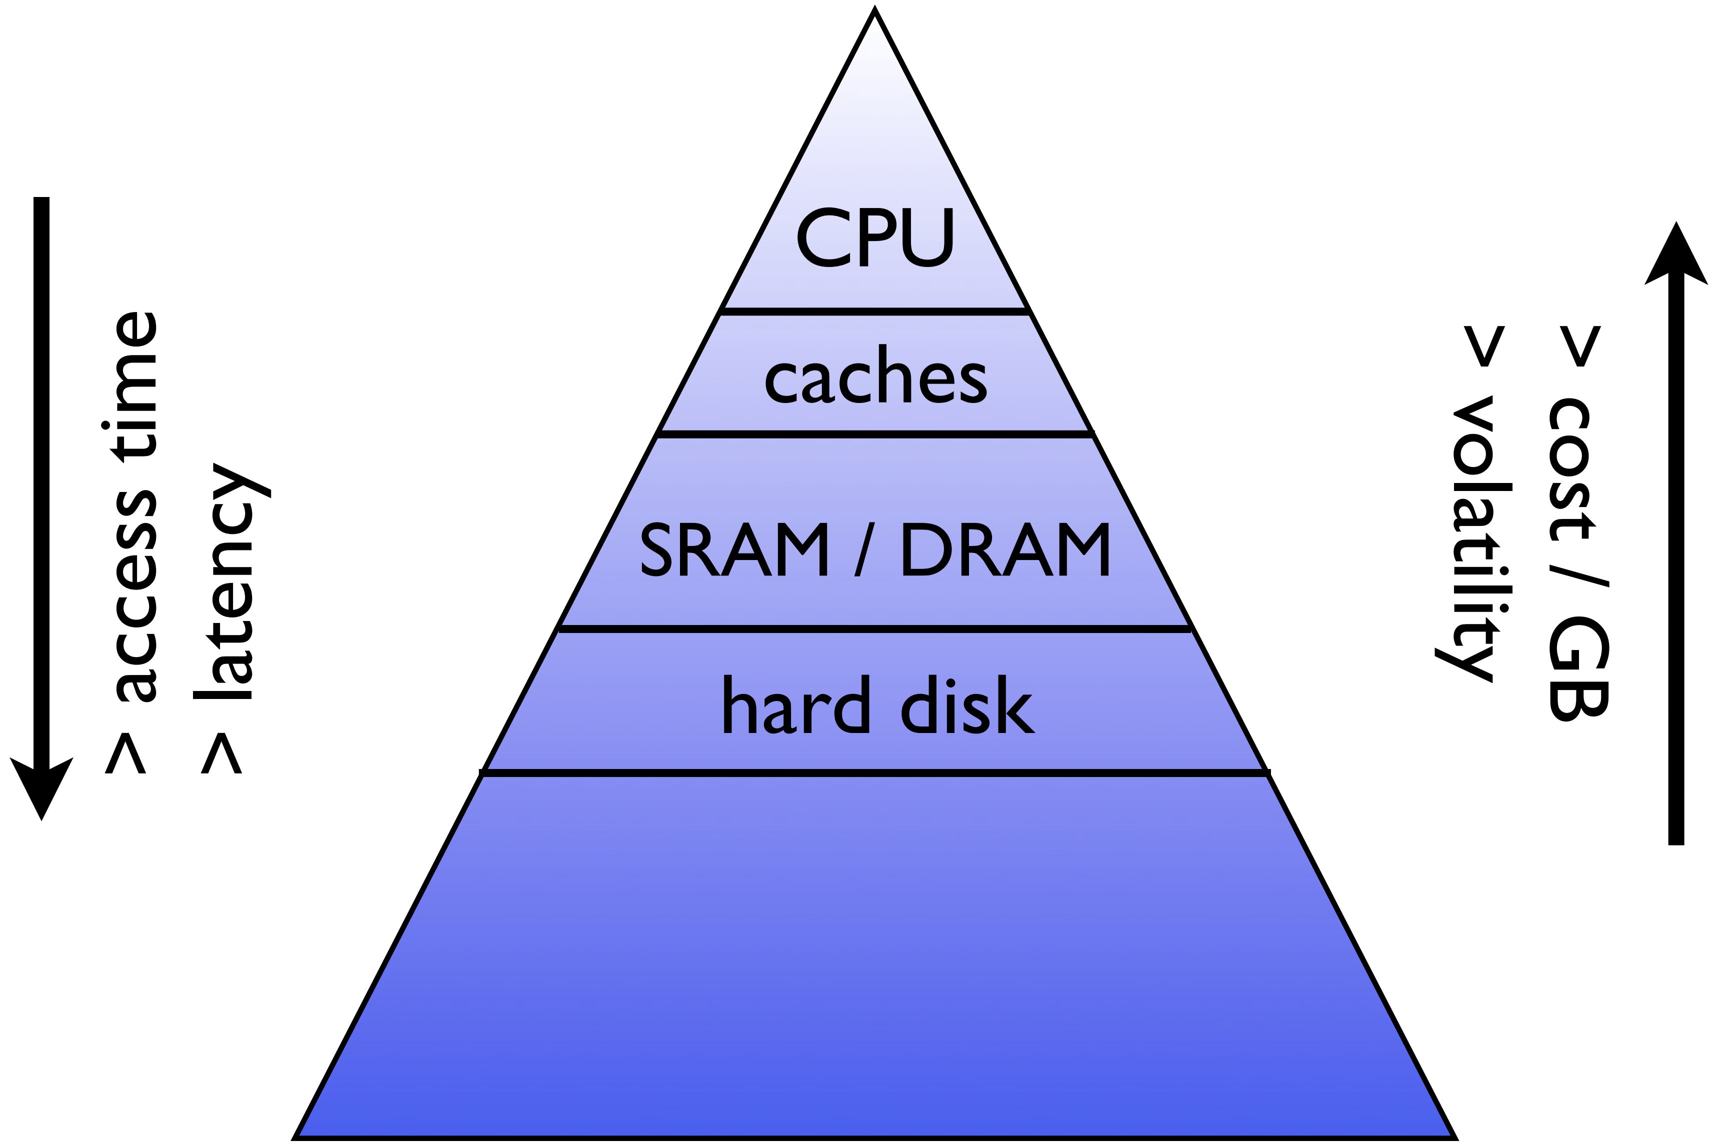

basic PC architecture

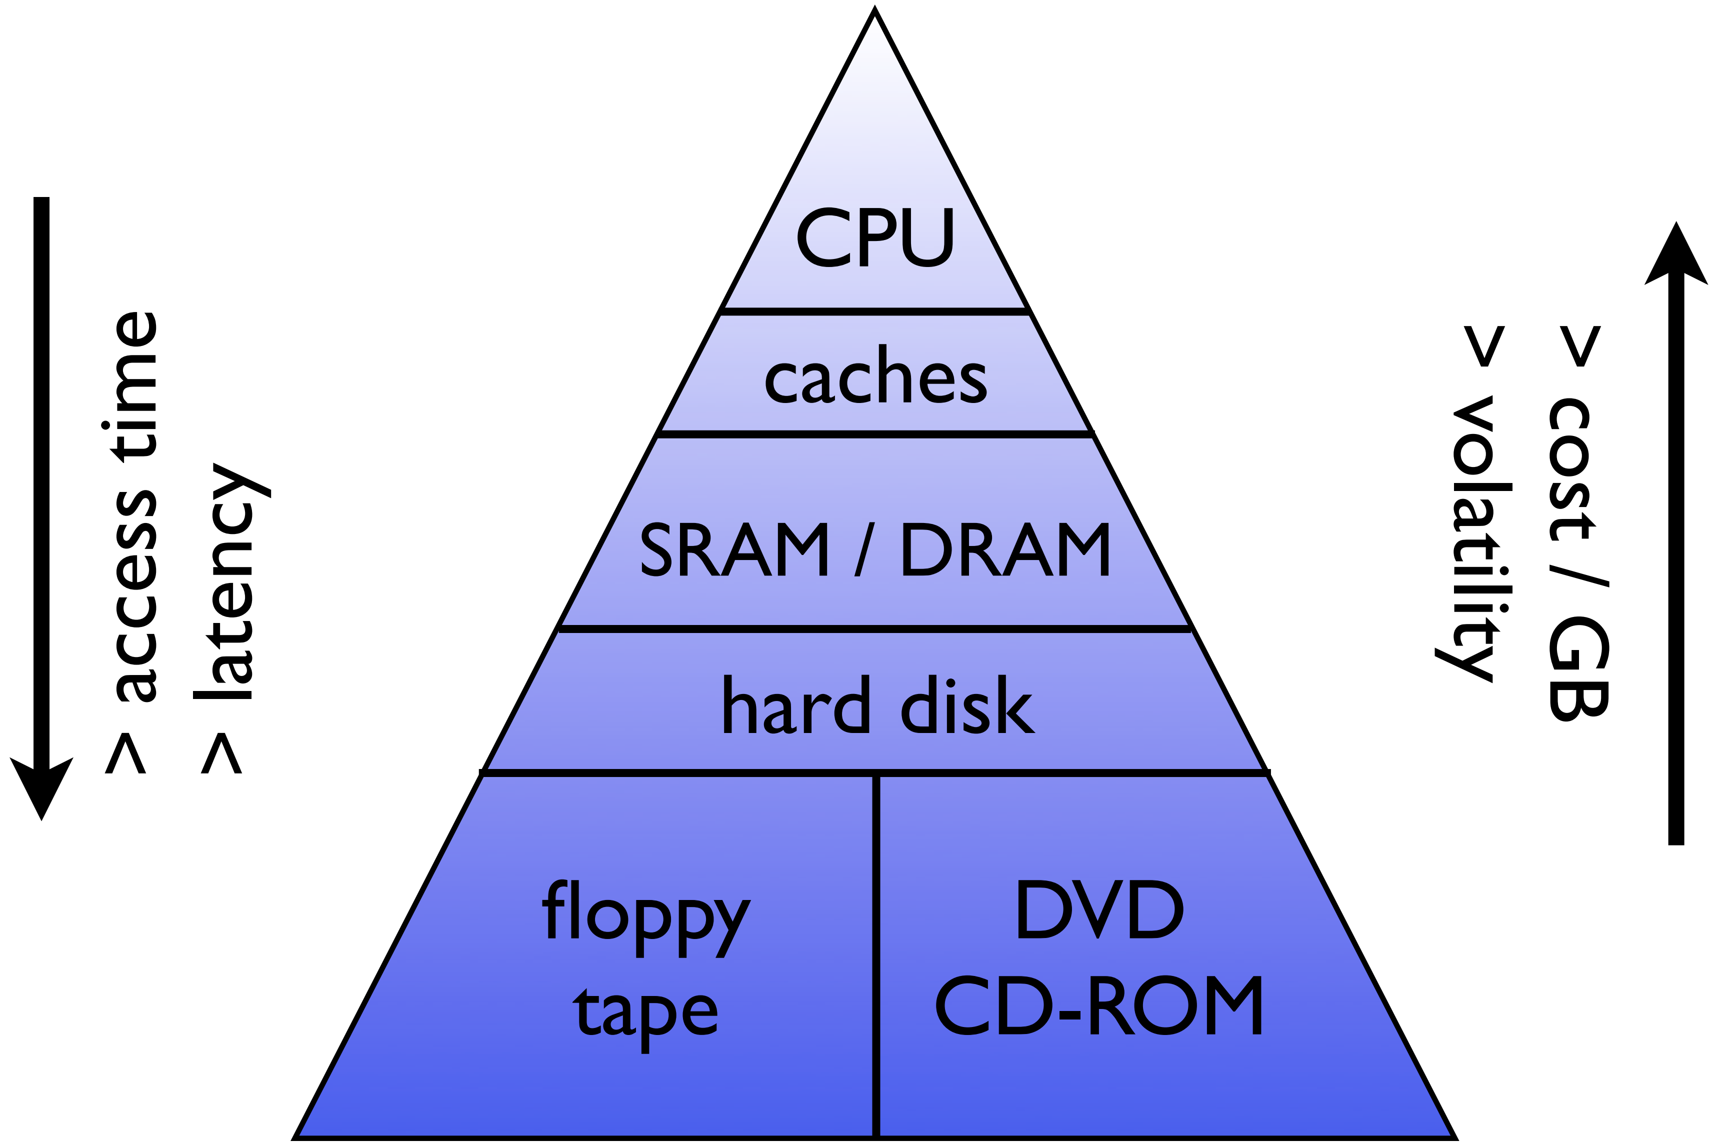

basic PC architecture

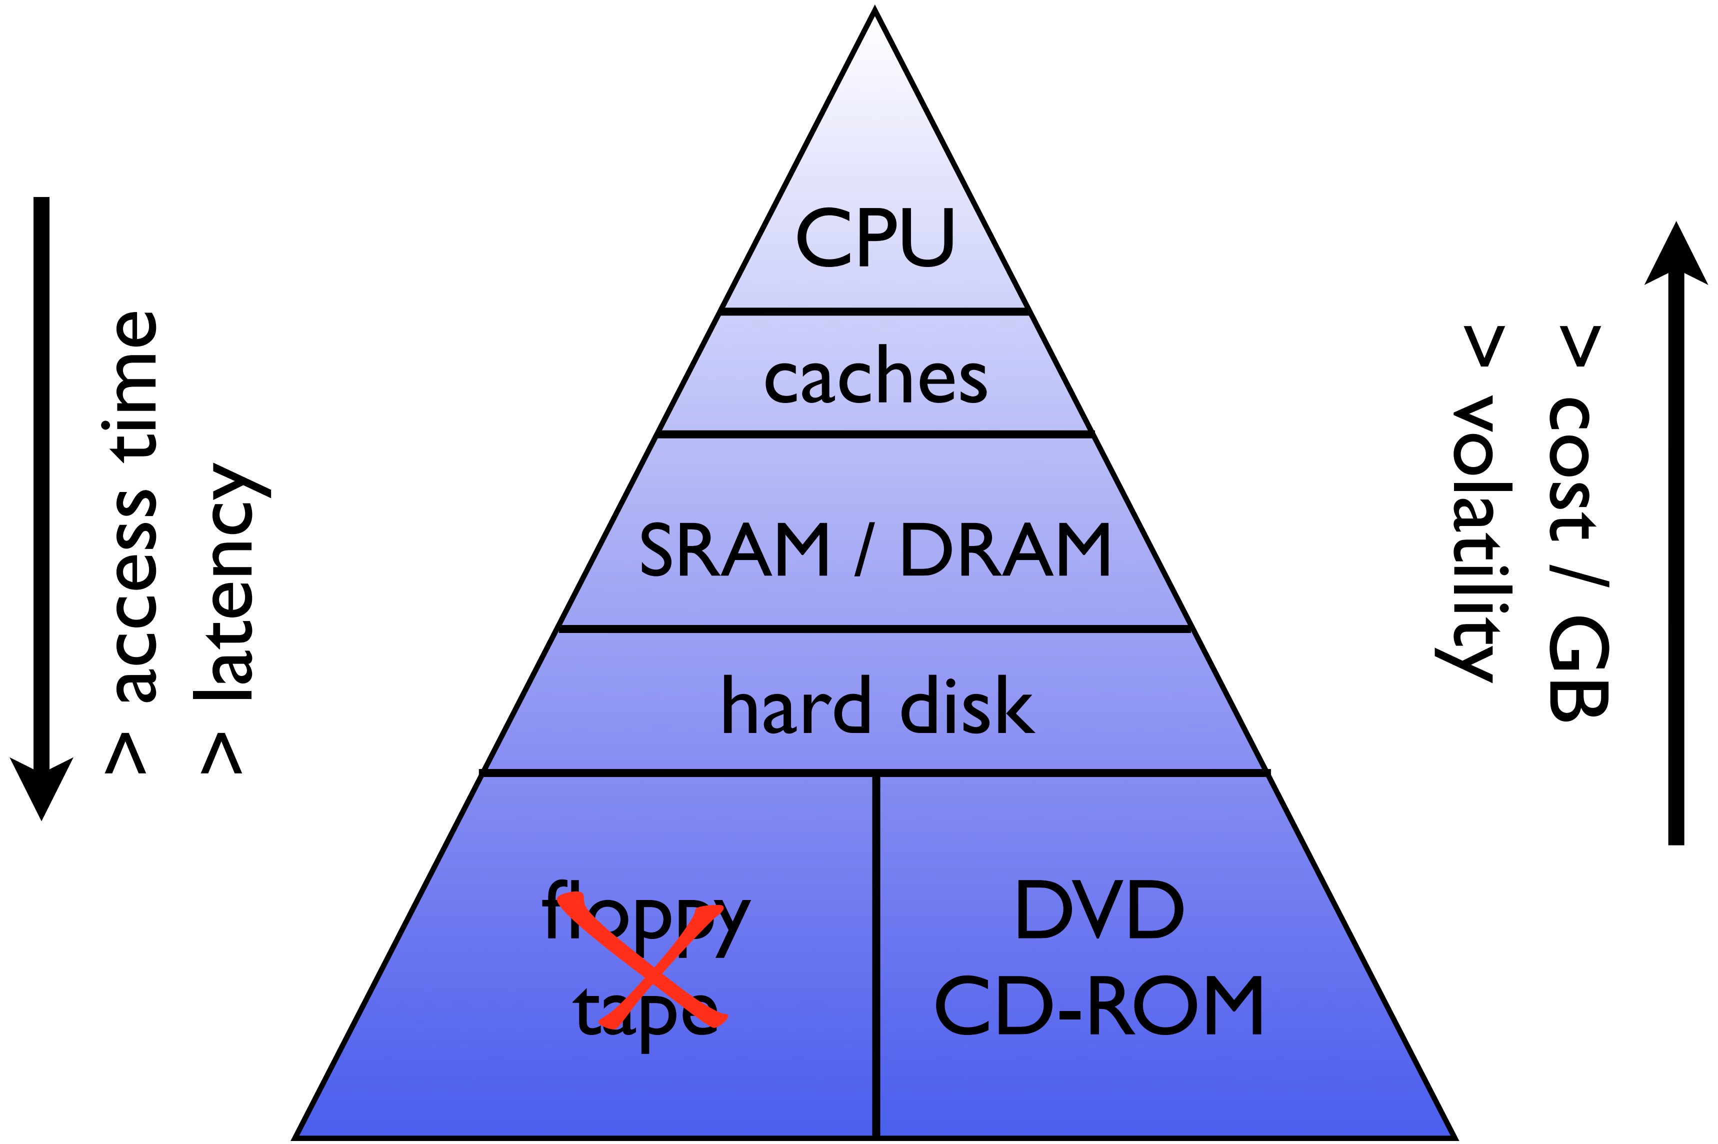

basic PC architecture

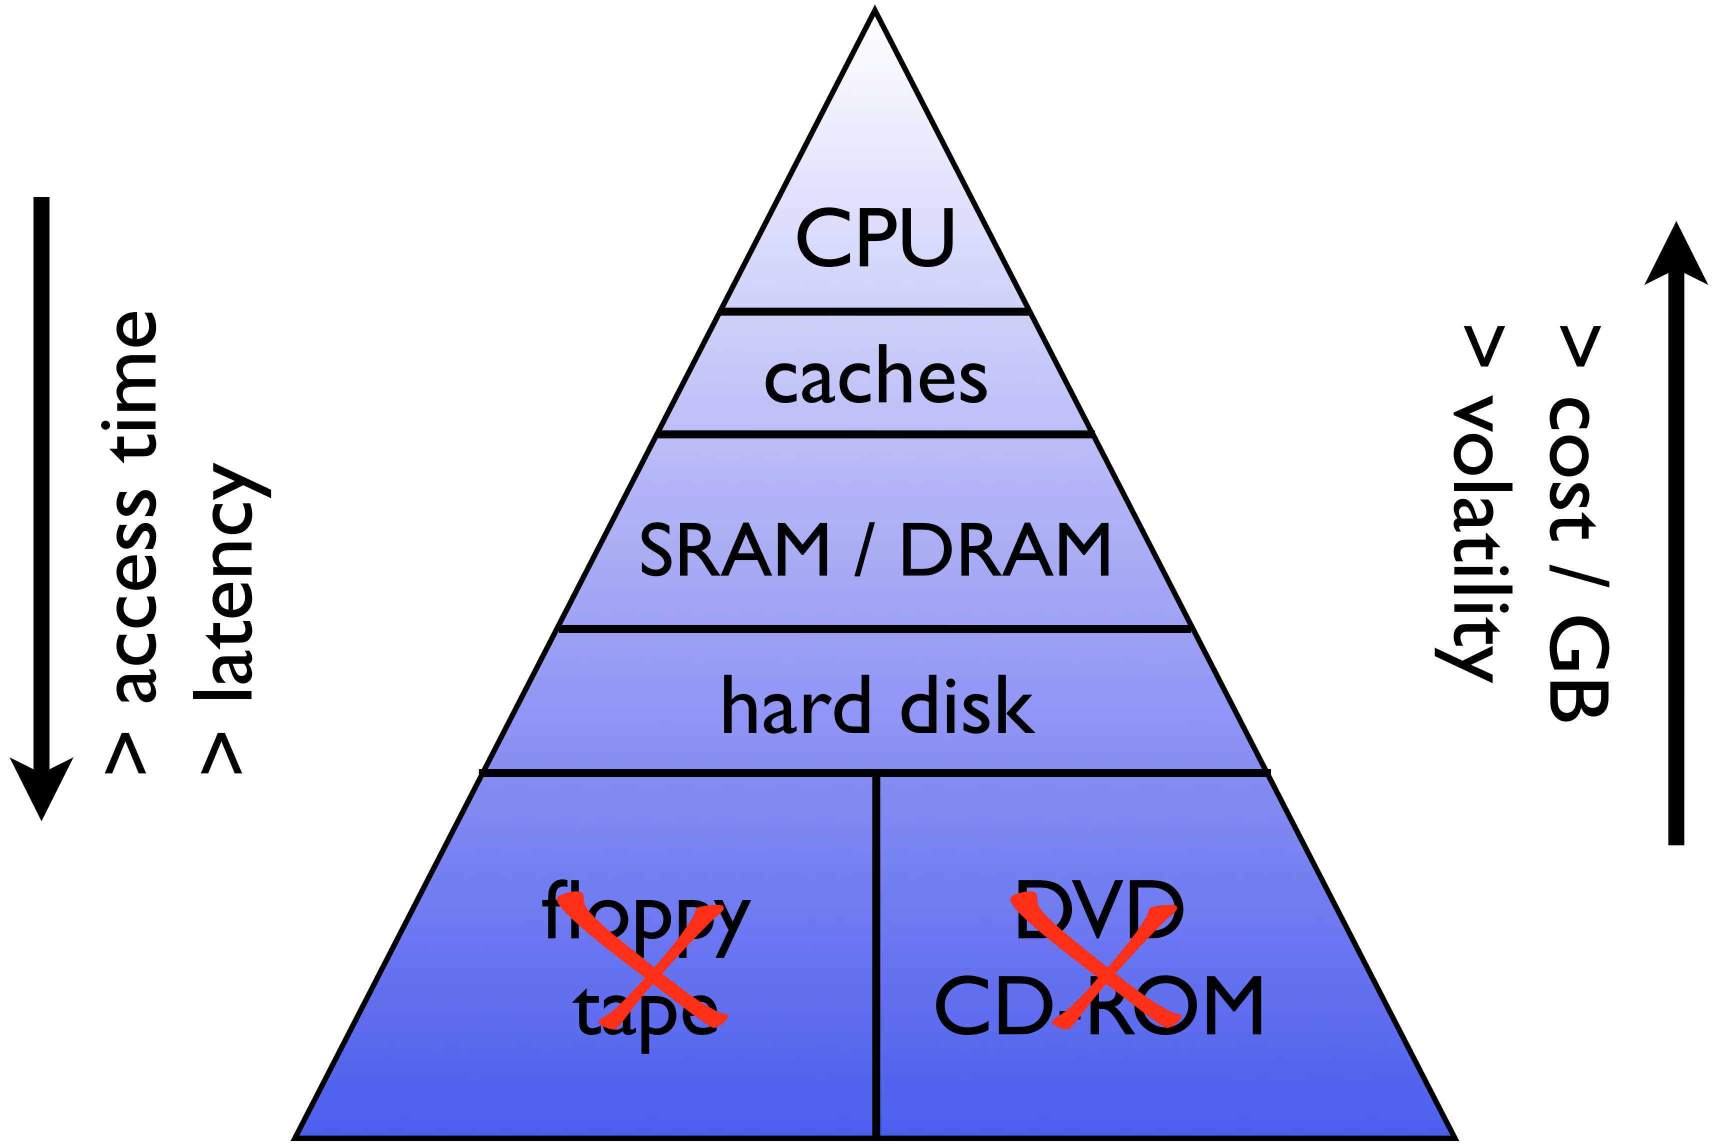

terminology

RAM

random access memory

ROM

read-only memory

access time & latency?

time between request for info & info returned

$ / GB

primary figure of merit.
most other things can be worked around

nonvolatility?

retains data without power

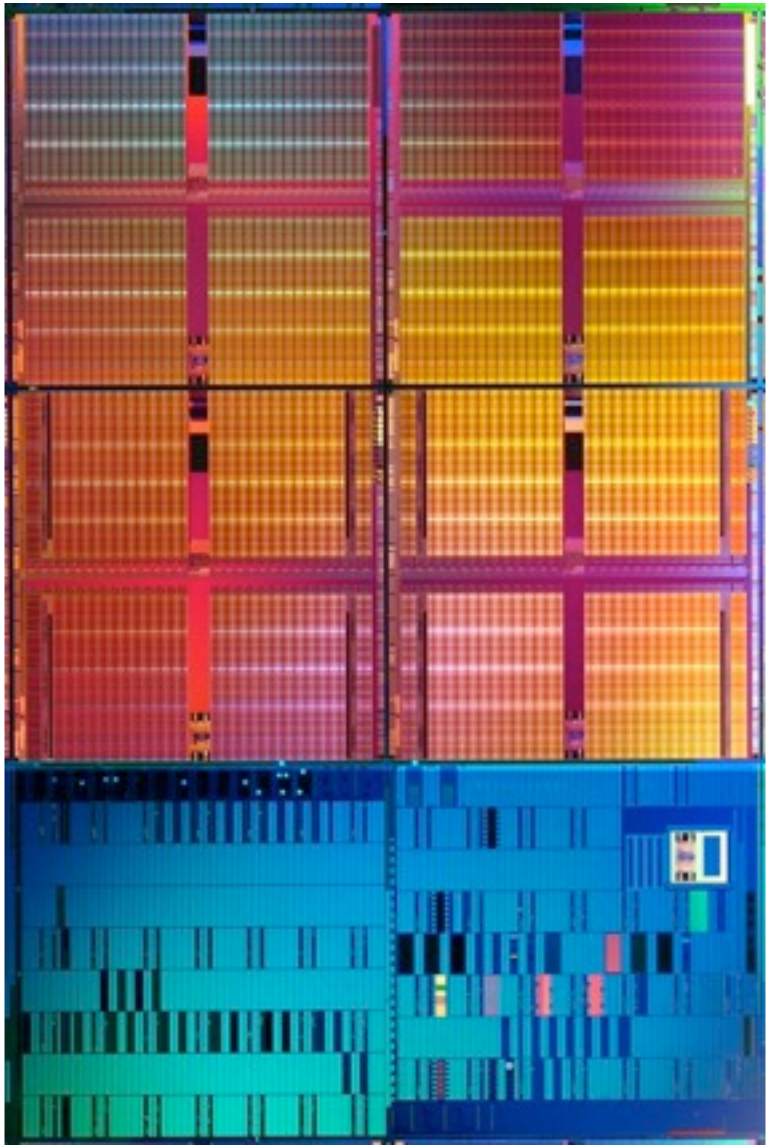

45nm SRAM die intel.com

every bit has a role

cache - *reduce latency to main memory*

- *small* memories close to CPU
- even faster than main memory
- temp storage of frequently accessed items

SRAM / DRAM - *main memory*

- blazingly fast
- relatively large
- voltaile!!

HDD - *mass storage*

- higher latency
- enormous capacity
- nonvolatile

removable

- portability
- backup
- large ROM

EDSAC / wikipedia.org

future paradigm shifts? distributed net storage?

the need for hard disks (tech)

volatility of semiconductor memories!

- *some* sort of nonvolatile storage necessary
- why not just battery backup of SRAM?

cost per GB

- SRAM/DRAM are too expensive
- Flash is too expensive
- cache RAM is more expensive

size & throughput

- higher latency, but bandwidth is huge
- enormous sizes

endurance

- essentially unlimited cycling
- radiation hard

punched cards are nonvolatile

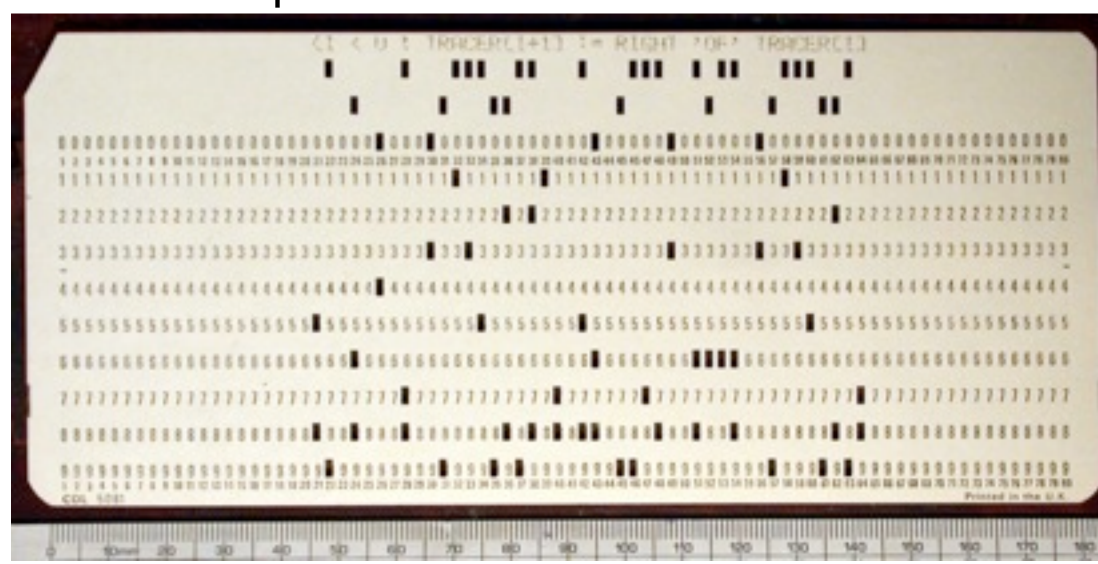

Back in the day, disks were expensive.

Sometimes, we would trick the system into using RAM as a disk to avoid swapping floppies.

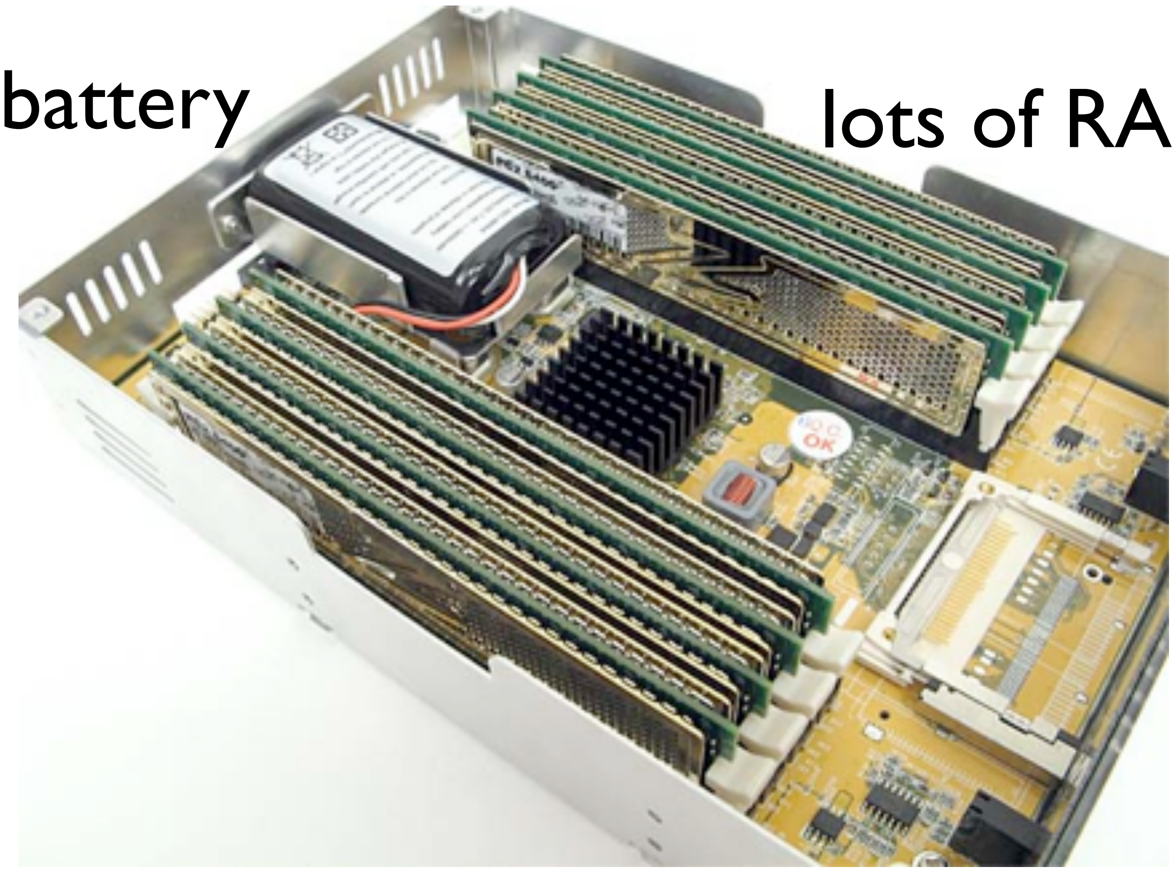

now RAM disks make a comeback ...

the need for mass storage (human)

sound

several MB per minute / lossy
tens of MB per minute / “lossless”

pictures

several MB per image

video

~ 1 MB per sec
several GB per movie
with lossy compression!

data mining

enormous sizes

apple.com

how do hard disks work, more or less?

wikipedia.org - “Hard_Disk”

spinning (~10^4 rpm) part holds data.
sliding part reads and writes data.

hard disk drives

160 Gbit 2.5" perpendicular drive for laptops

images from M. Coey

hard disk drives

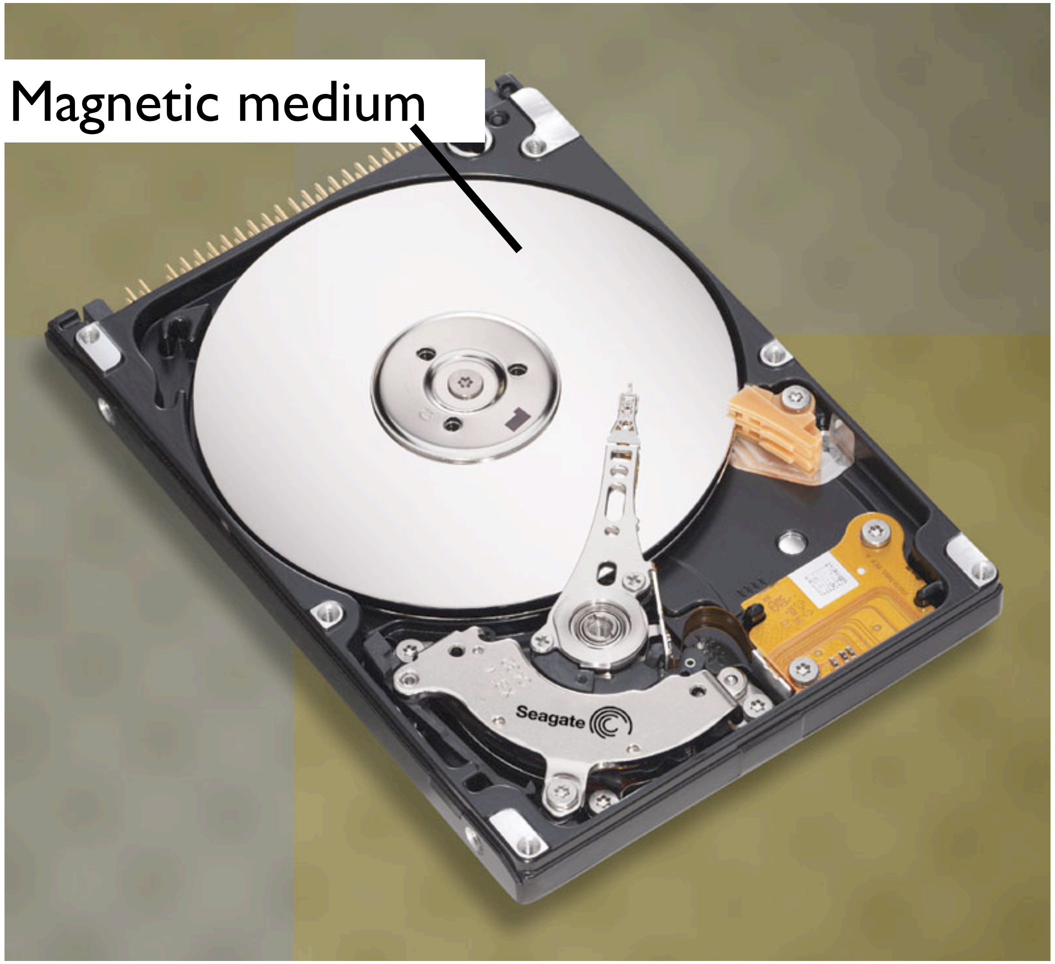

160 Gbit 2.5" perpendicular drive for laptops

images from M. Coey

hard disk drives

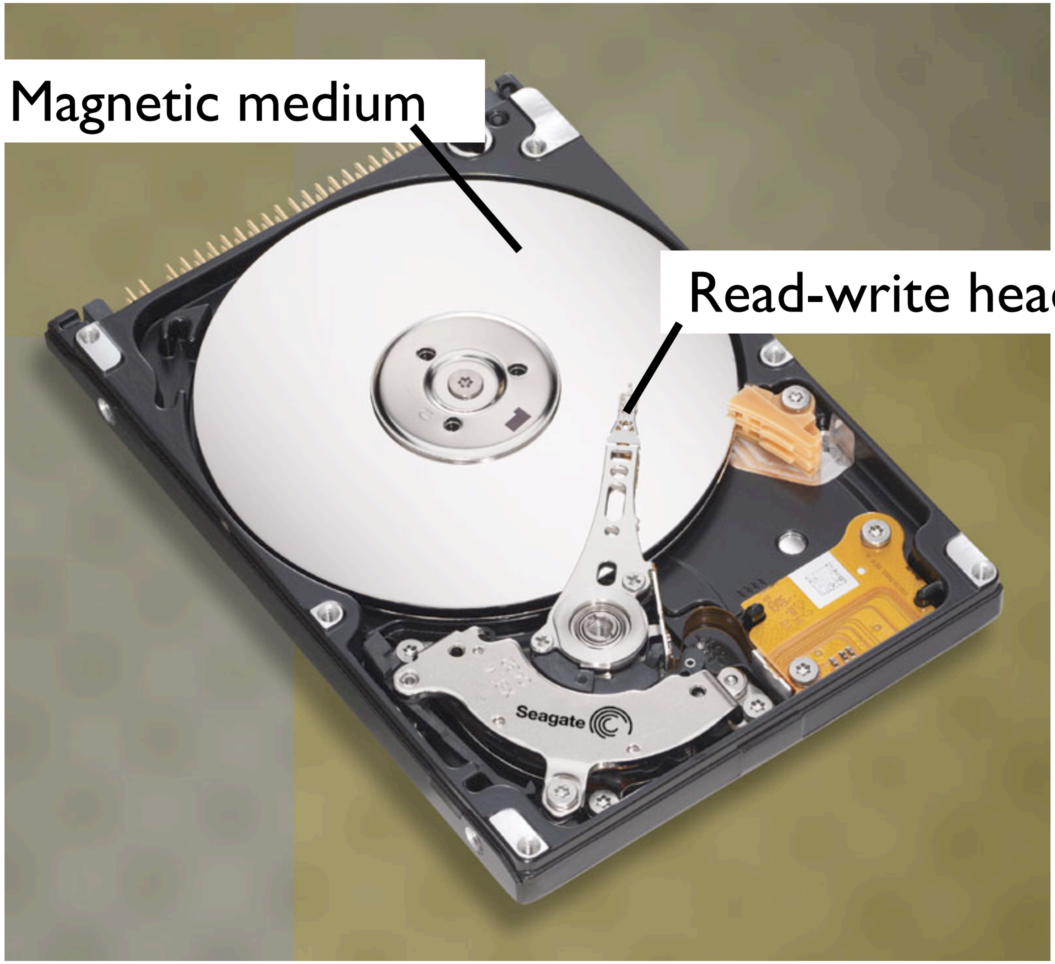

160 Gbit 2.5" perpendicular drive for laptops

images from M. Coey

hard disk drives

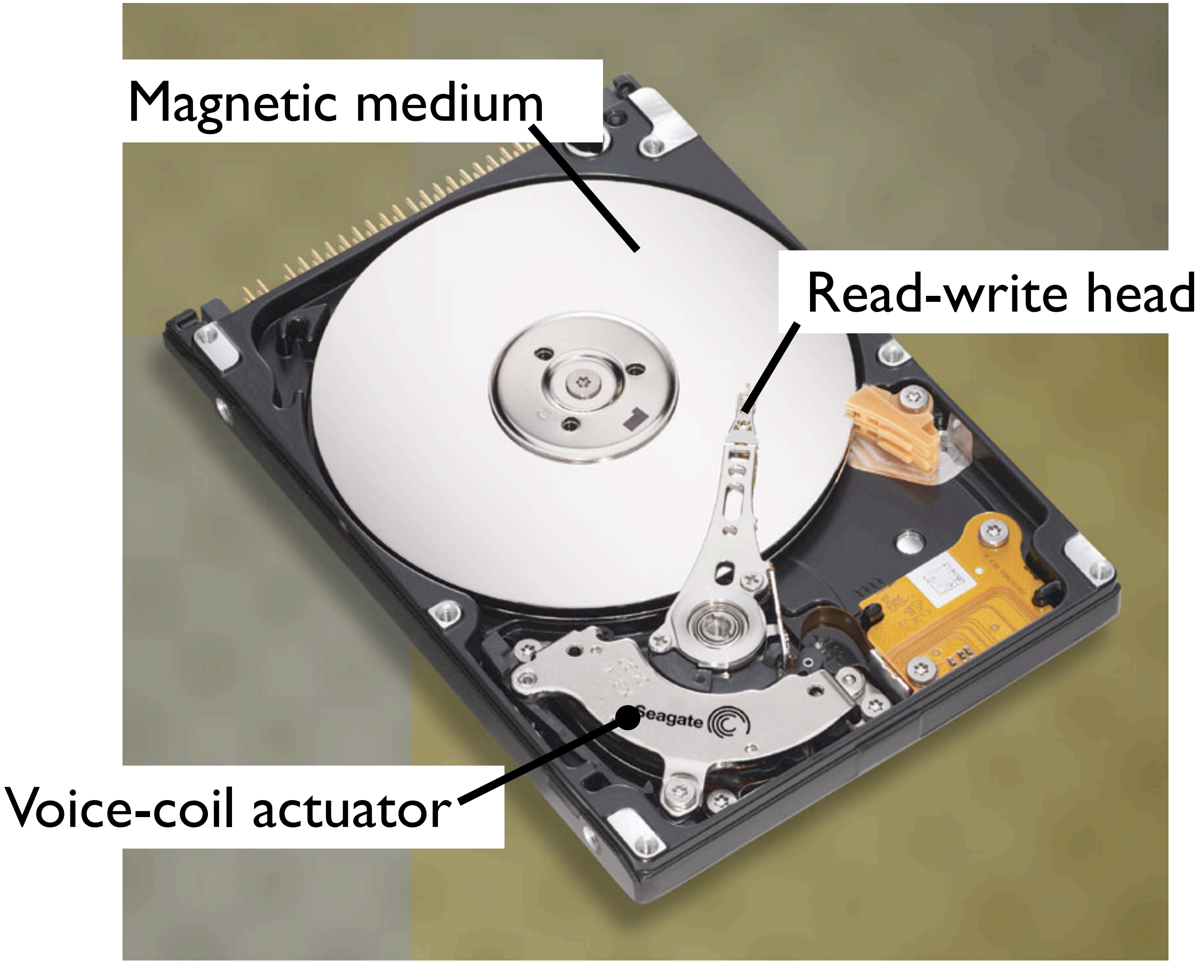

160 Gbit 2.5” perpendicular drive for laptops

images from M. Coey

hard disk drives

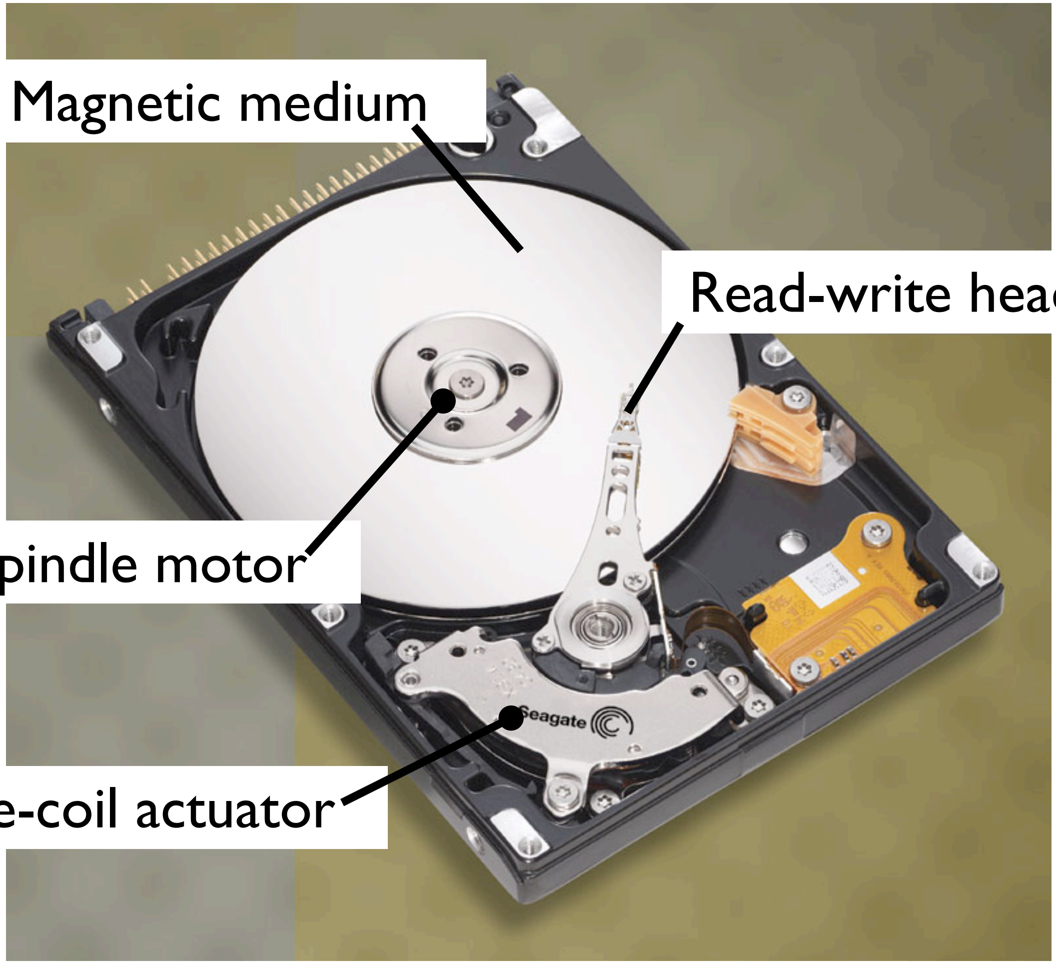

160 Gbit 2.5” perpendicular drive for laptops

images from M. Coey

hard disk drives

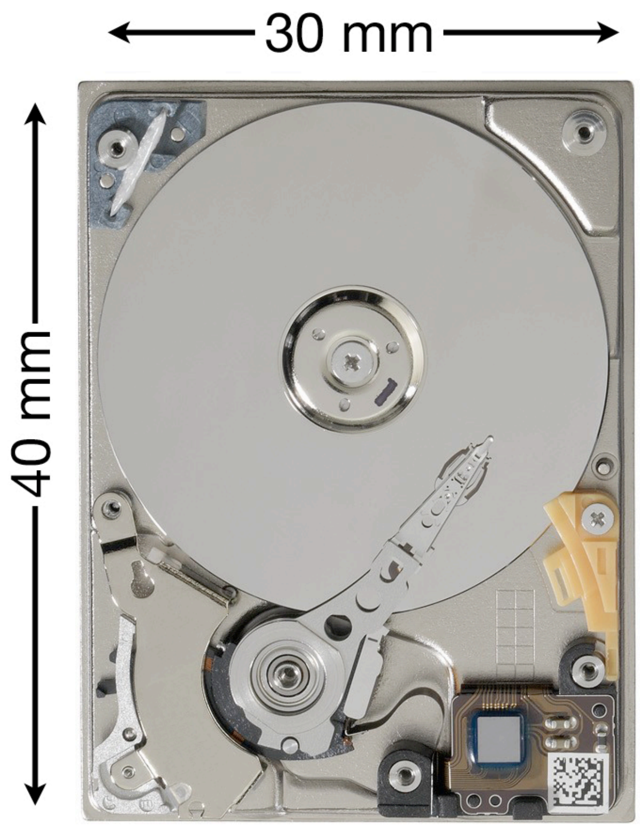

8 Gbit 1” drive for cameras

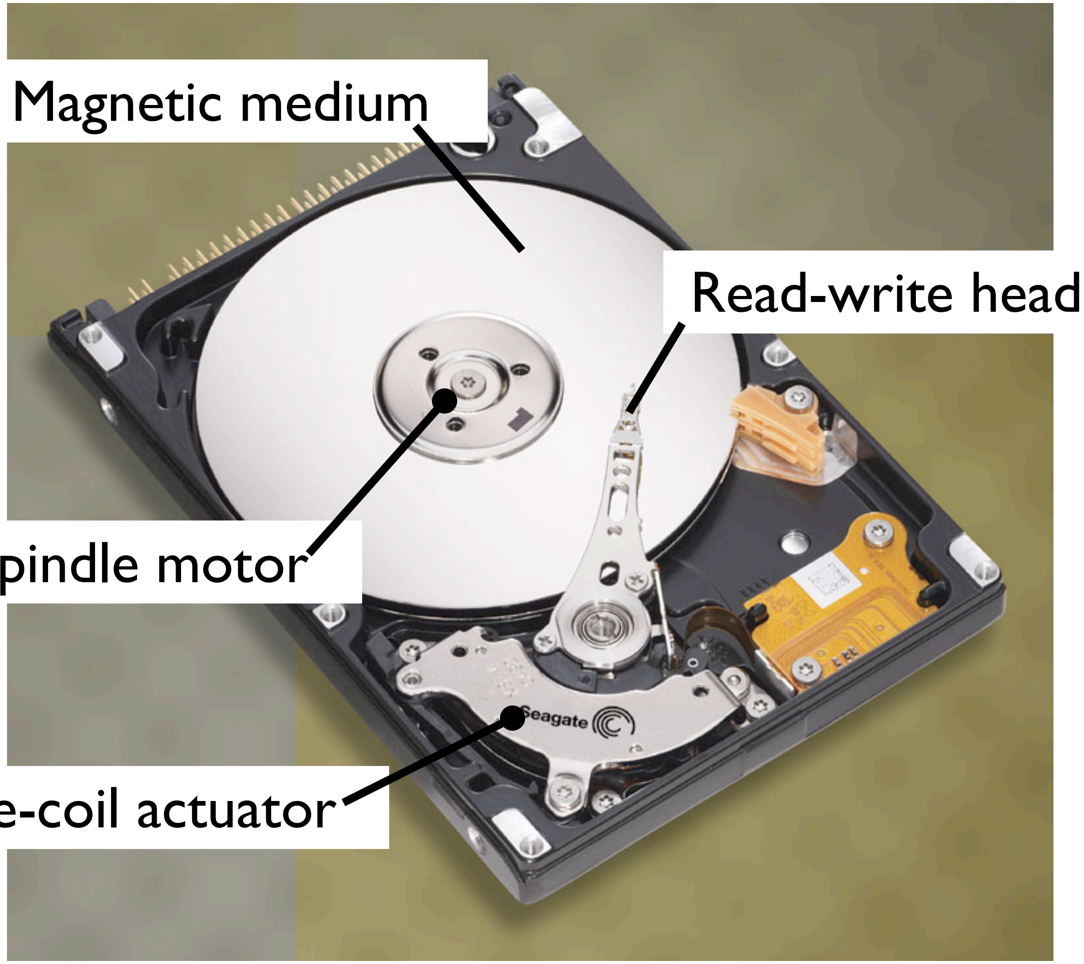

160 Gbit 2.5” perpendicular drive for laptops

images from M. Coey

media basics

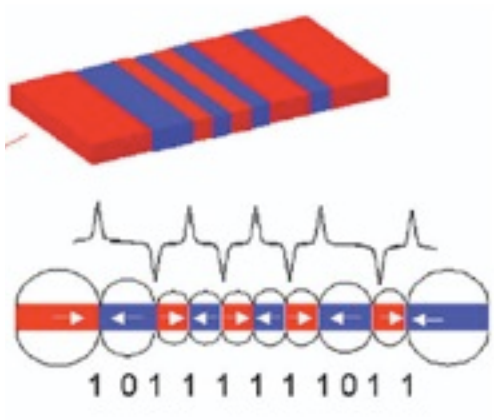

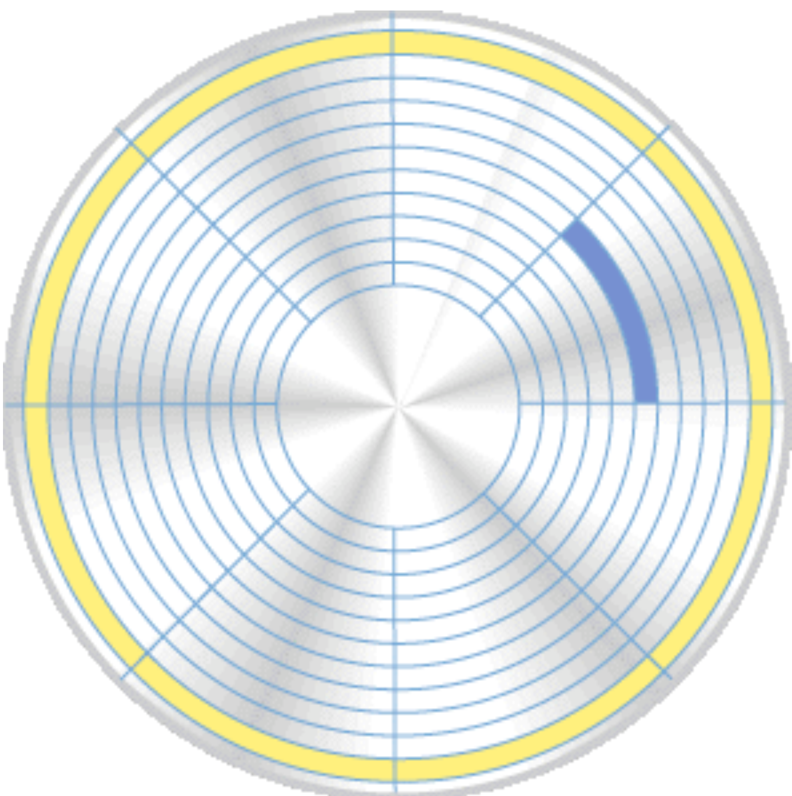

Hard disk

tiny magnetized regions

direction (N/S) stores bit

magnetic sensor reads bits

LP records

tiny bumps

needle moves

actual record grooves

CDs

pits store bits

optical reflectivity

actual CD surface

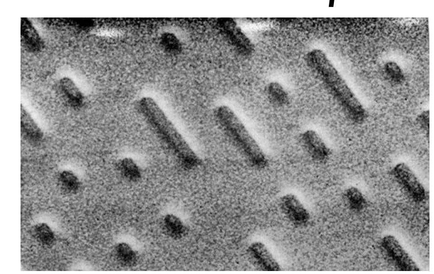

media basics

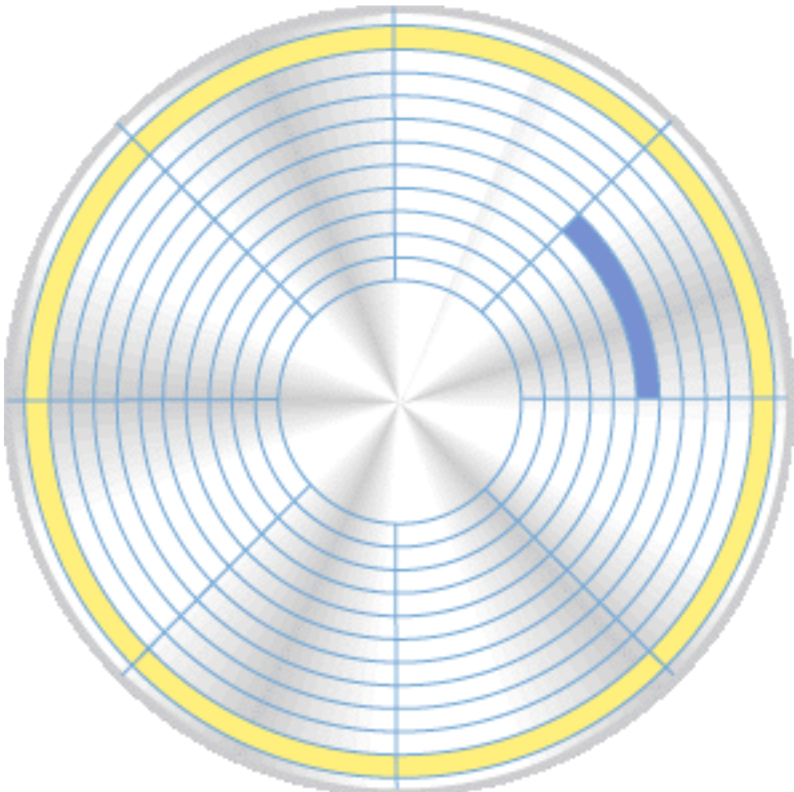

hard disk platters are round.

so how is data arranged?

tracks = concentric circles
sectors = wedge of a track

sector has fixed # bytes

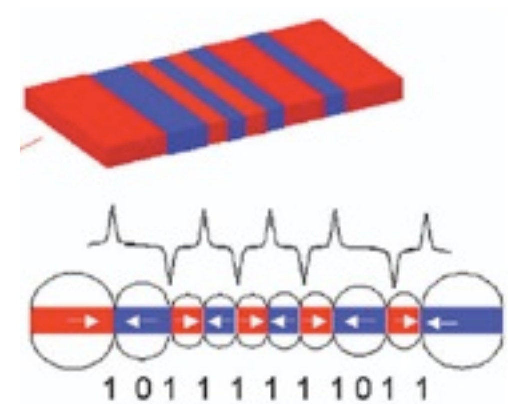

media basics

mfm image

sees transition field

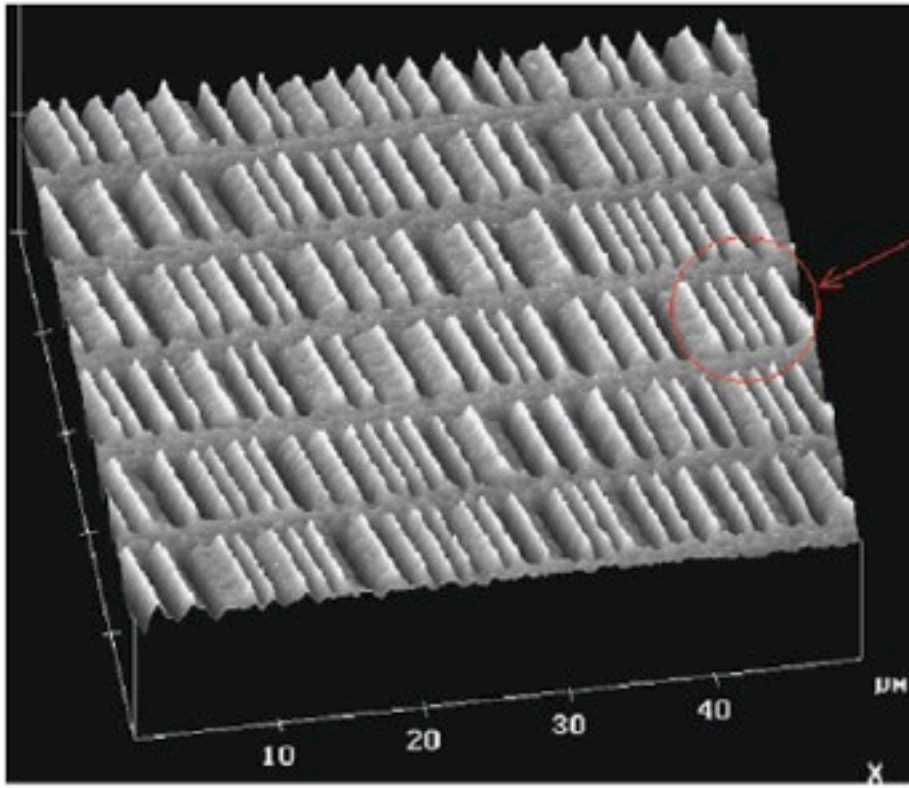

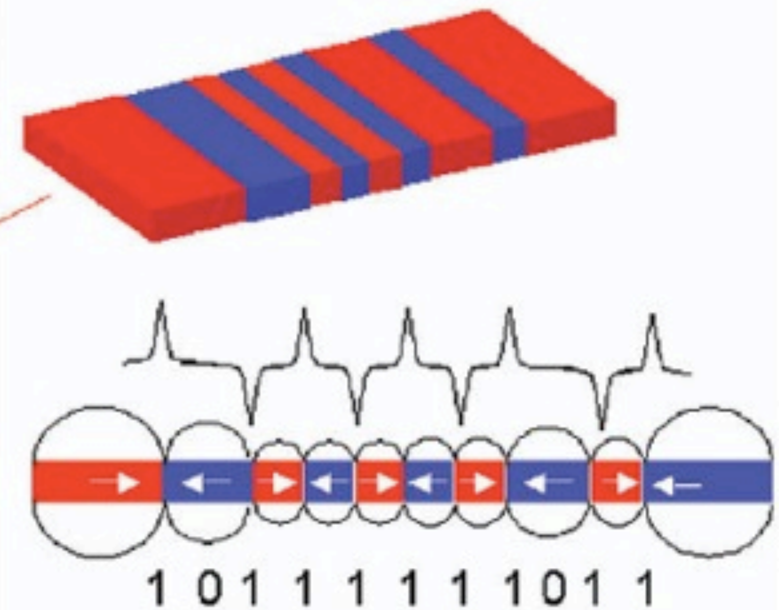

Jimmy Zhu, *Materials Today*, July/Aug 2003

CoCrPt alloy

platters - Al or glass substrate

typical magnetic region

~200-250 nm wide, ~25-30 nm down-track

100 billion bits (Gigabits) per in^2

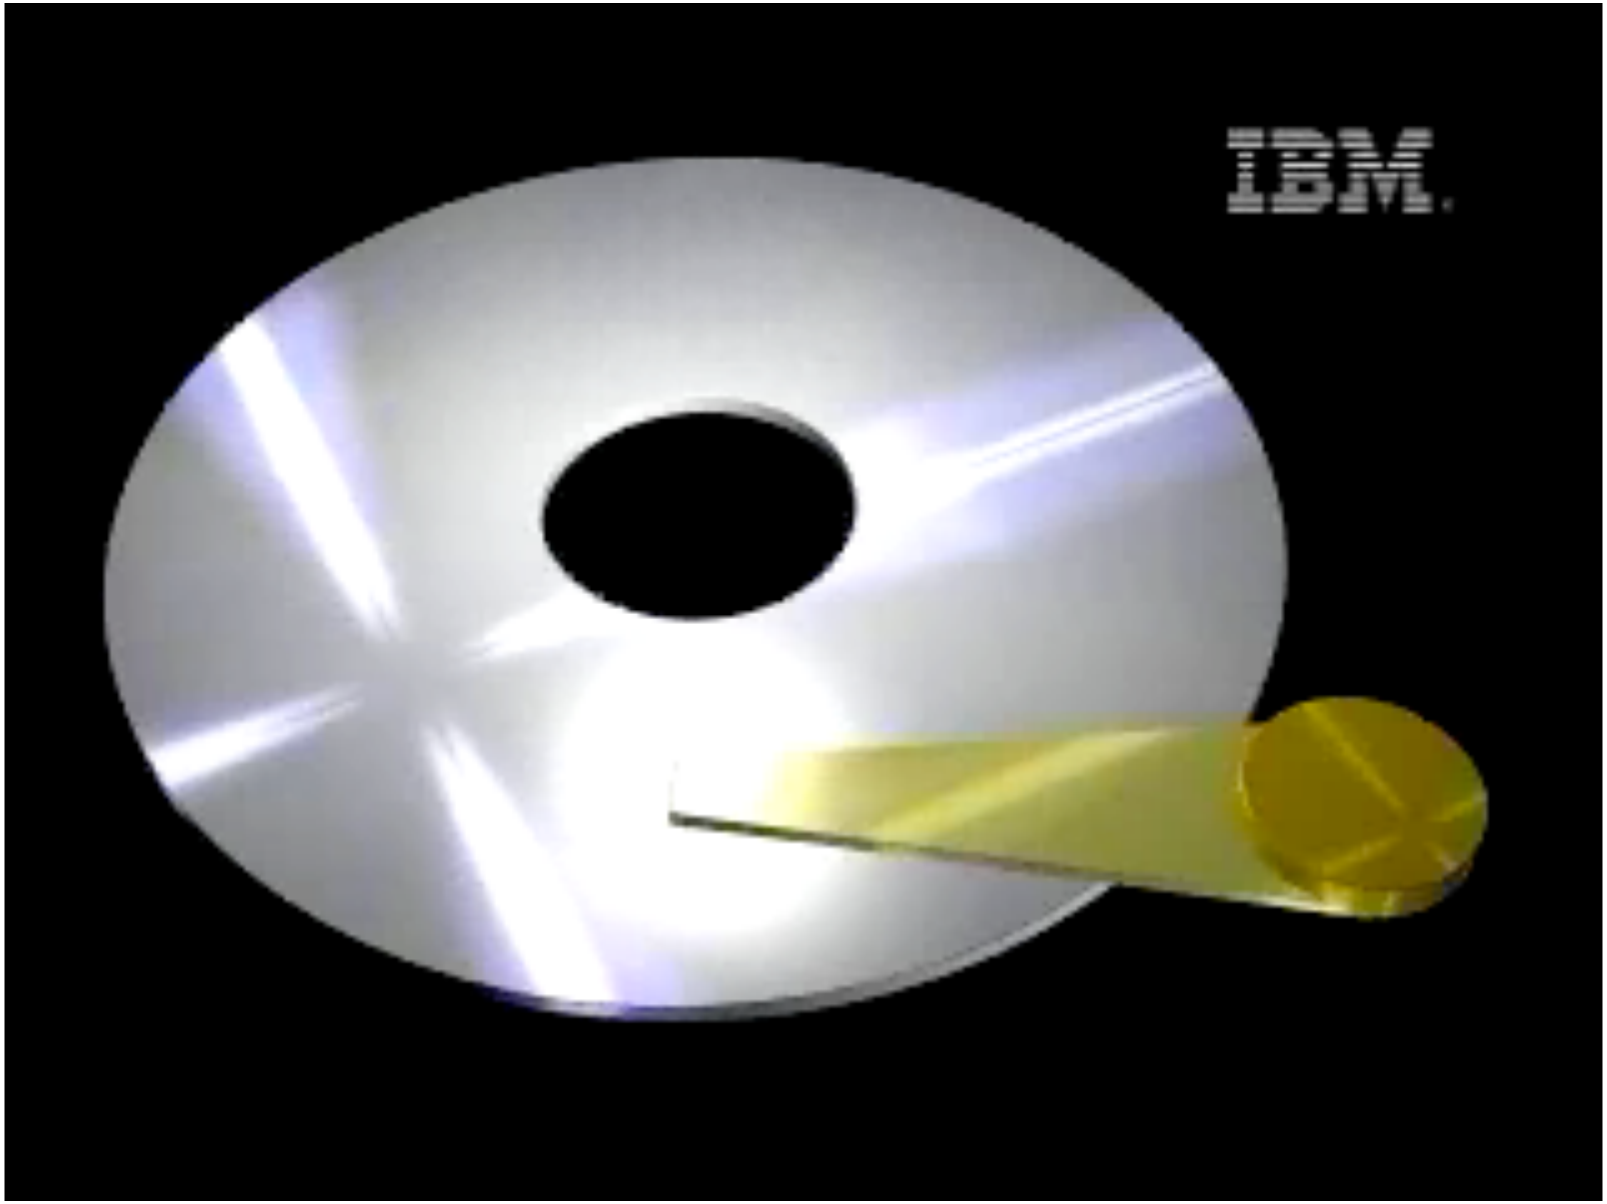
IBM.

reading and writing basics

(longitudinal recording)

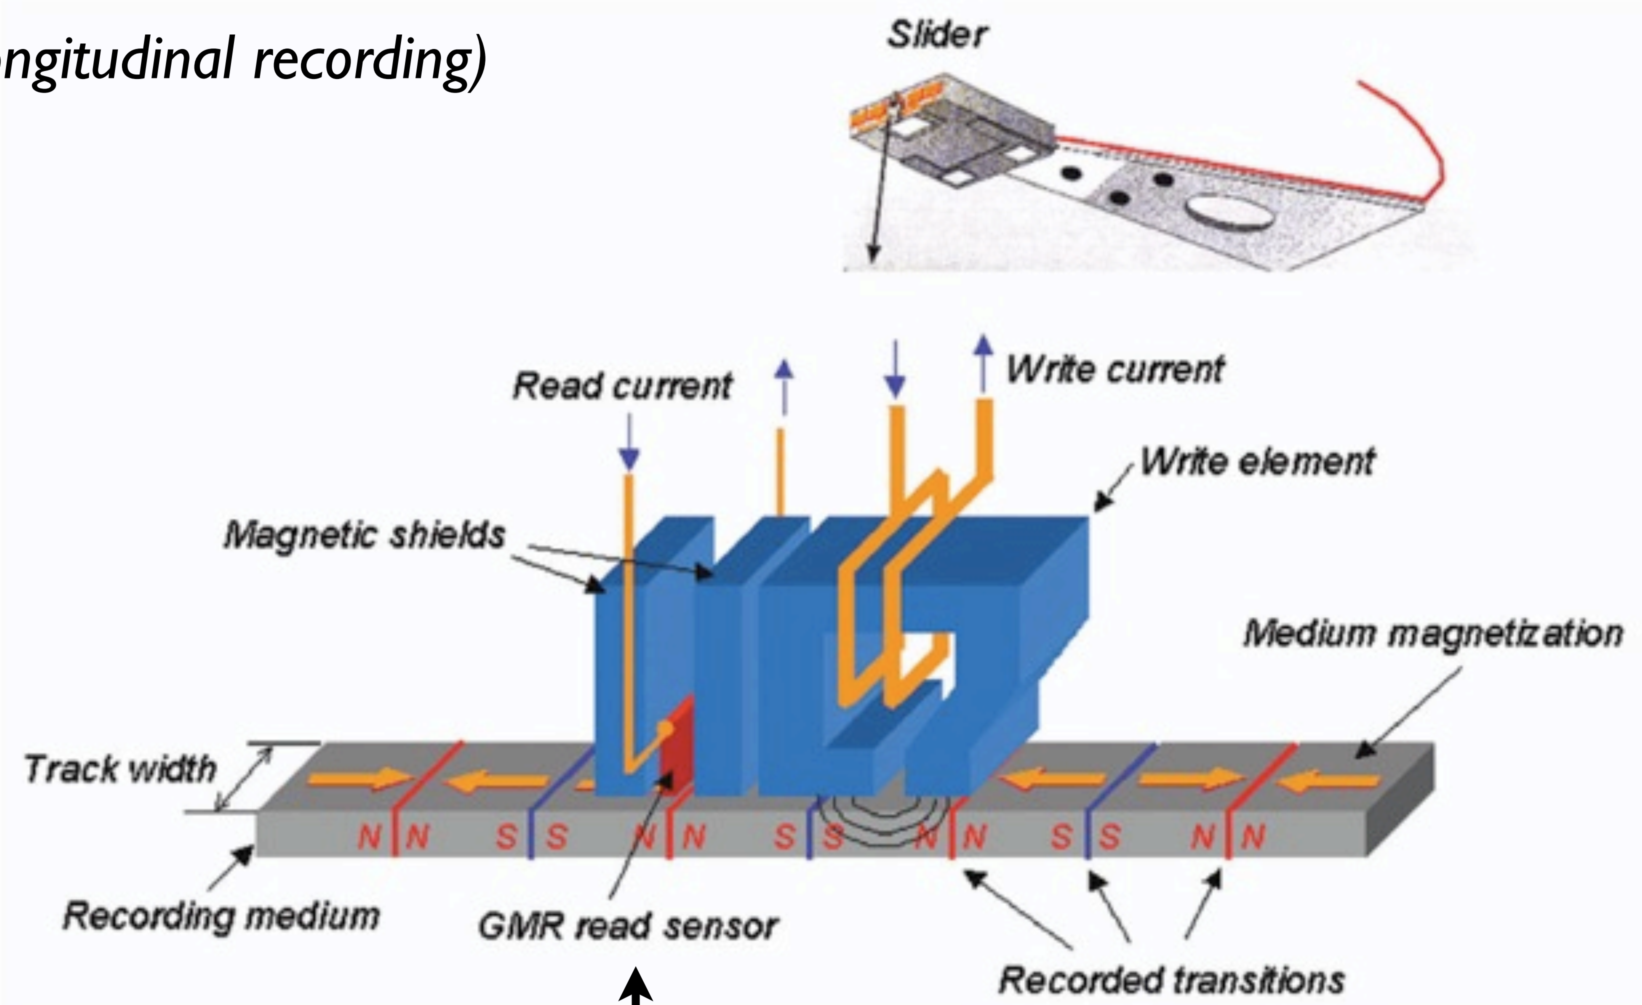

sensor - *magnetoresistive*

Jimmy Zhu, *Materials Today*, July/Aug 2003

reading and writing basics

(perpendicular recording)

Jimmy Zhu, *Materials Today*, July/Aug 2003

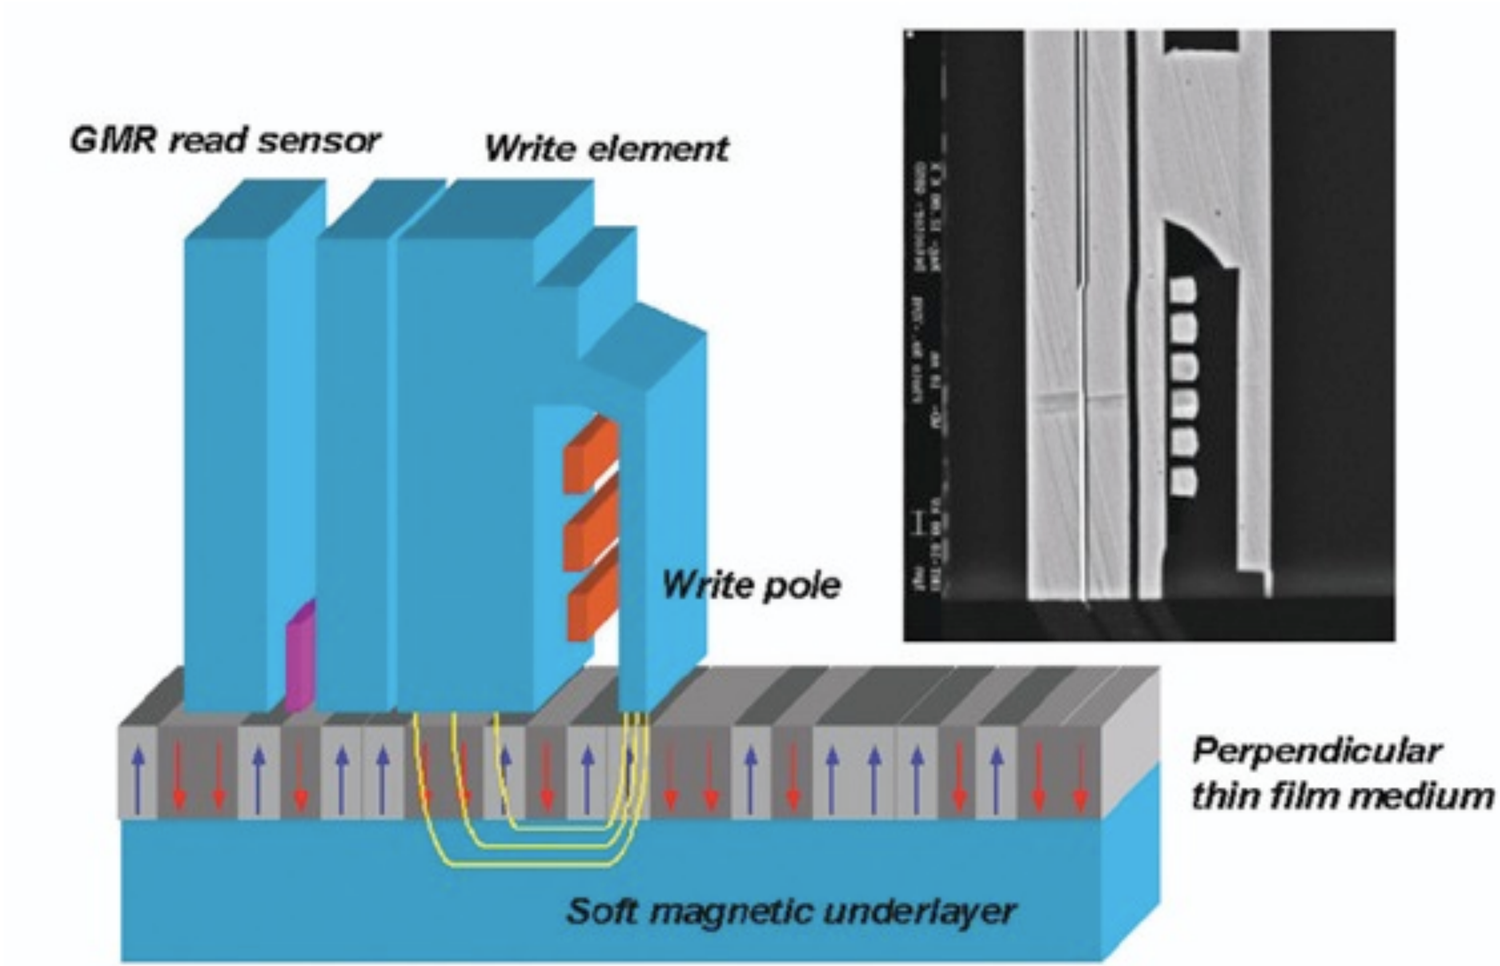

soft underlayer becomes part of the flux guide
... careful concentration of flux ...

read head (and its reflection)


wikipedia.org - “Hard_Disk”

positioning basics

- current powers voice coil†
- field generated moves head L or R
- more precise than stepper motor

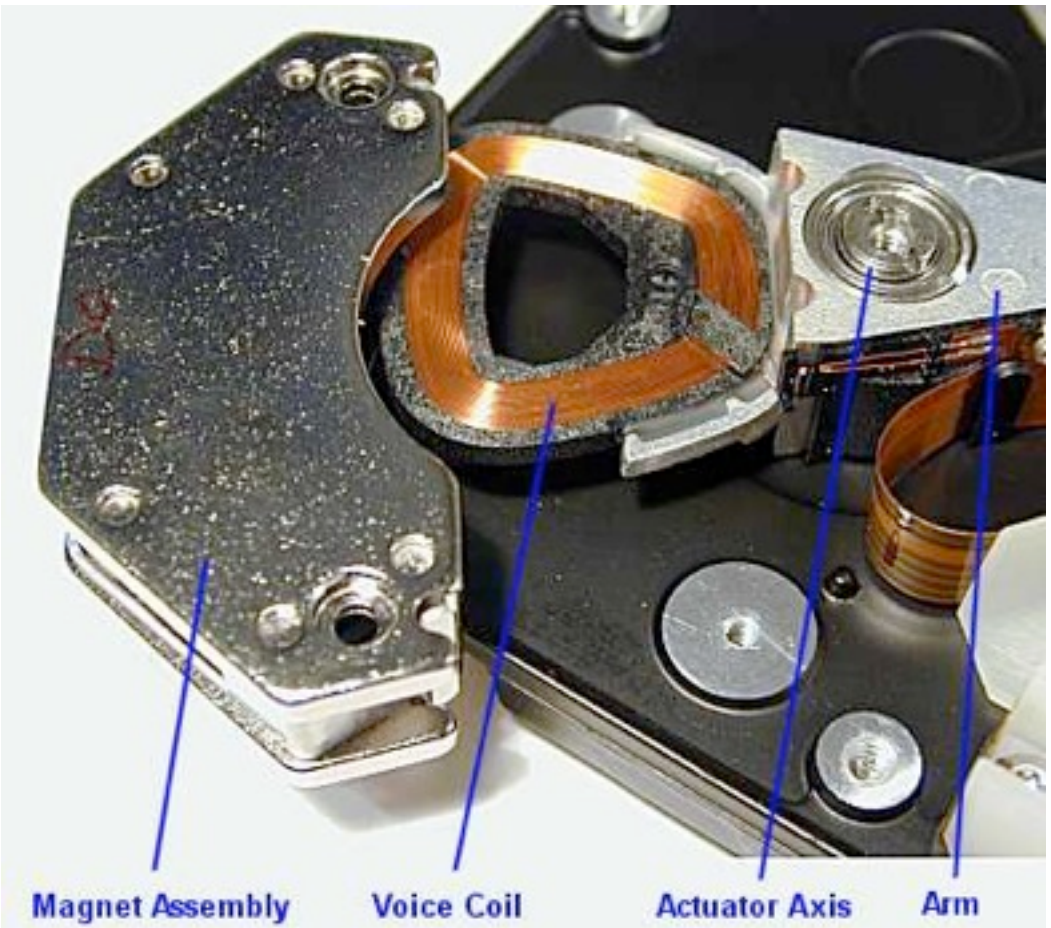

www.pcguide.com/ref/hdd/op/actActuator-c.html

† this is the same way a speaker cone moves

positioning basics

- current powers voice coil†
- field generated moves head L or R
- more precise than stepper motor

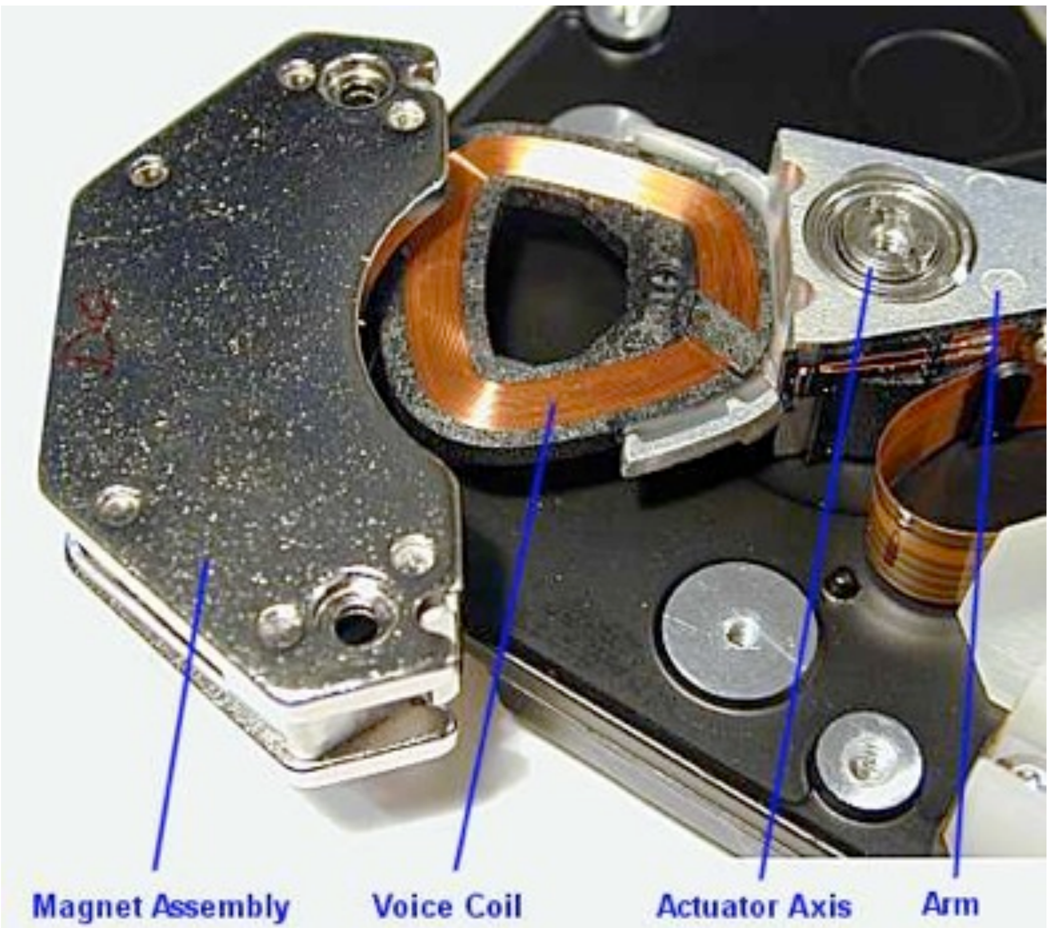

www.pcguide.com/ref/hdd/op/actActuator-c.html

IBM 62PC "Piccolo" HDD, ~1979 - an early 8" disk

wikipedia.org - “Hard_Disk”

† this is the same way a speaker cone moves

why magnets?

microscopic view

magnets remember their state

once magnetized, they stay that way

with a little bit of energy, we can control them

switch from N to S

why magnets?

microscopic view

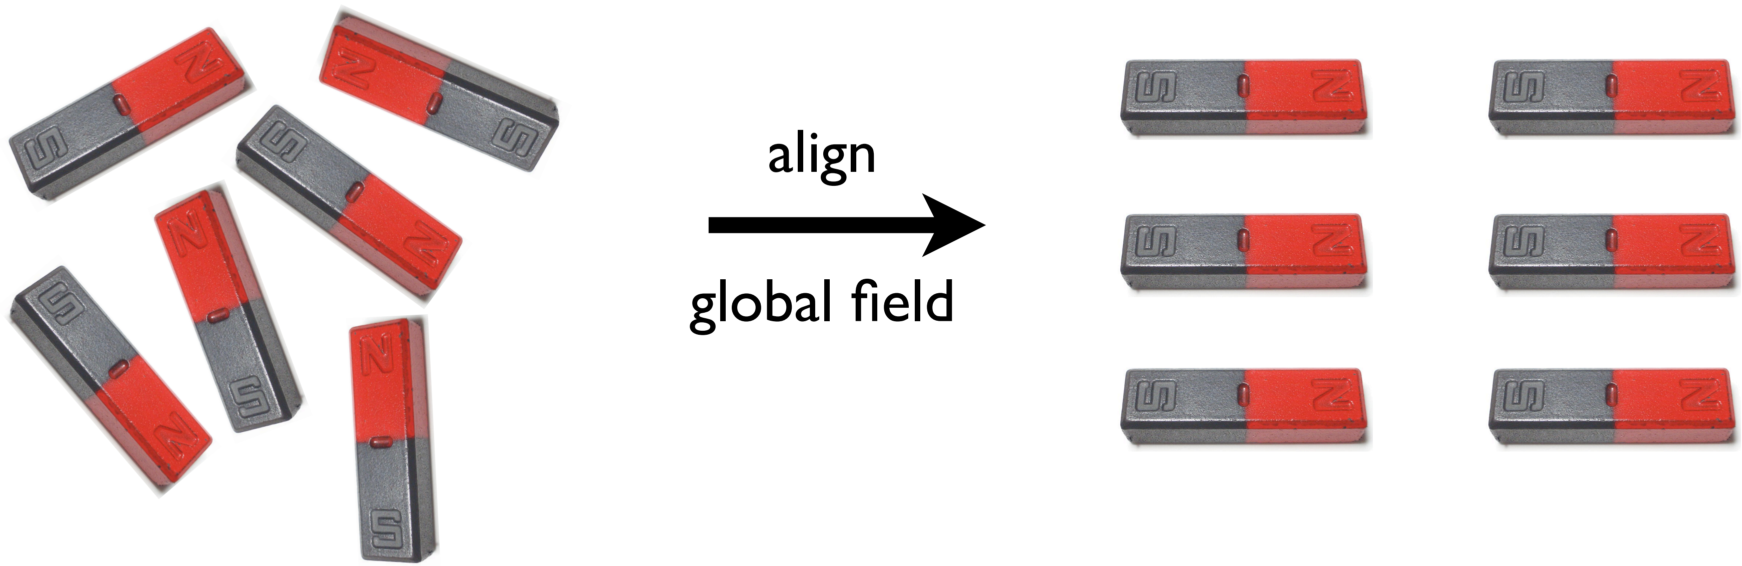

magnets remember their state

once magnetized, they stay that way

with a little bit of energy, we can control them

switch from N to S

why magnets?

microscopic view

magnets remember their state

once magnetized, they stay that way

with a little bit of energy, we can control them

switch from N to S

why magnets?

microscopic view

magnets remember their state

once magnetized, they stay that way

with a little bit of energy, we can control them

switch from N to S

why magnets?

what happens when you break a magnet?

you get two magnets

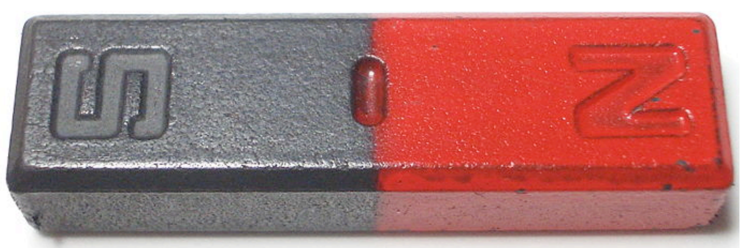

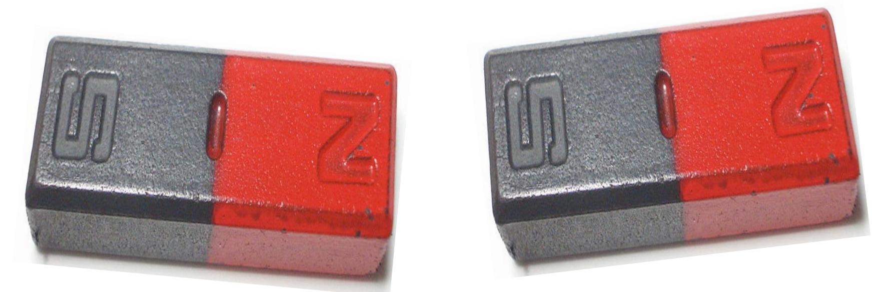

now: do this 25 more times

➠ 33 million magnets, all 50nm across

about 1,000 times thinner than a hair

we can make *really tiny* magnets

smaller is *better*, to a point

technology timeline

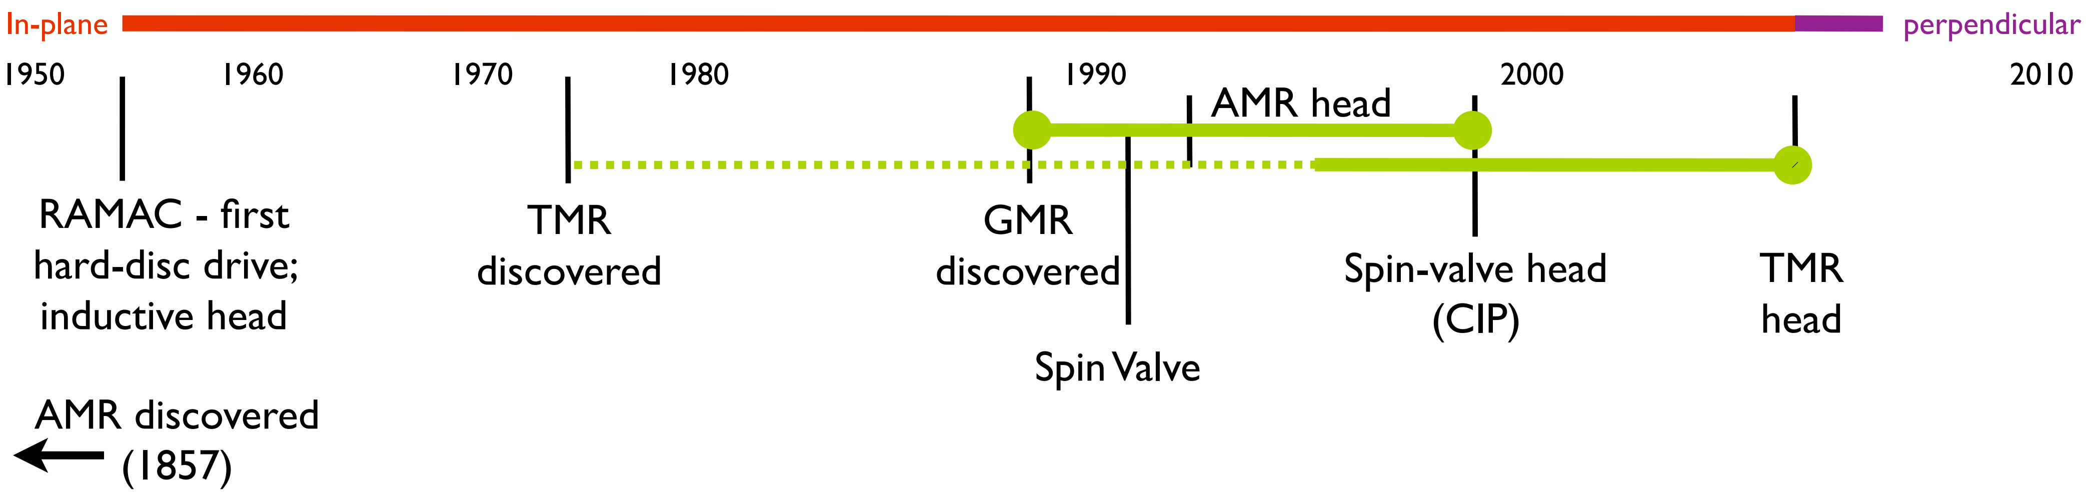

images and text from M. Coey

technology timeline

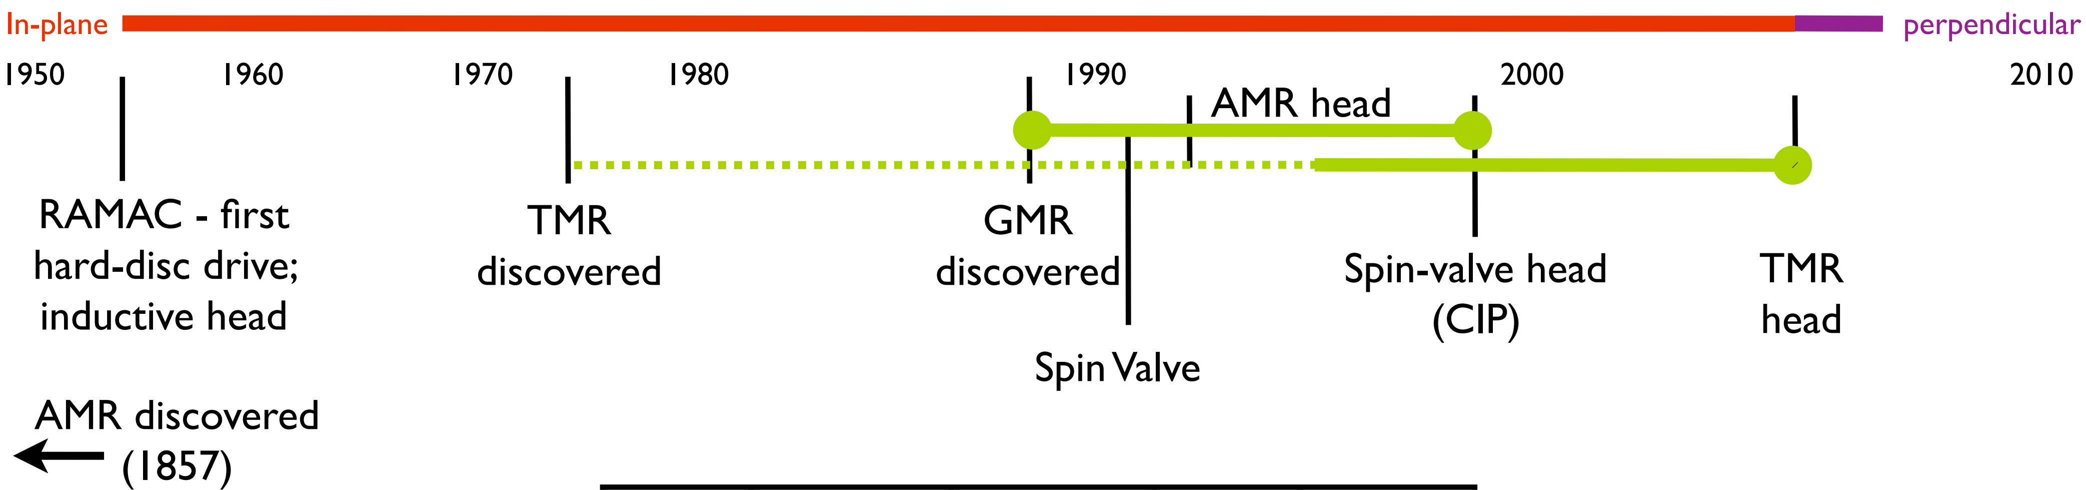

year	capacity	platters	size	rpm
1955	40 Mb	50x2	24”	1200
2005	160 Gb	1	2.5”	18000

images and text from M. Coey

the incredible shrinking hard disk

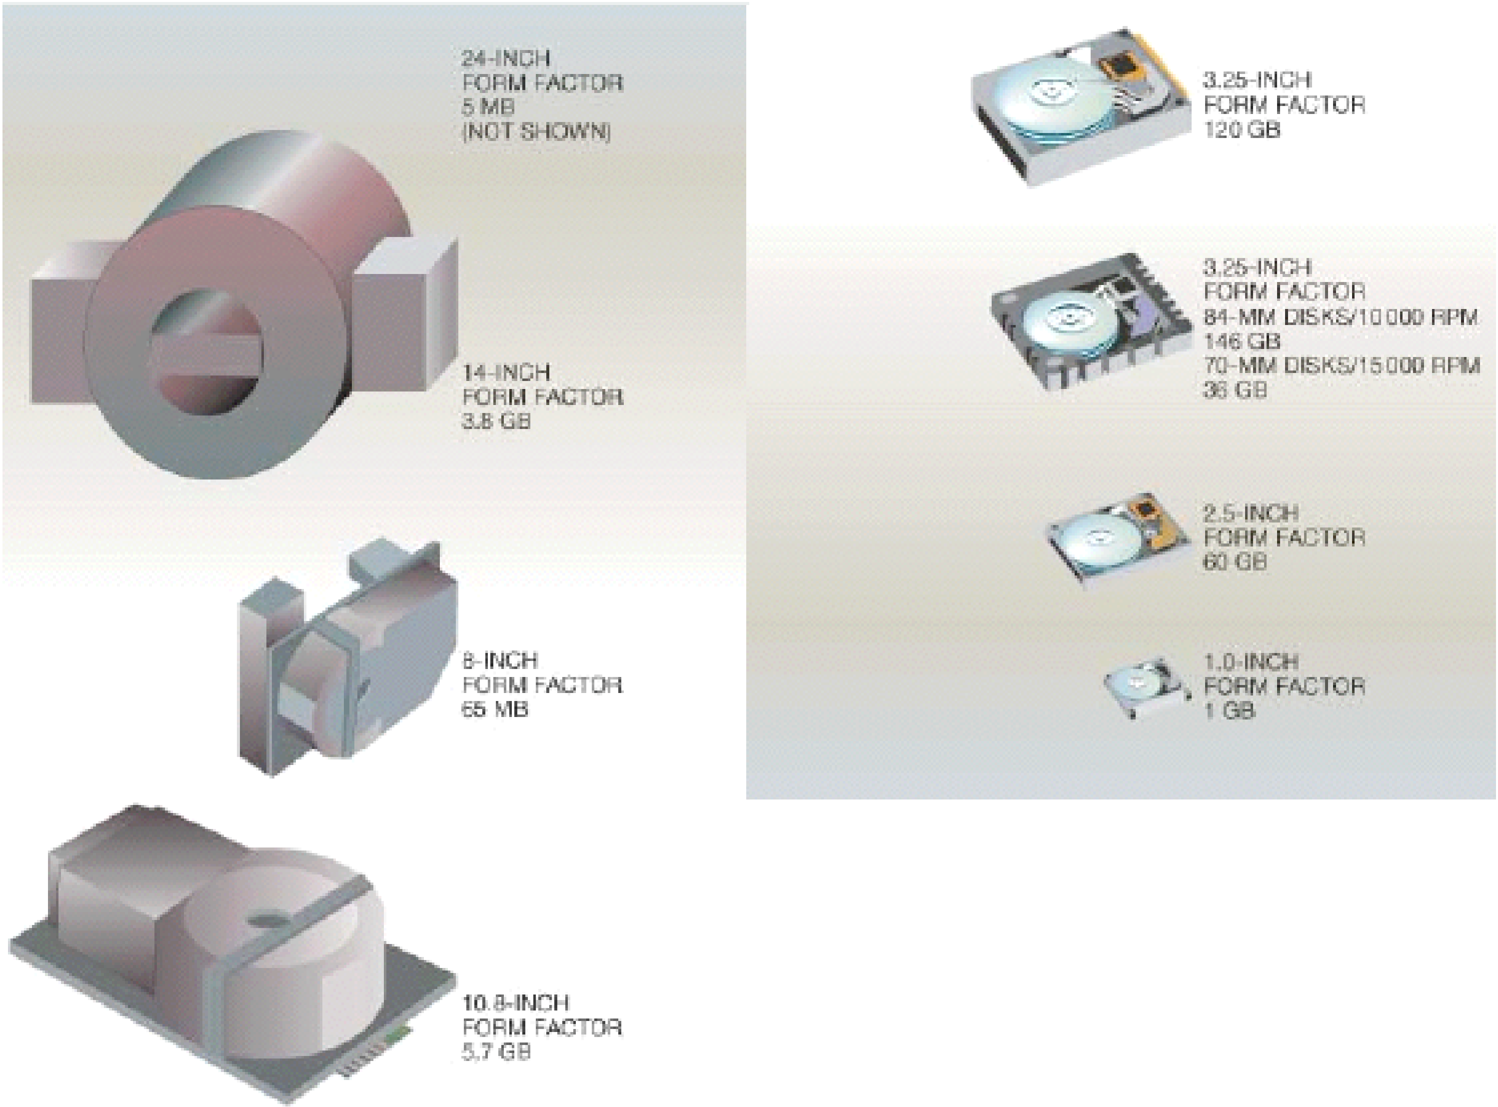

E. Grochowski and R.D. Halem
IBM Systems Journal, **42** (2), 2003

the incredible shrinking hard disk

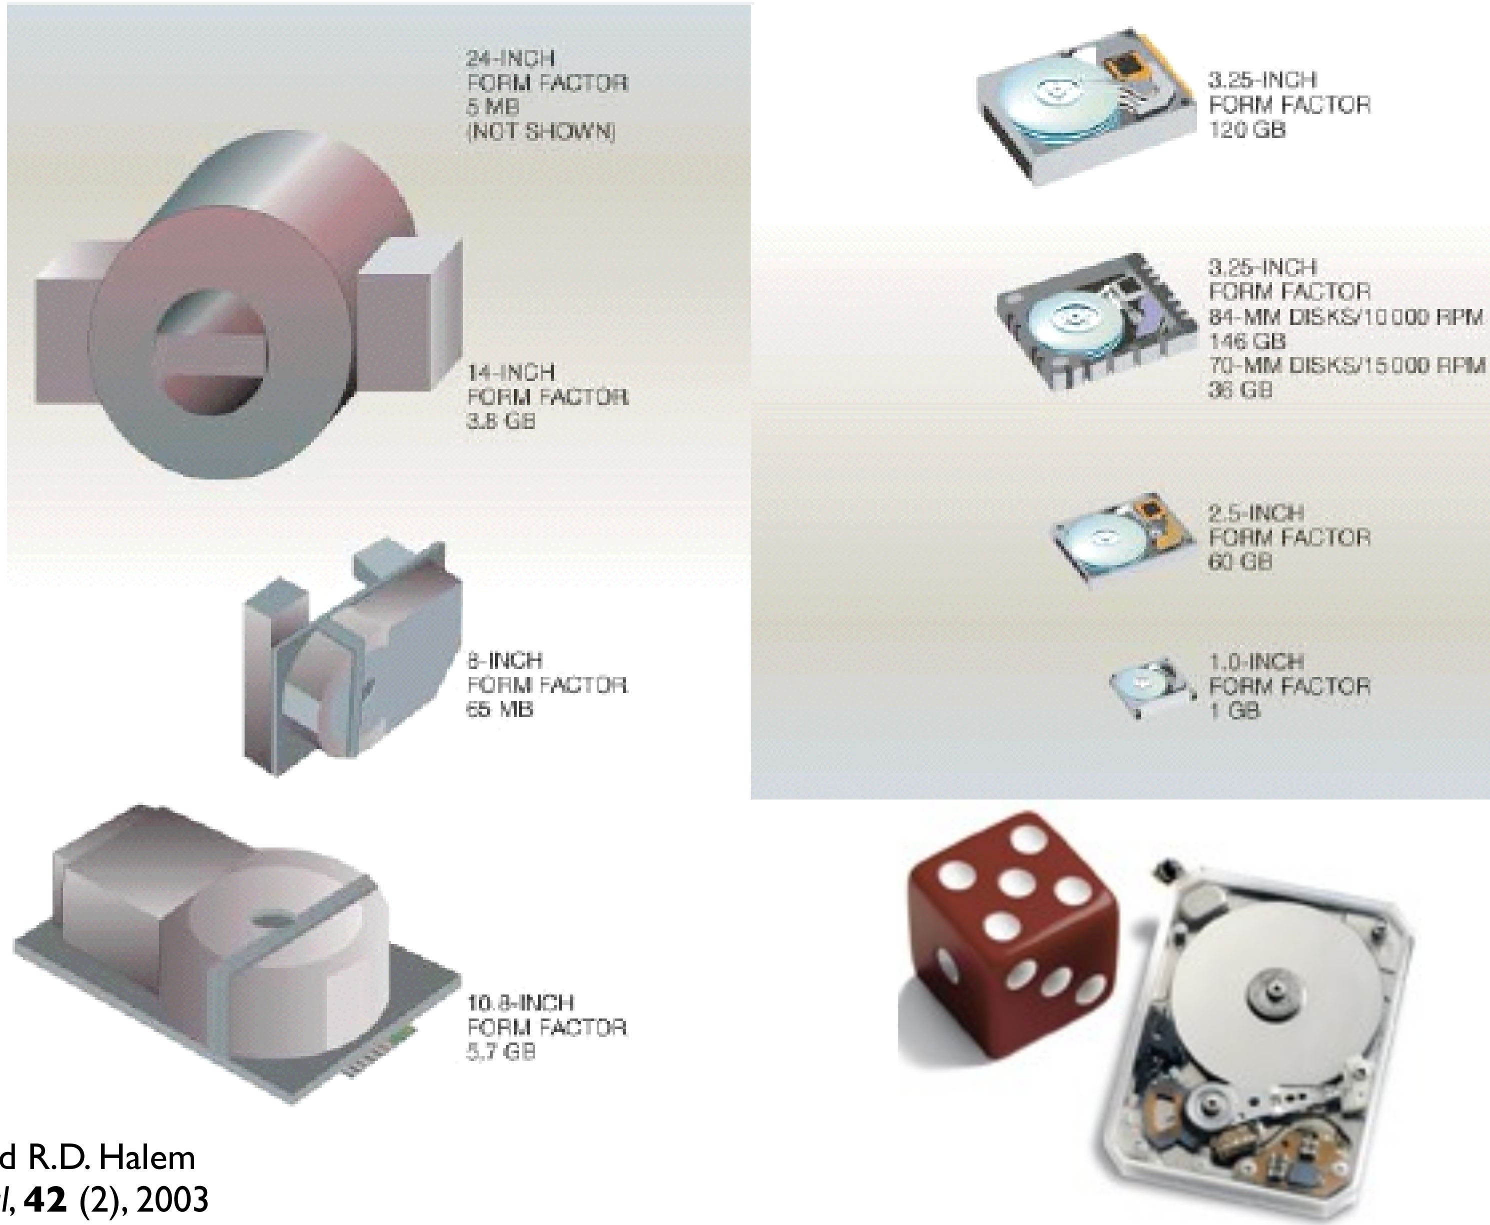

E. Grochowski and R.D. Halem
IBM Systems Journal, **42** (2), 2003

growth of areal density

E. Grochowski and R.D. Halem
IBM Systems Journal, **42** (2), 2003

growth of areal density

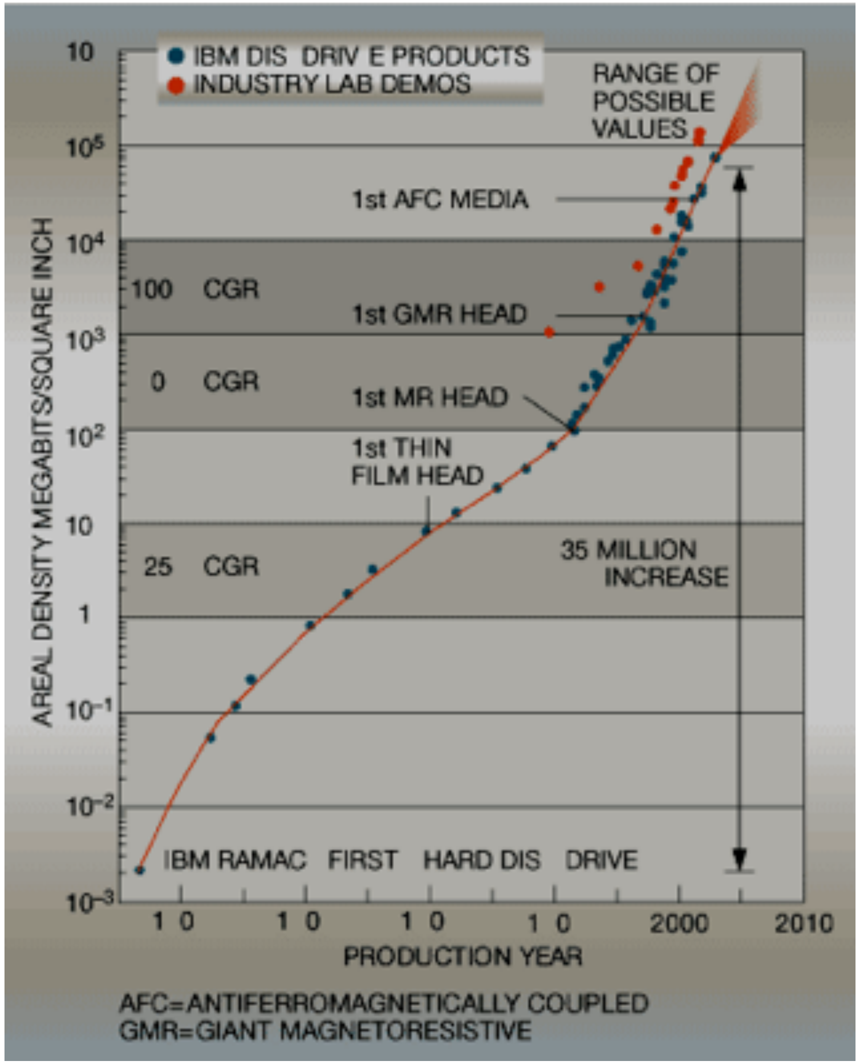

E. Grochowski and R.D. Halem
IBM Systems Journal, **42** (2), 2003

growth of areal density

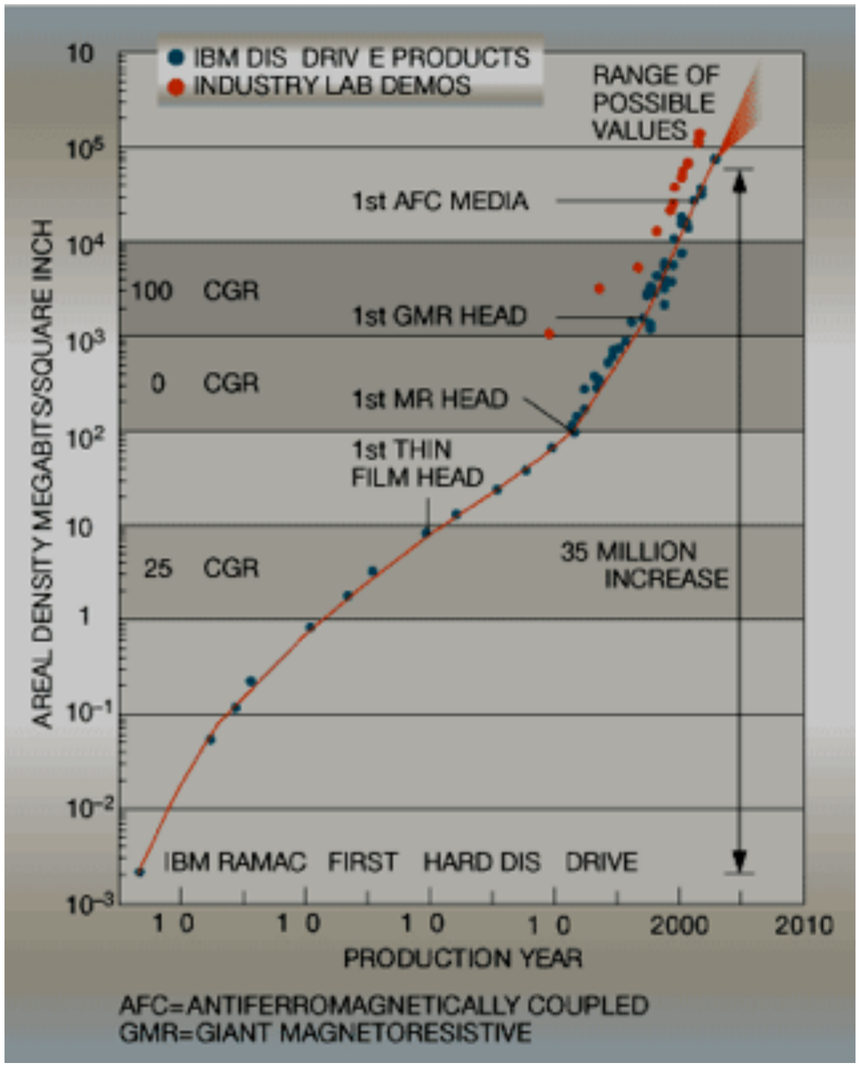

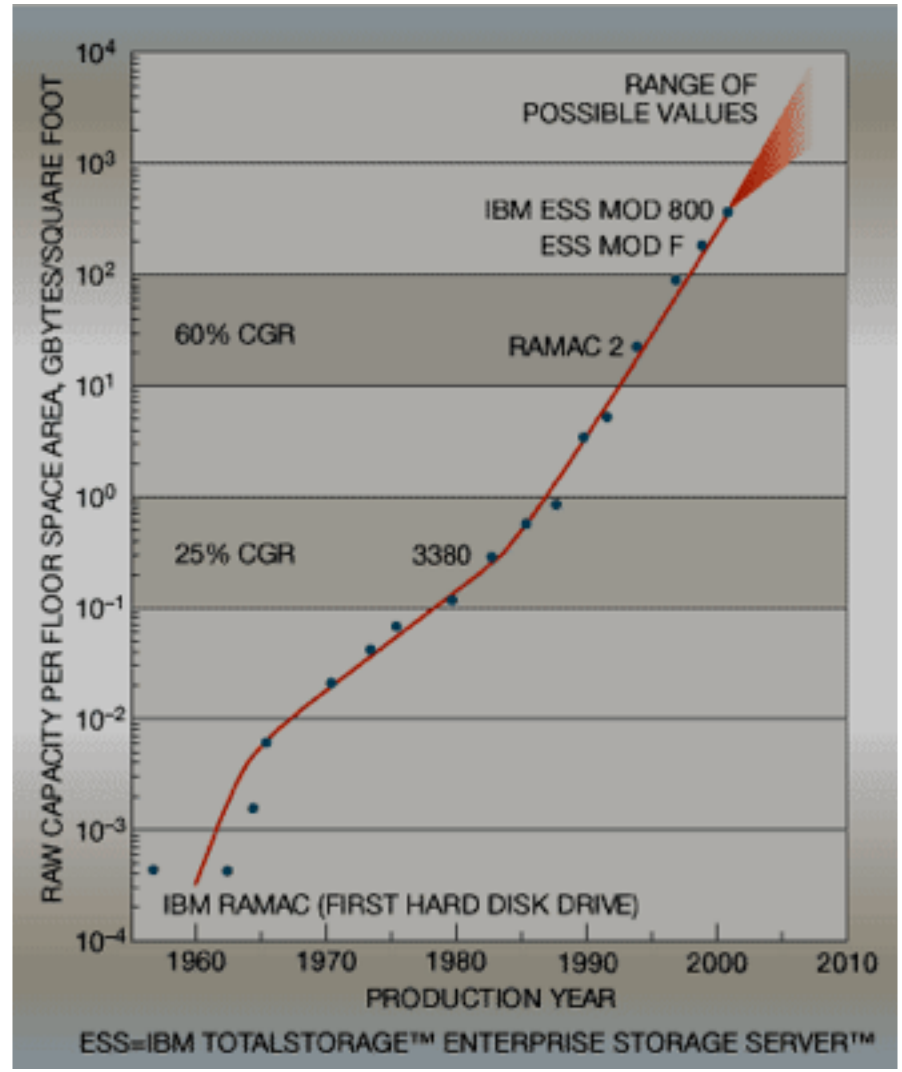

E. Grochowski and R.D. Halem
IBM Systems Journal, **42** (2), 2003

areal density *vs.* DRAM

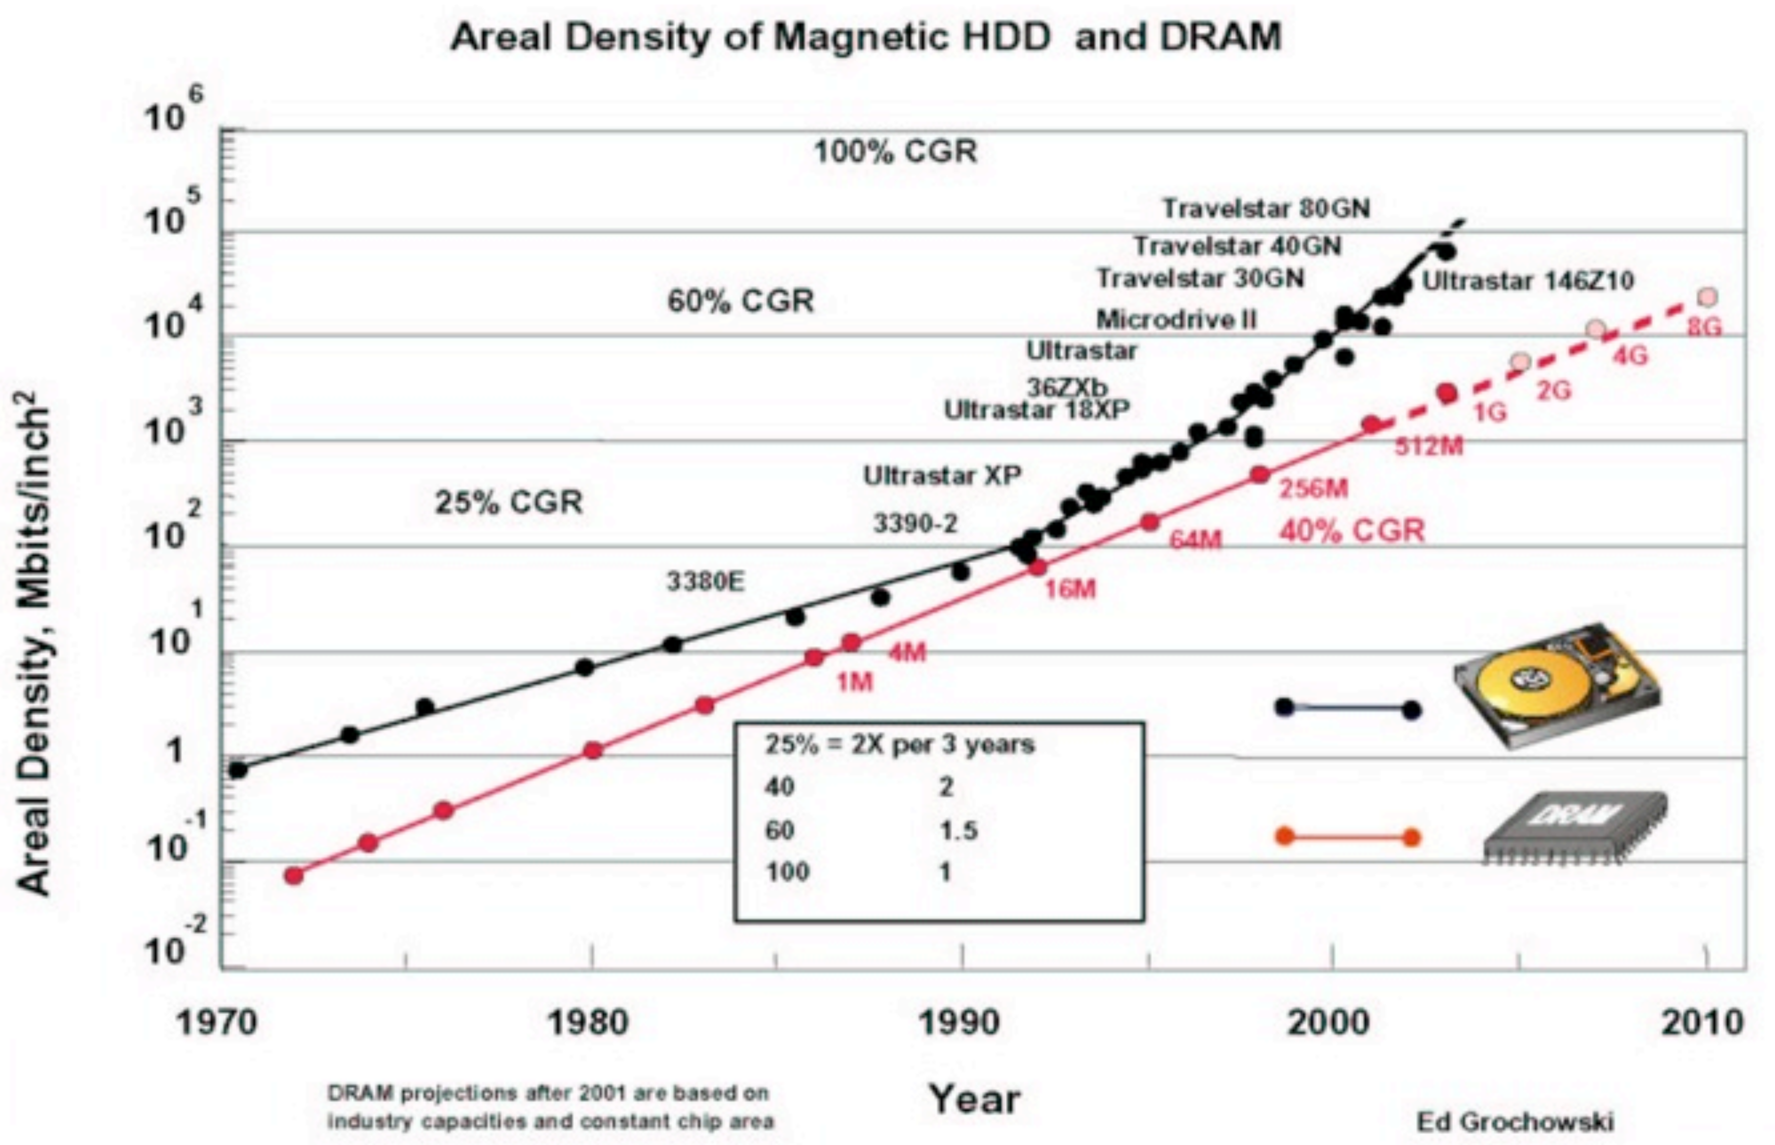

The incredible shrinking bit!
Predicted relative sizes of HDD storage bits

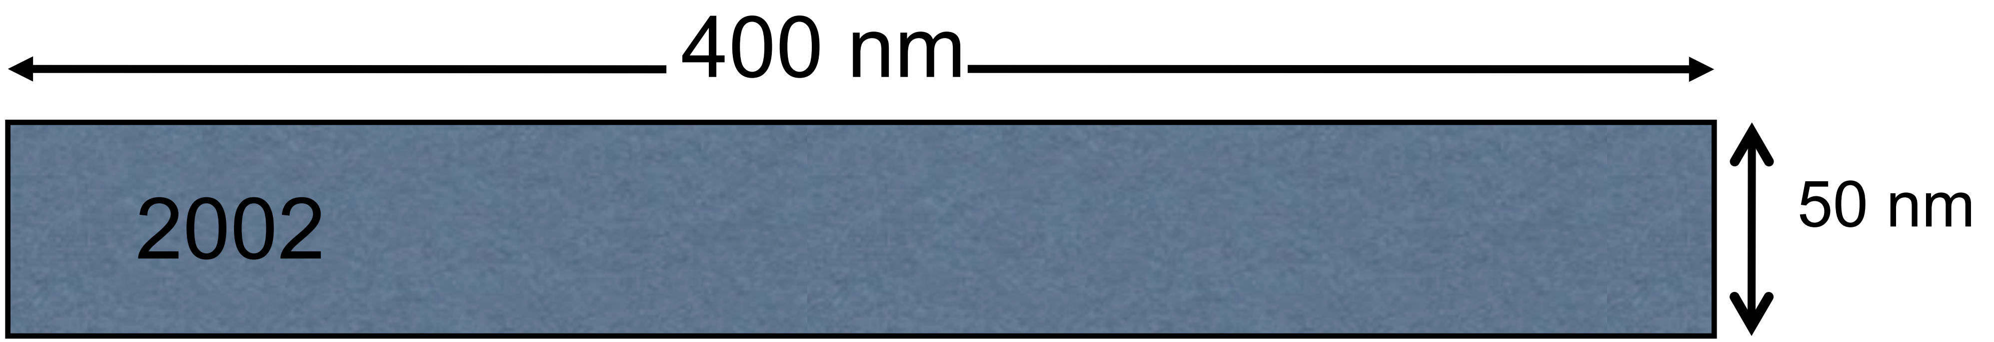

(assumes areal densities continue to double yearly)

The incredible shrinking bit!
Predicted relative sizes of HDD storage bits

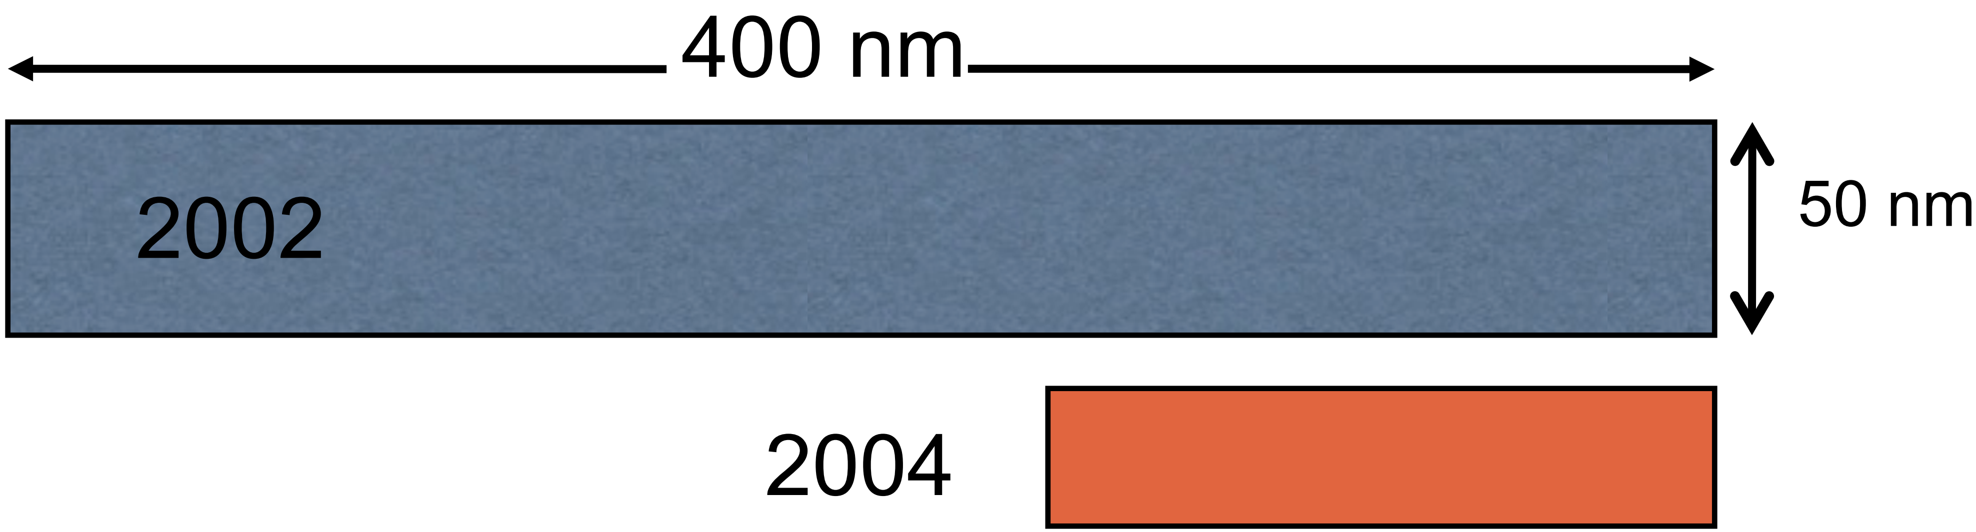

(assumes areal densities continue to double yearly)

The incredible shrinking bit!
Predicted relative sizes of HDD storage bits

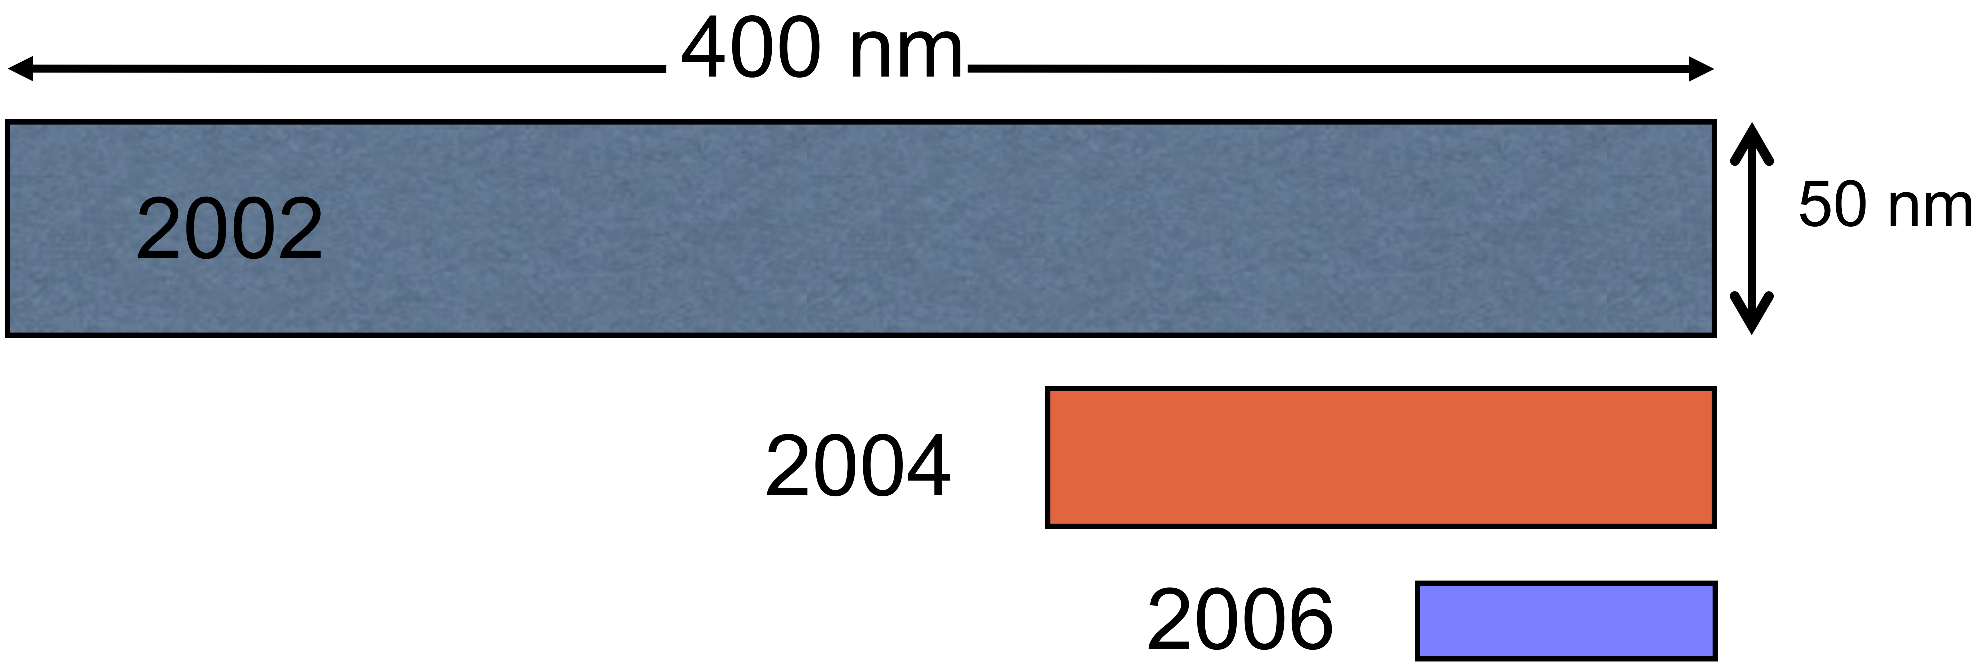

(assumes areal densities continue to double yearly)

The incredible shrinking bit!
Predicted relative sizes of HDD storage bits

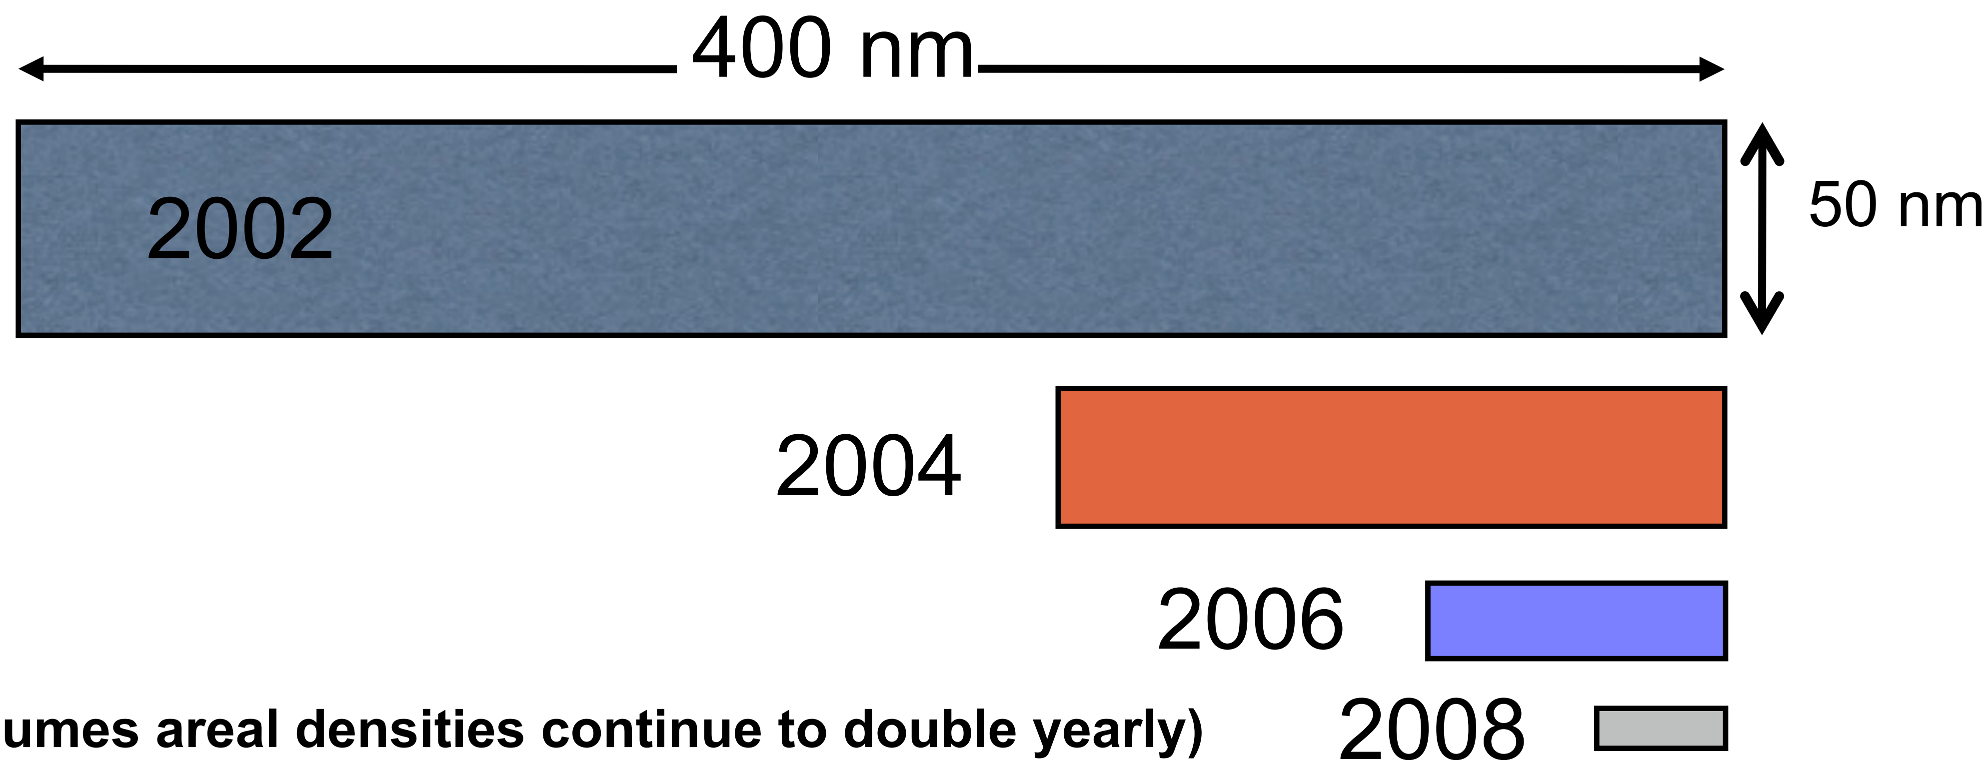

(assumes areal densities continue to double yearly)

The incredible shrinking bit!
Predicted relative sizes of HDD storage bits

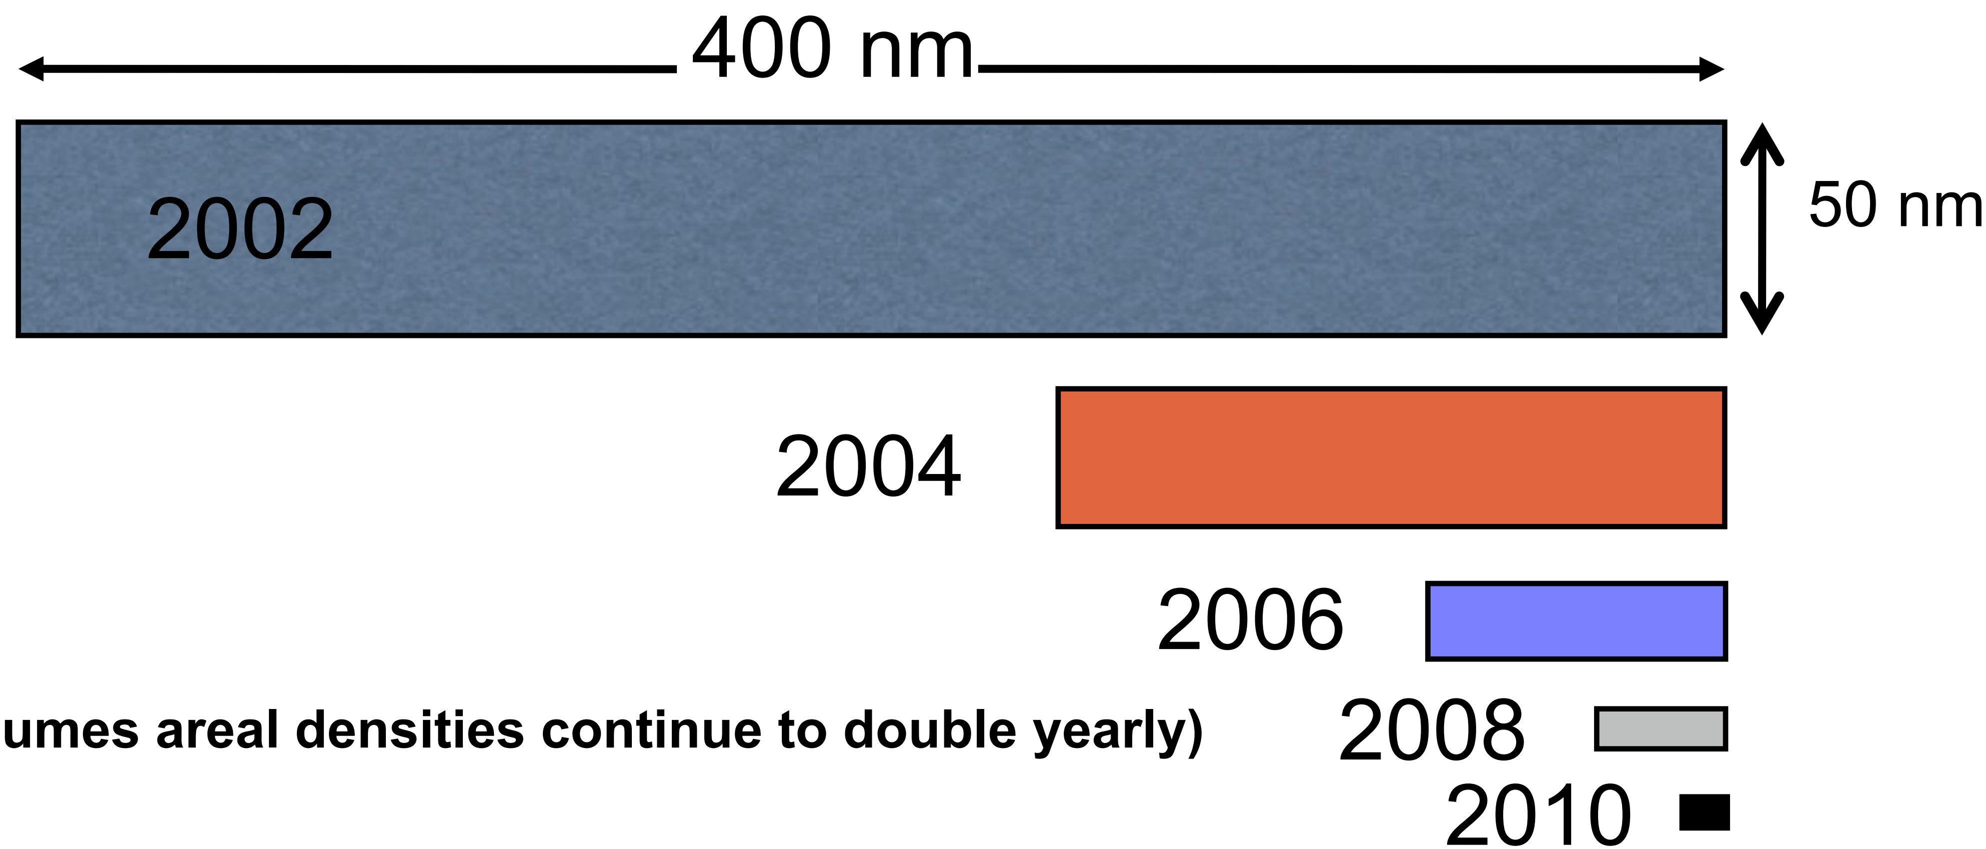

(assumes areal densities continue to double yearly)

The incredible shrinking bit!
Predicted relative sizes of HDD storage bits

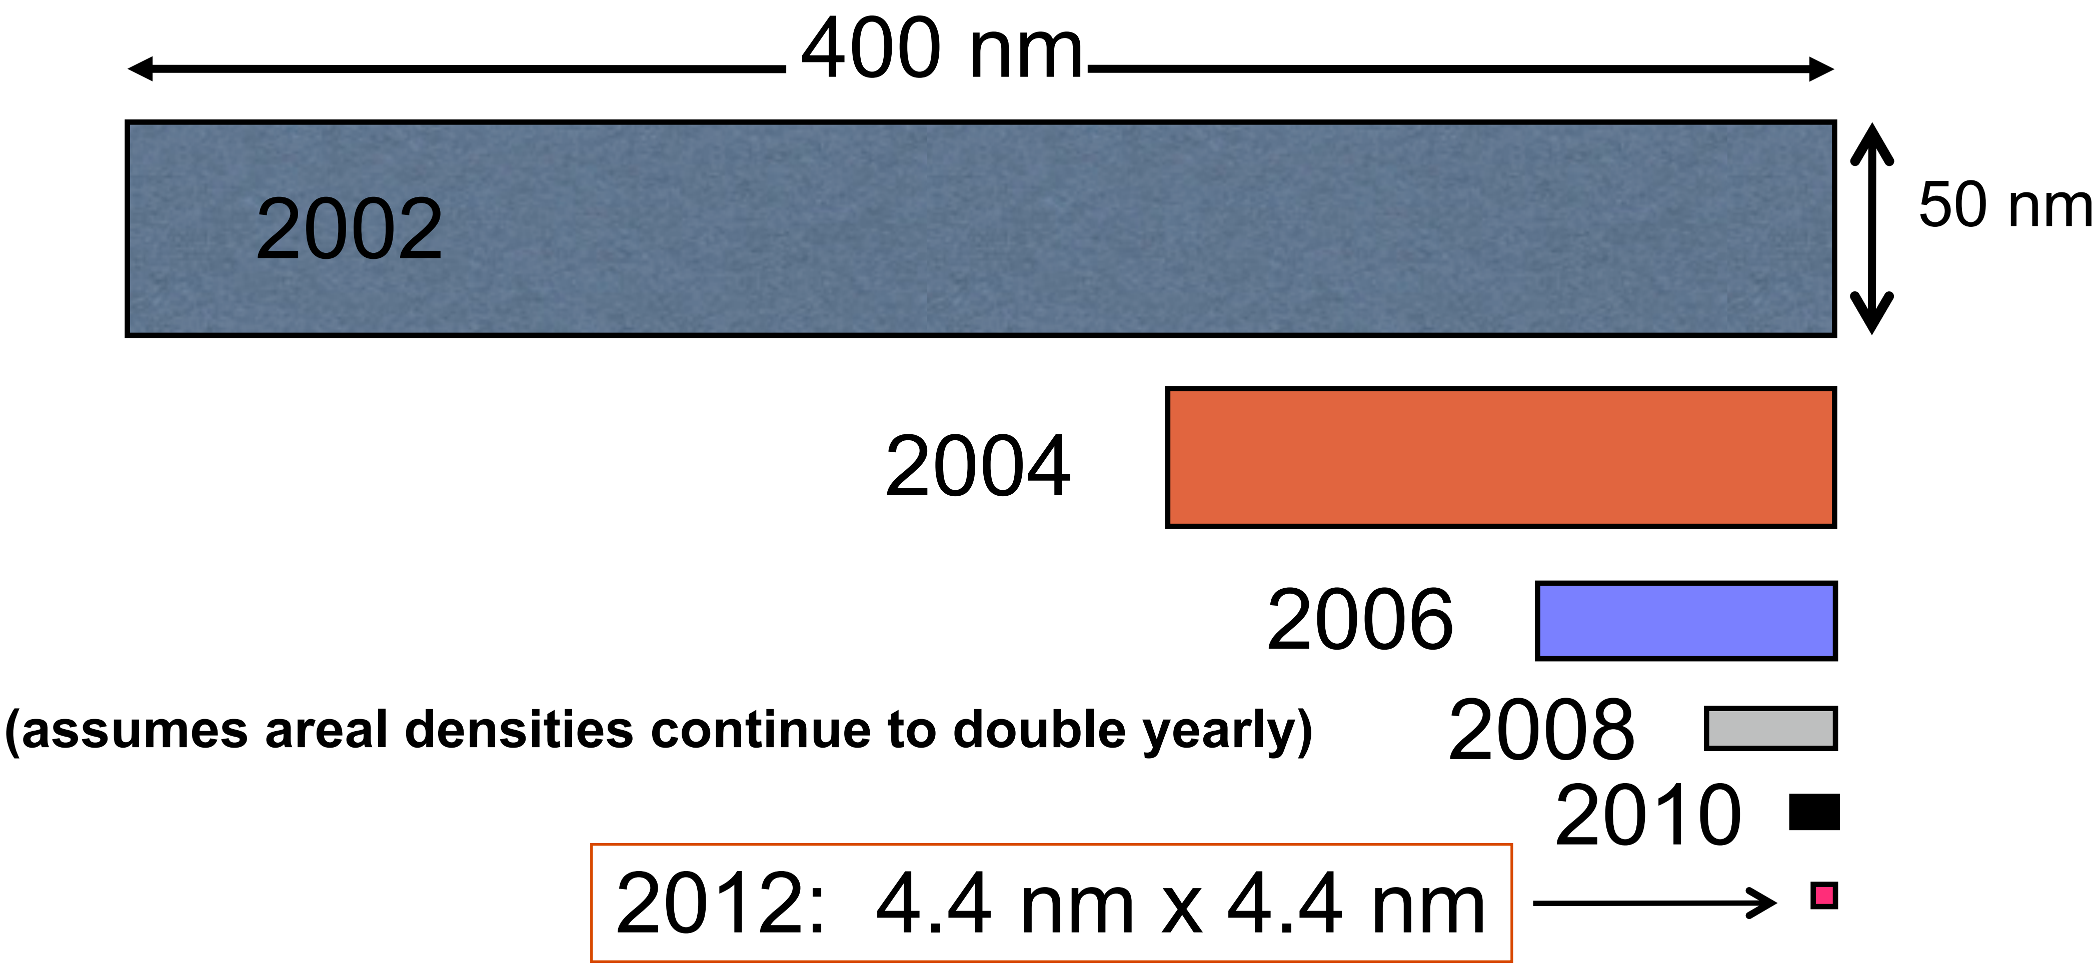

The incredible shrinking bit!
Predicted relative sizes of HDD storage bits

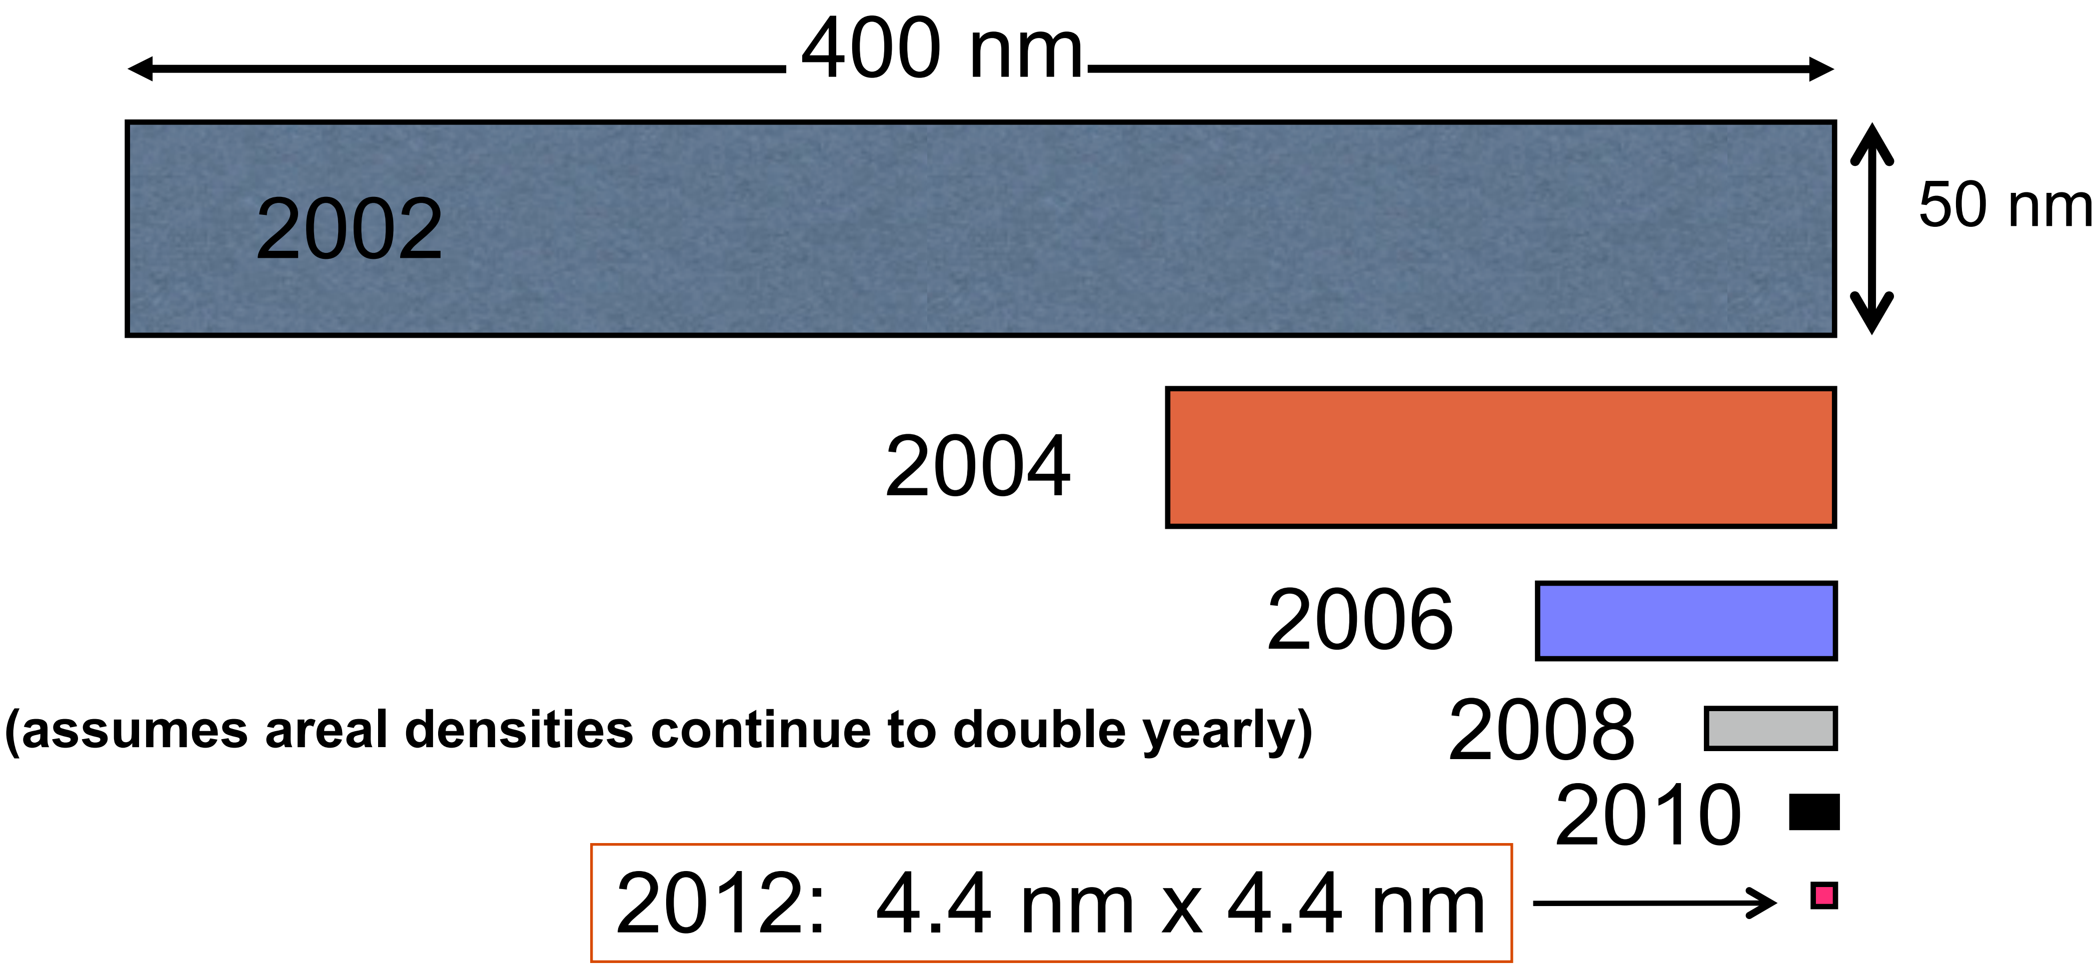

this will probably not happen

50 TB per square inch on a quarter ...

- over 3.4 million high-resolution photos, or …
- 2,800 audio CDs, or…
- 1,600 hours of television, or …

50 TB per square inch on a quarter ...

- over 3.4 million high-resolution photos, or …
- 2,800 audio CDs, or…
- 1,600 hours of television, or …

- the entire printed collection of the U.S. Library of Congress

Library of Congress, Jefferson building

so what's the problem?

at some point, they are no longer stable

heat makes them 'wiggle'
    like drops of water on a griddle

bits are no longer reliable

so we need stronger magnets ...
    ... which need more field to magnetize
    ... which needs more power

HUGE challenge in nanoscale materials science!

1 bit needs $k_B T$ ...

$$$ vs flash and DRAM

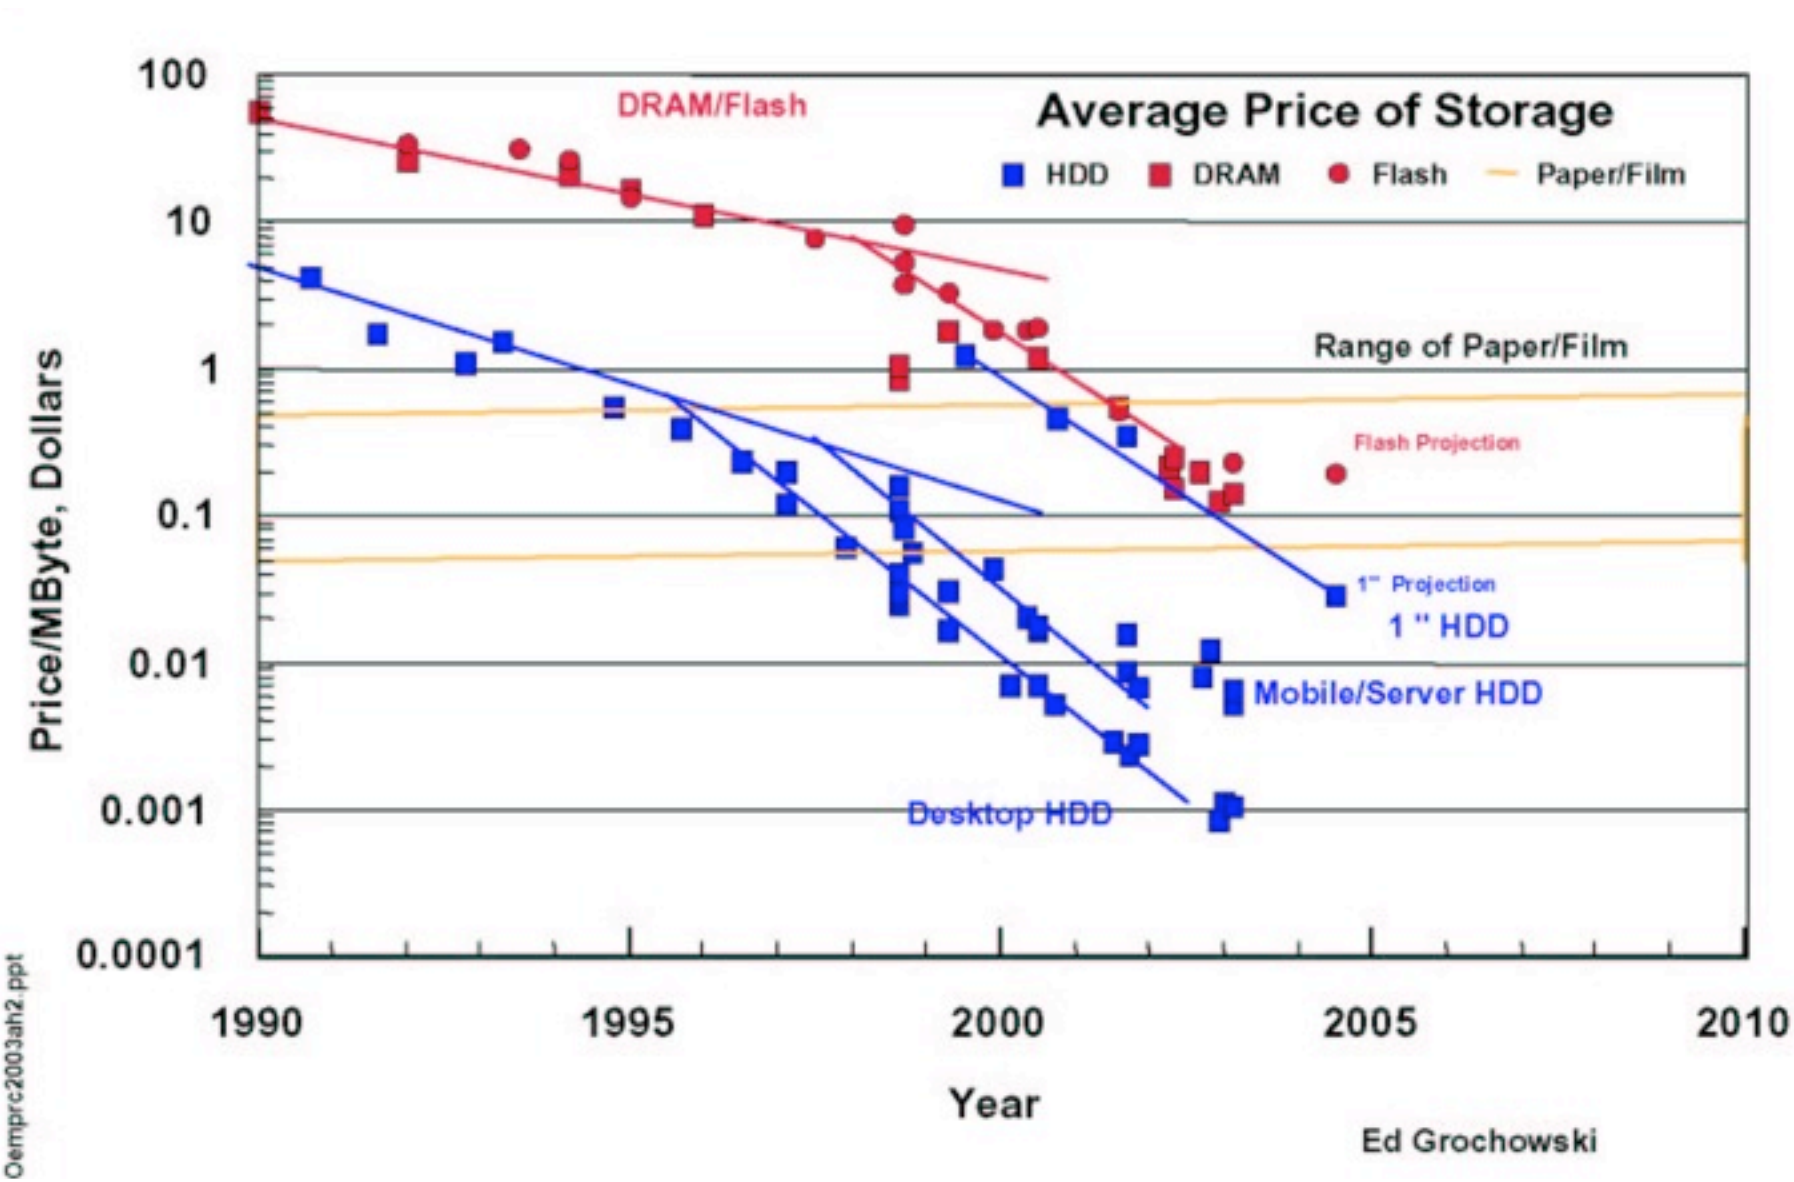

$$$ vs flash and DRAM

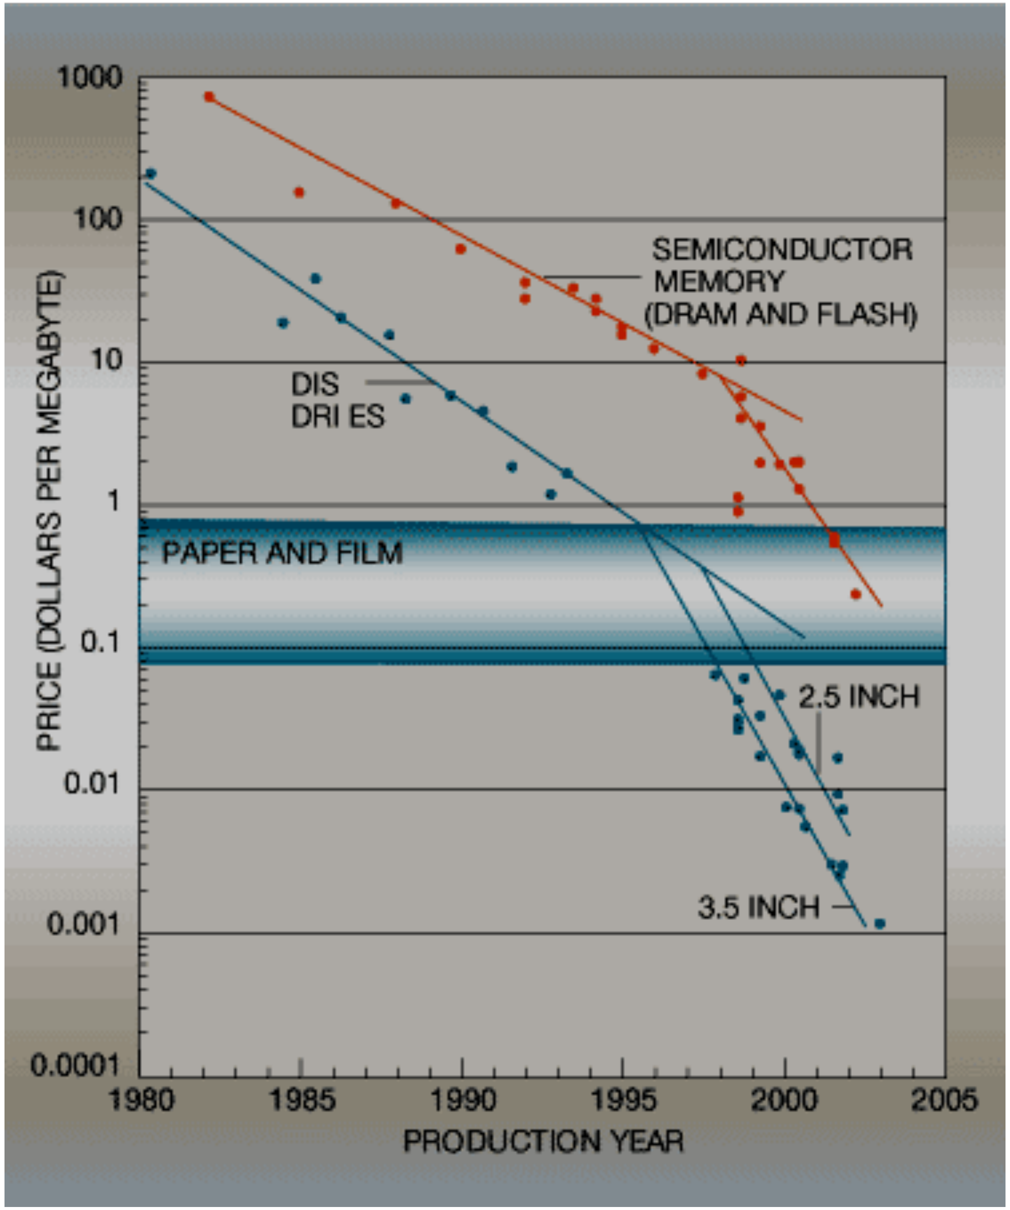

E. Grochowski and R.D. Halem
IBM Systems Journal, **42** (2), 2003

seek time

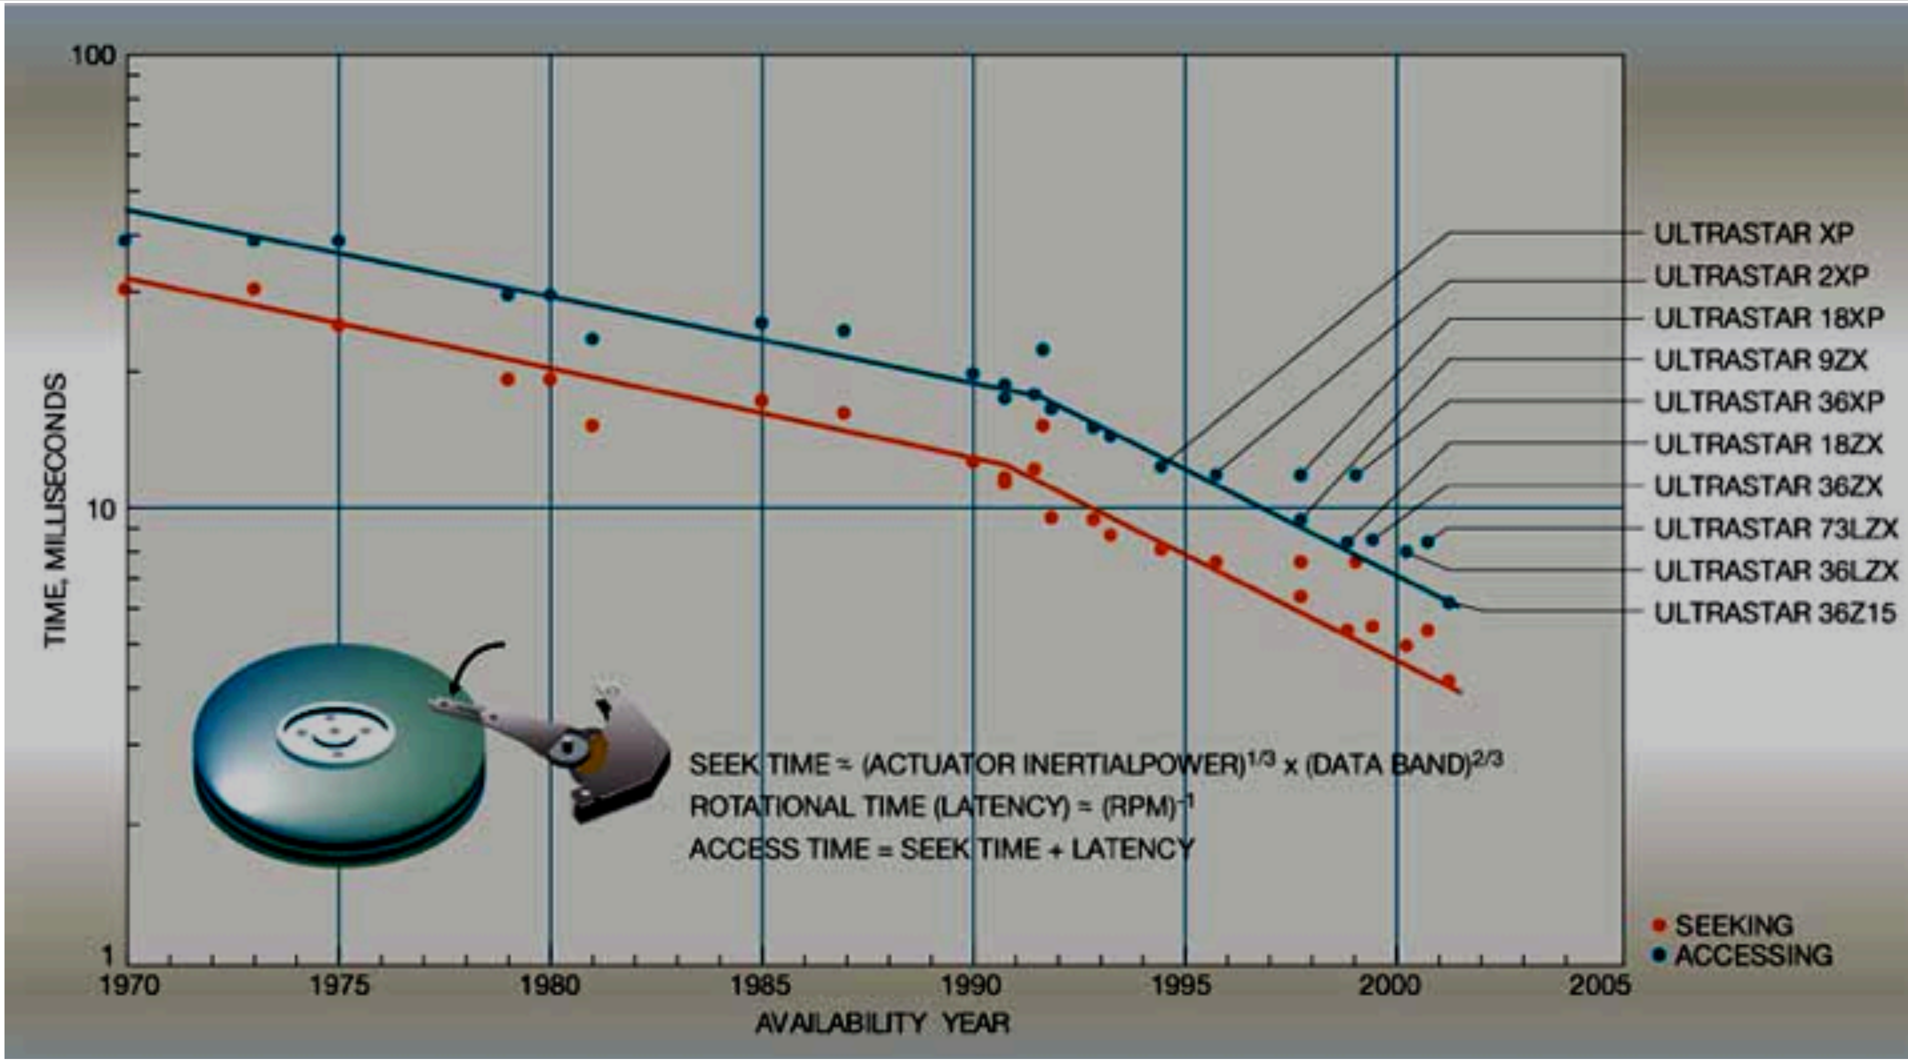

E. Grochowski and R.D. Halem
IBM Systems Journal, **42** (2), 2003

seek time

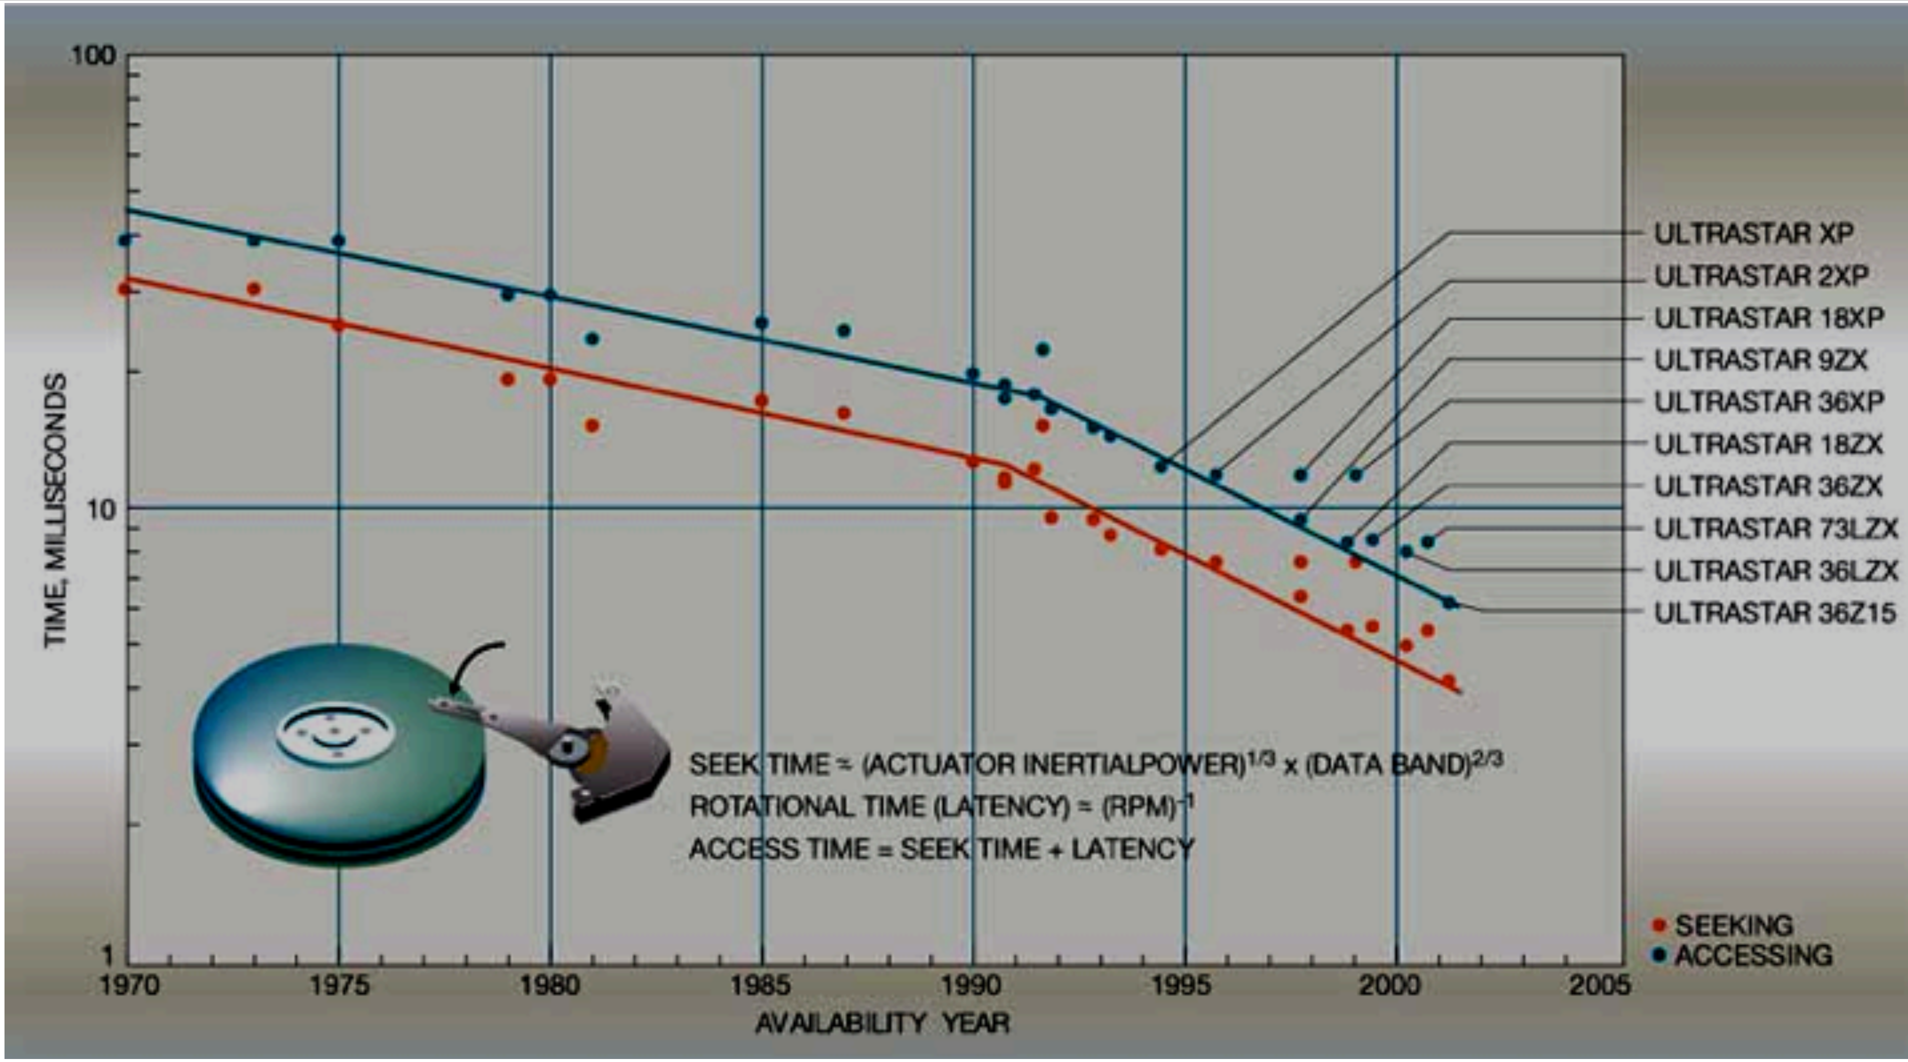

E. Grochowski and R.D. Halem
IBM Systems Journal, **42** (2), 2003

far higher than DRAM/SRAM (~nsec)
reduction limited by mechanics!

throughput

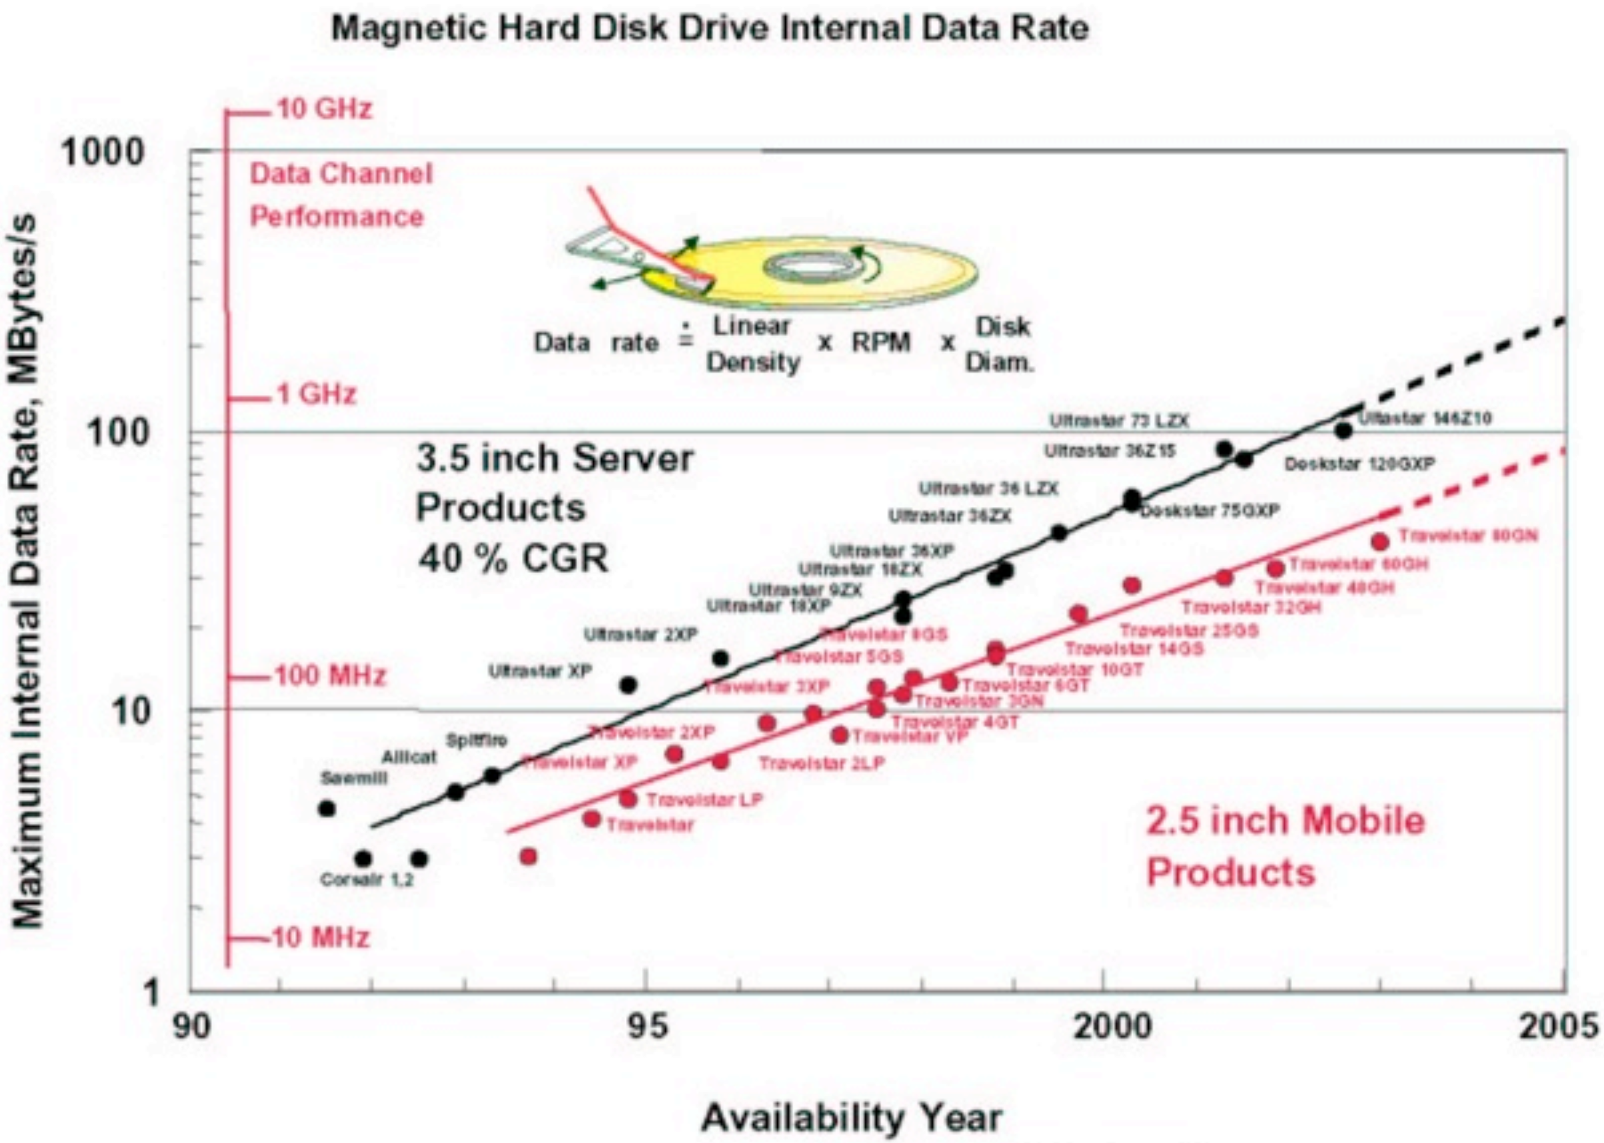

Magnetic Hard Disk Drive Internal Data Rate
Maximum Internal Data Rate, MBytes/s
Availability Year
Data Channel Performance
10 GHz
1 GHz
100 MHz
10 MHz
Data rate = Linear Density x RPM x Disk Diam.
3.5 inch Server Products 40 % CGR
2.5 inch Mobile Products
1000
100
10
1
90
95
2000
2005

throughput

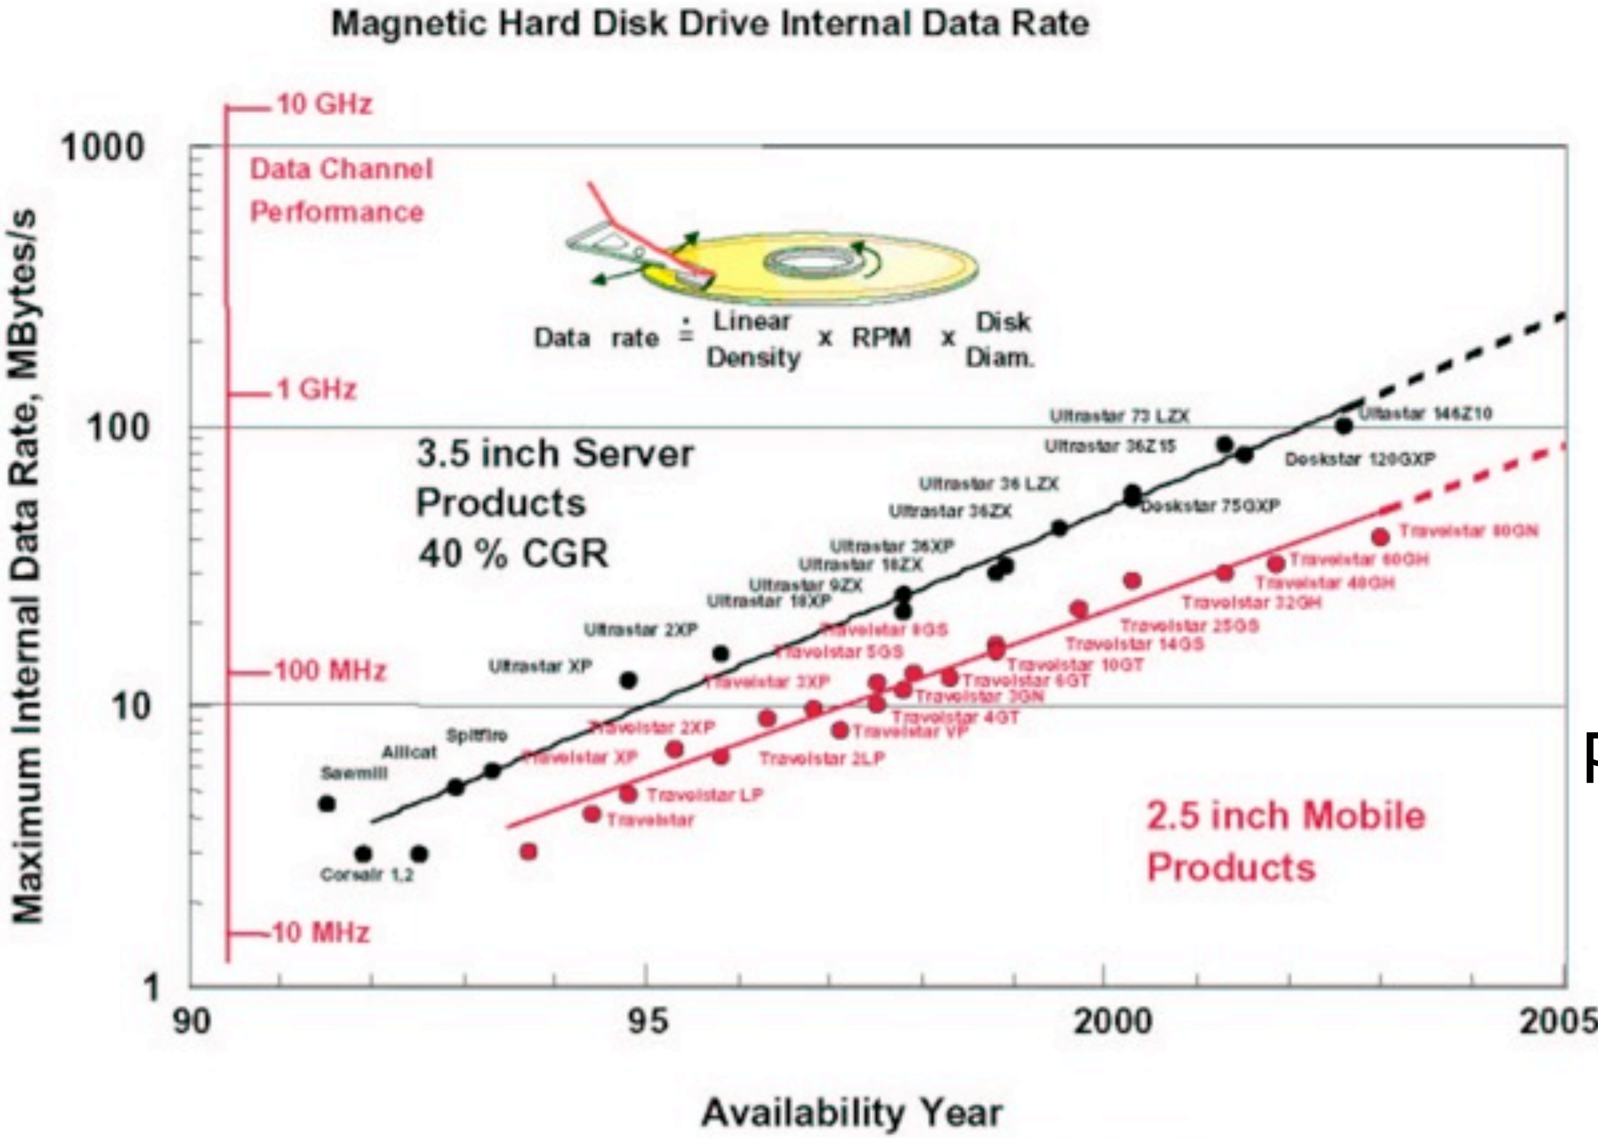

areal density growing *faster* than data rate

problem, or no?

density metrics

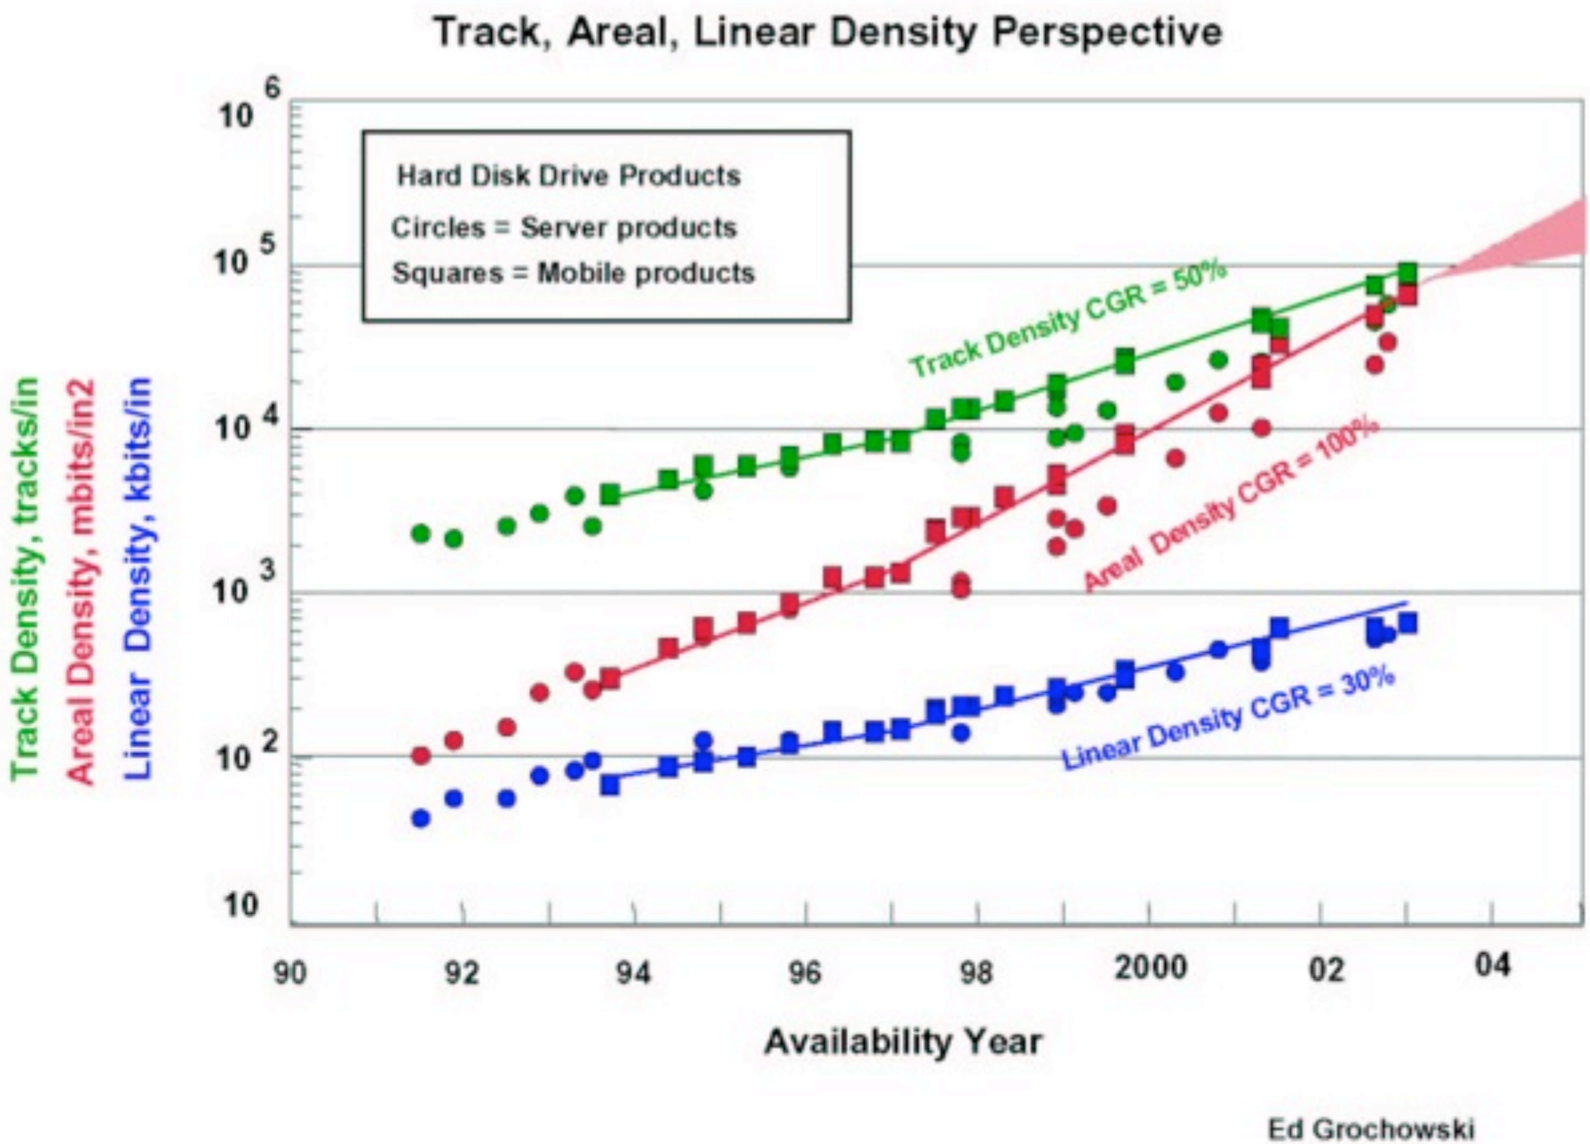

density metrics

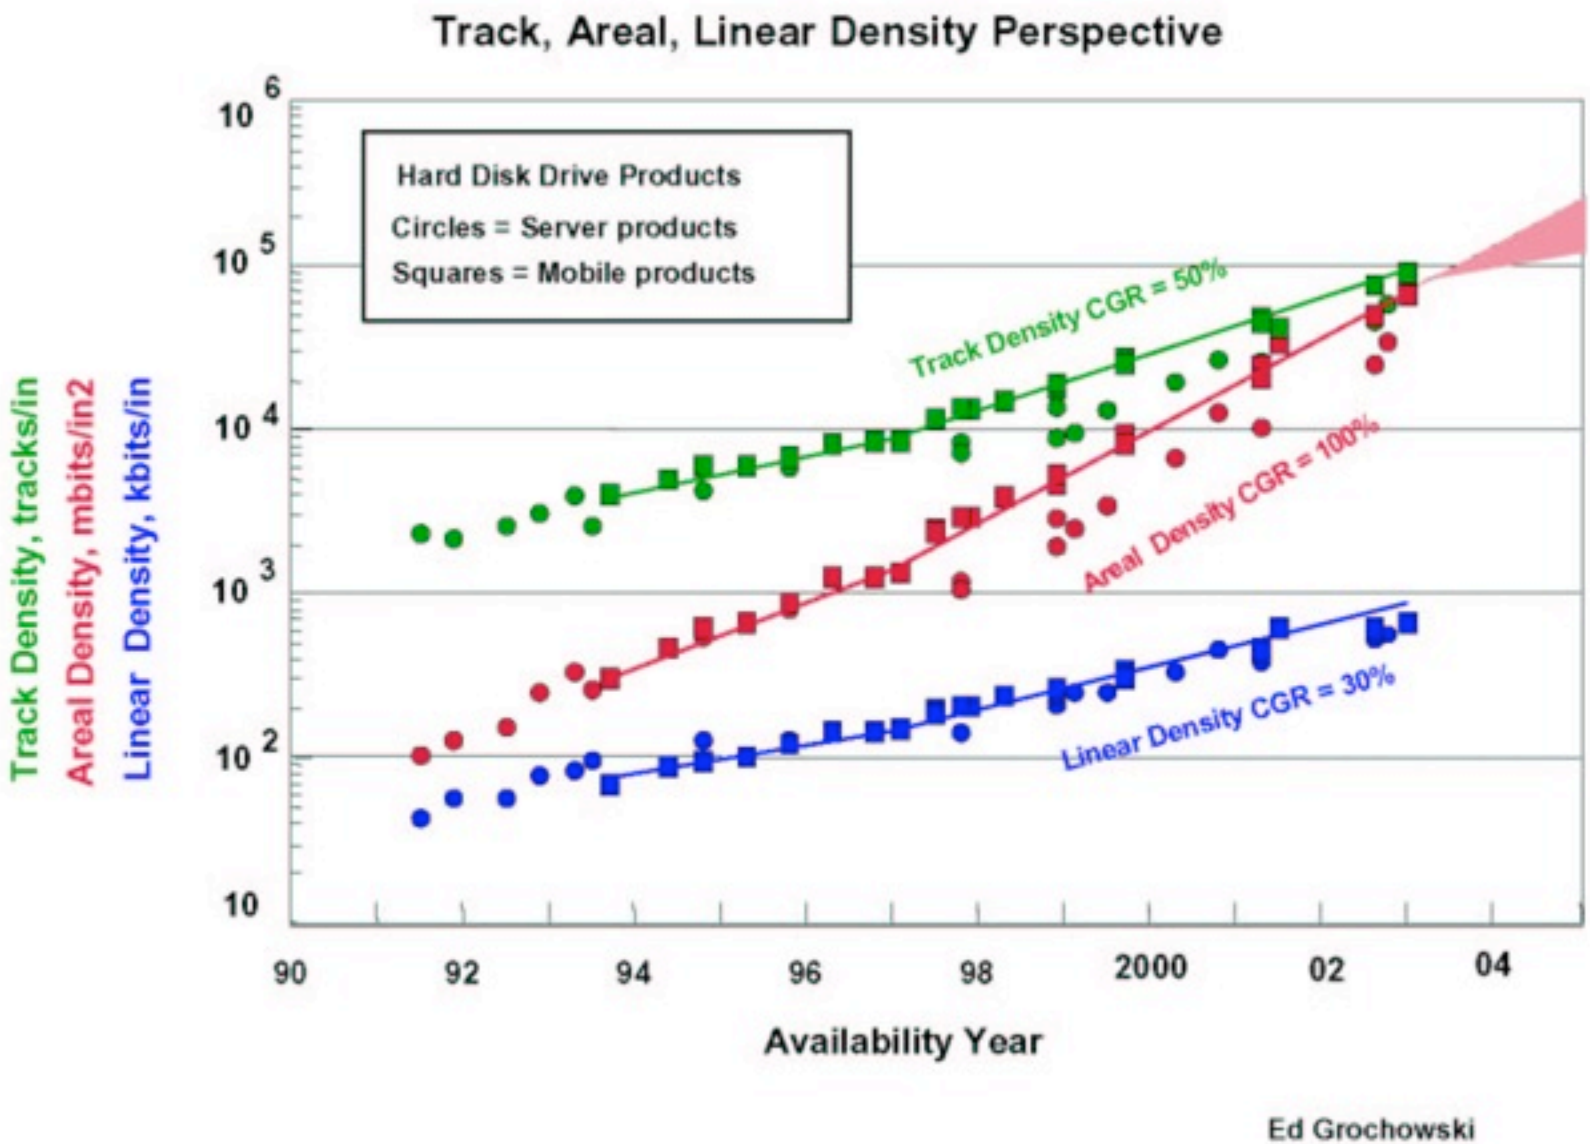

this is why we quote *areal* density.

head-media spacing

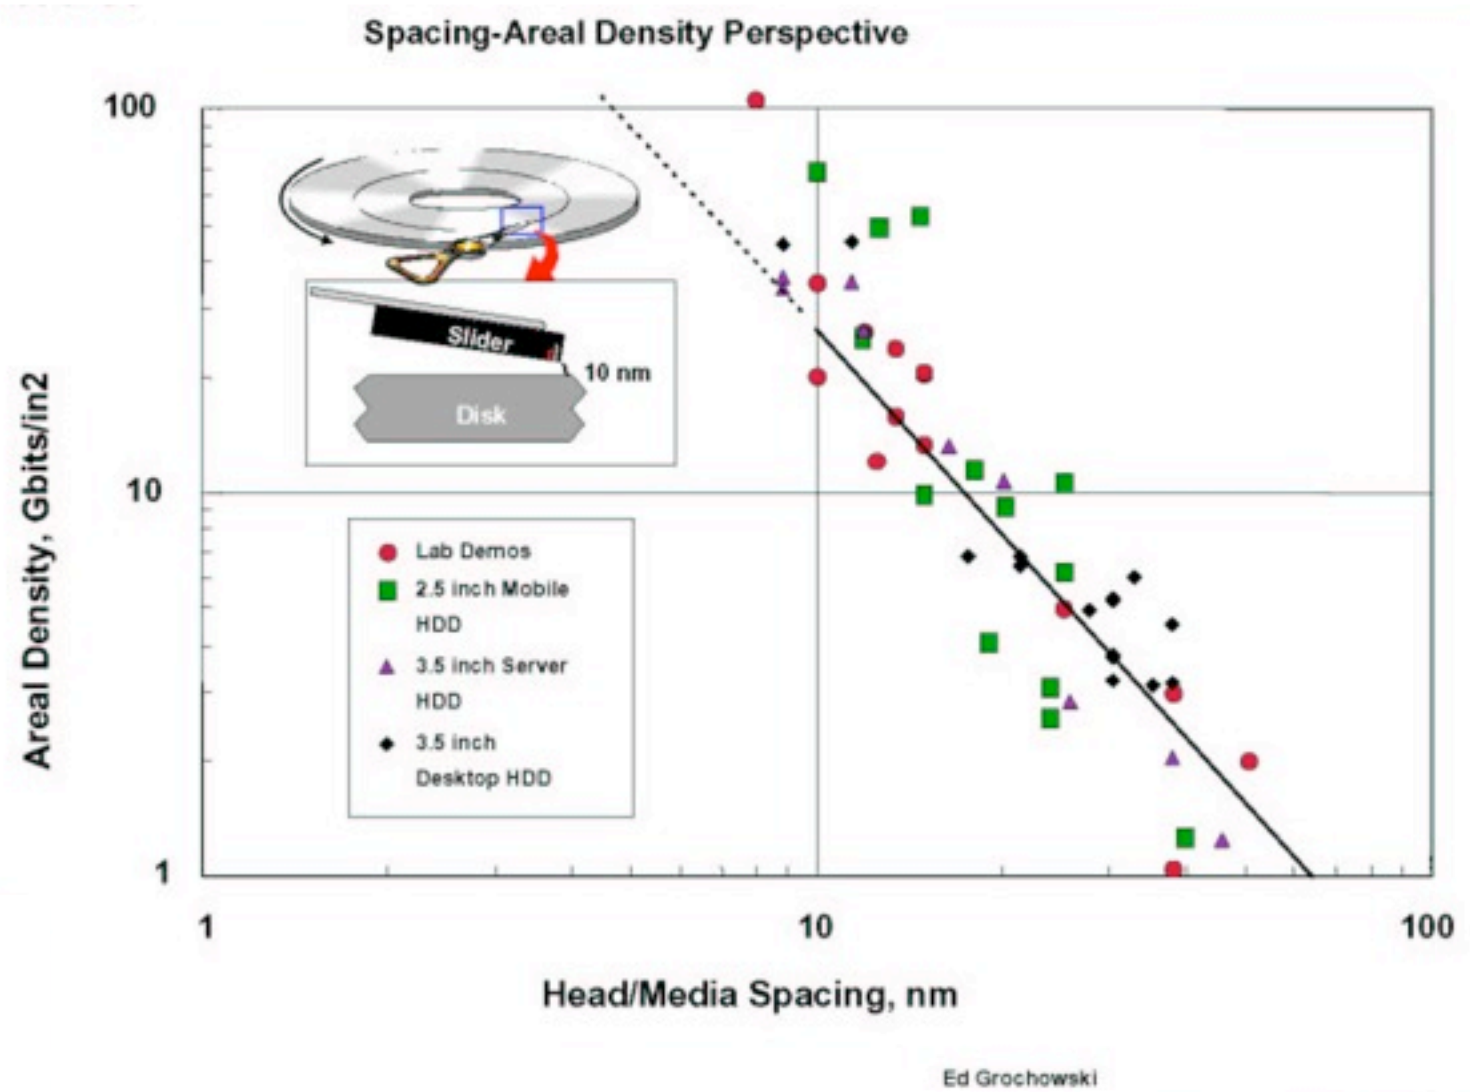

head-media spacing

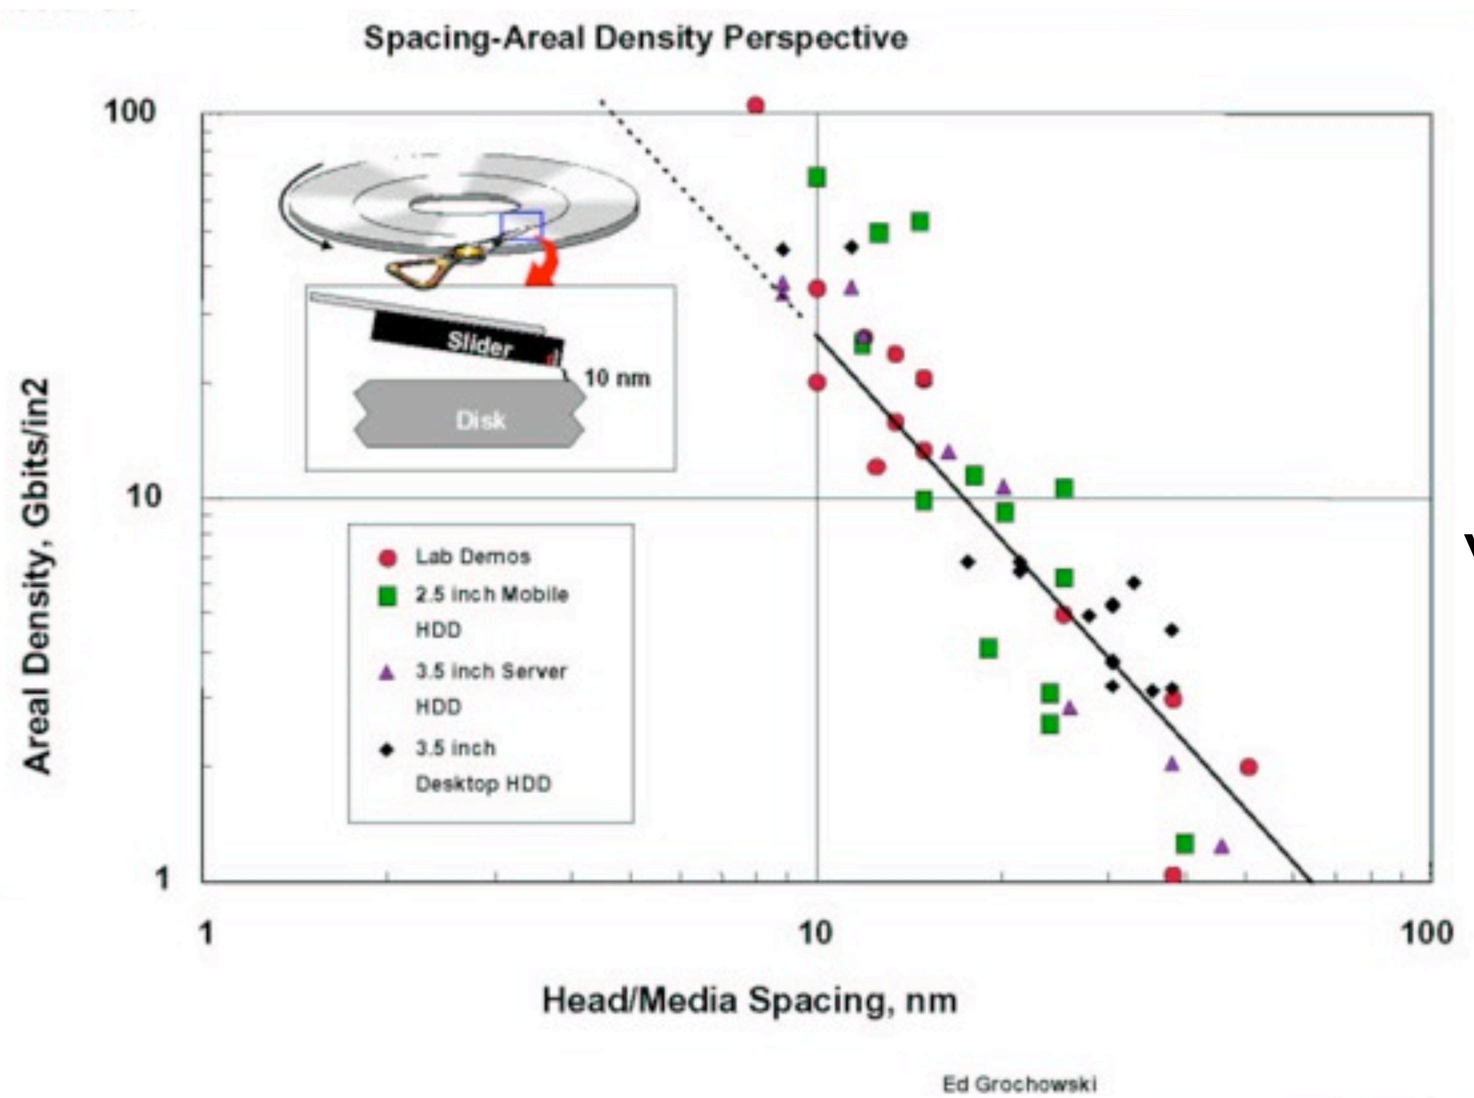

tribology ...

we are basically in contact

physical spacing

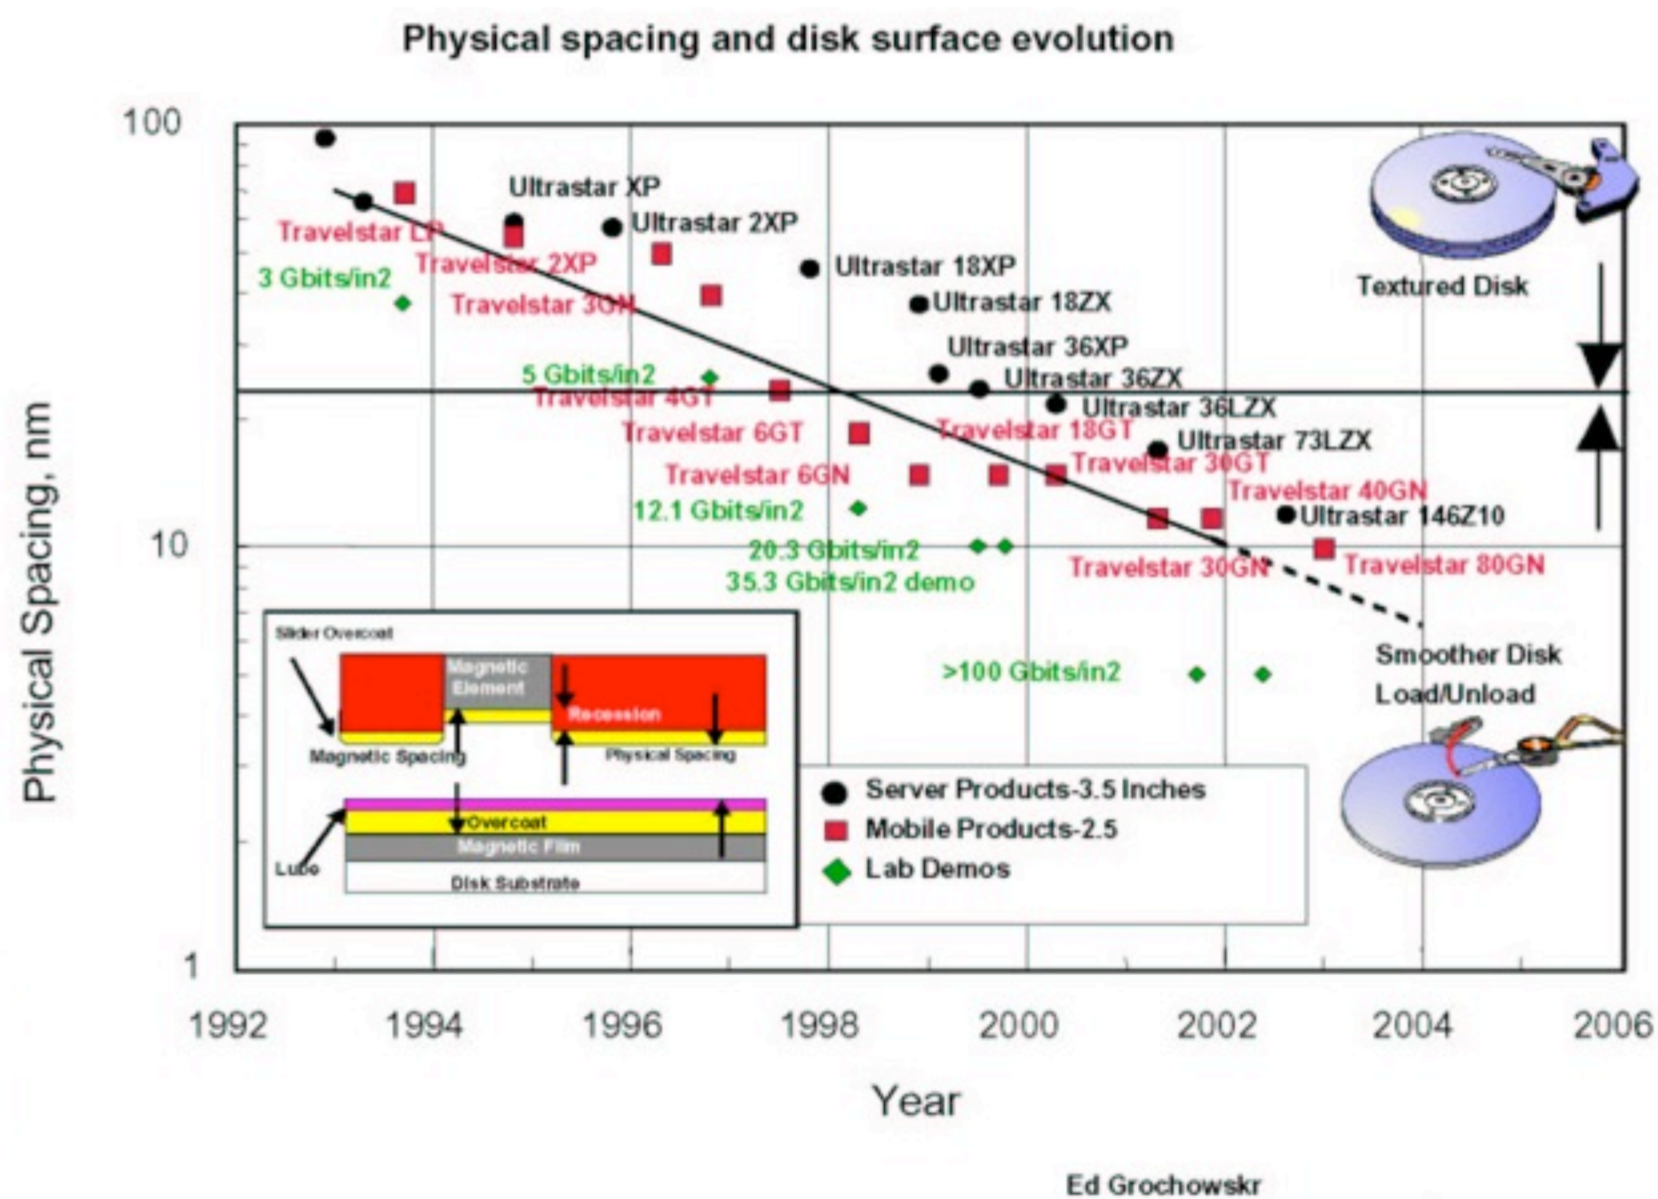

competition?

	SRAM	DRAM	Flash	FeRAM	MRAM	PCRAM
Read speed	Fastest	Medium	Fast	Fast	Fast	Fast
Write speed	Fastest	Medium	Slow	Medium	Fast	Medium
Array efficiency	High	High	Medium/low	Medium	Medium/high	Medium/high
Future scalability	Good	Limited	Limited	Limited	Good	Good
Cell density	Low	High	High	Medium	Medium/high	Medium/high
Nonvolatility	No	No	Yes	Yes	Yes	Yes
Endurance	Infinite	Infinite	Limited	Limited	Infinite	Infinite
Cell leakage	Increasing	High	Low	Low	Low	Low
Low voltage	Yes	Limited	Limited	Limited	Yes	Yes
Complexity	Low	Medium	Medium	Medium	Medium	Medium

competition?

	SRAM	DRAM	Flash	FeRAM	MRAM	PCRAM
Read speed	Fastest	Medium	Fast	Fast	Fast	Fast
Write speed	Fastest	Medium	Slow	Medium	Fast	Medium
Array efficiency	High	High	Medium/low	Medium	Medium/high	Medium/high
Future scalability	Good	Limited	Limited	Limited	Good	Good
Cell density	Low	High	High	Medium	Medium/high	Medium/high
Nonvolatility	No	No	Yes	Yes	Yes	Yes
Endurance	Infinite	Infinite	Limited	Limited	Infinite	Infinite
Cell leakage	Increasing	High	Low	Low	Low	Low
Low voltage	Yes	Limited	Limited	Limited	Yes	Yes
Complexity	Low	Medium	Medium	Medium	Medium	Medium

price is the real advantage.

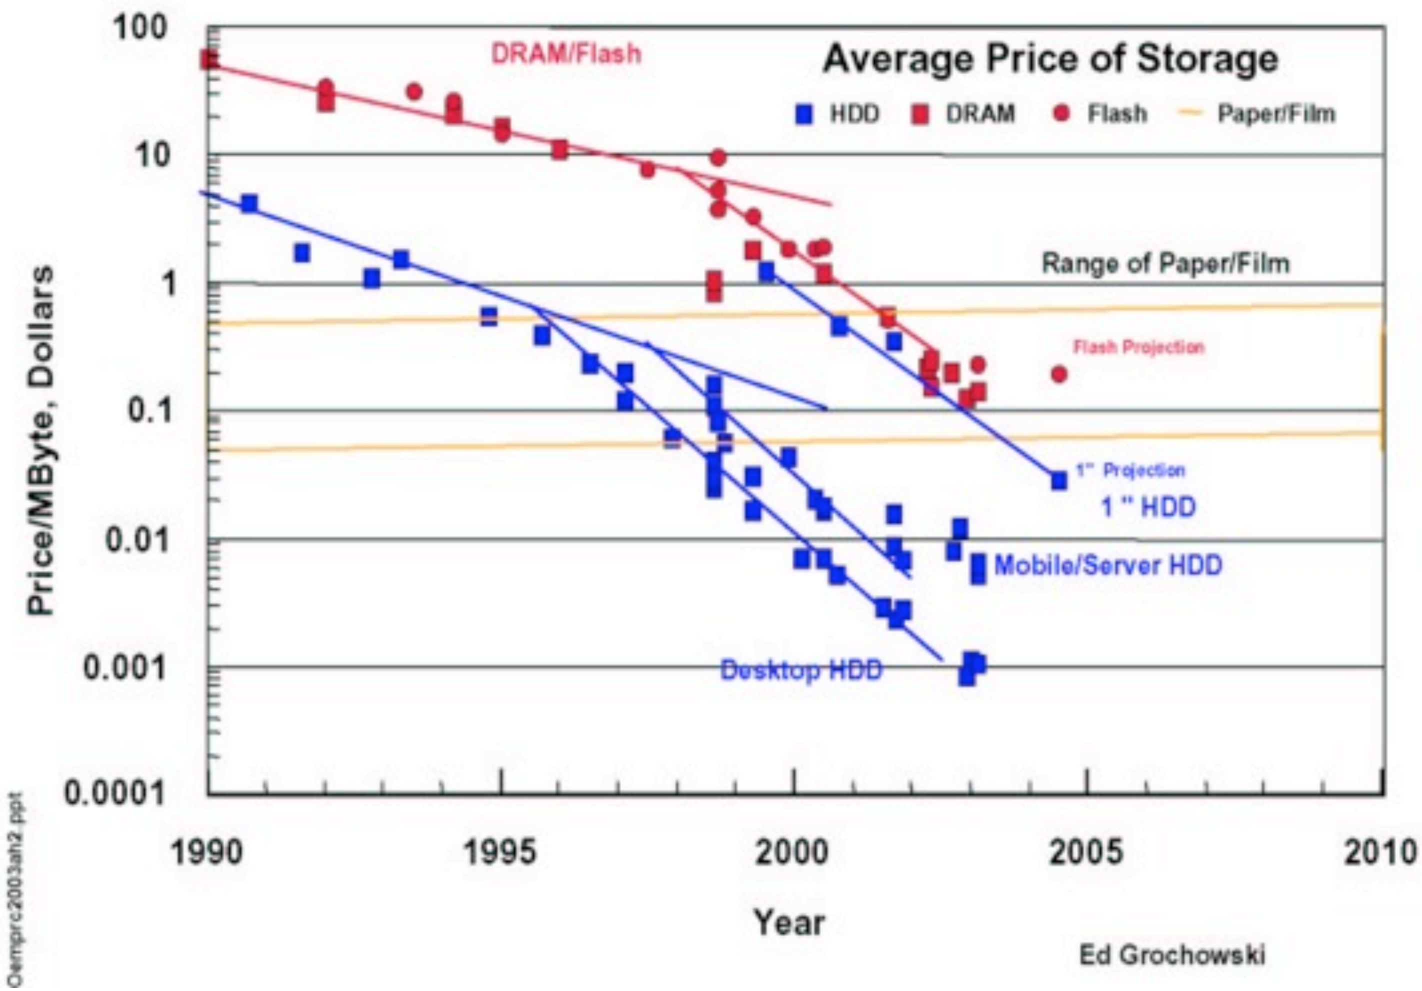

price is the real advantage.

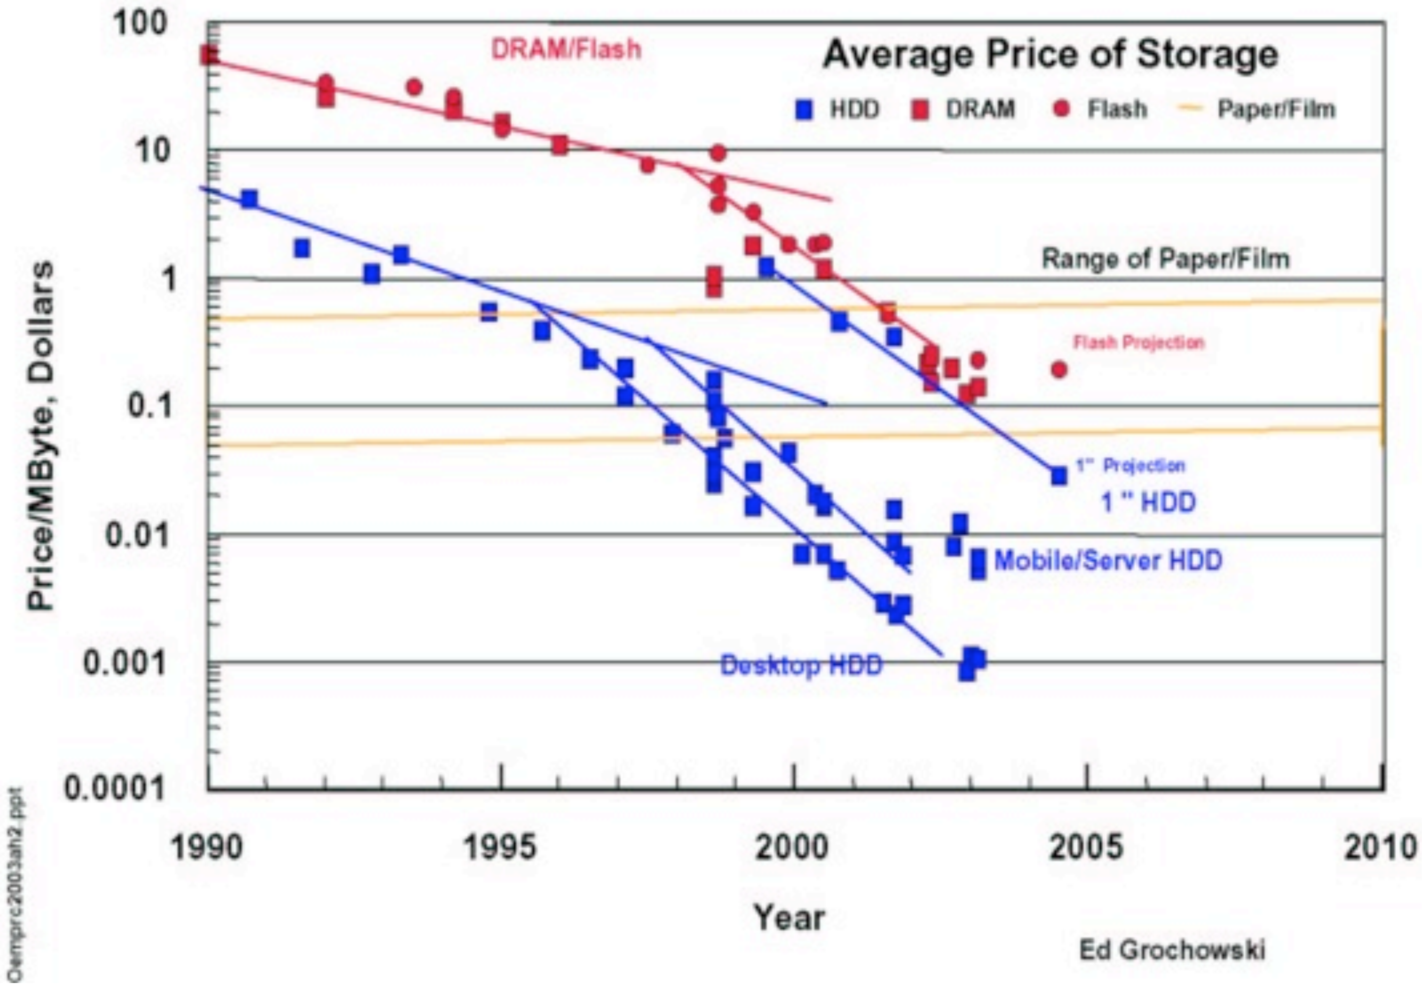

flash is beating the 1" HDD in some apps

e.g., mp3, camera

price is the real advantage.

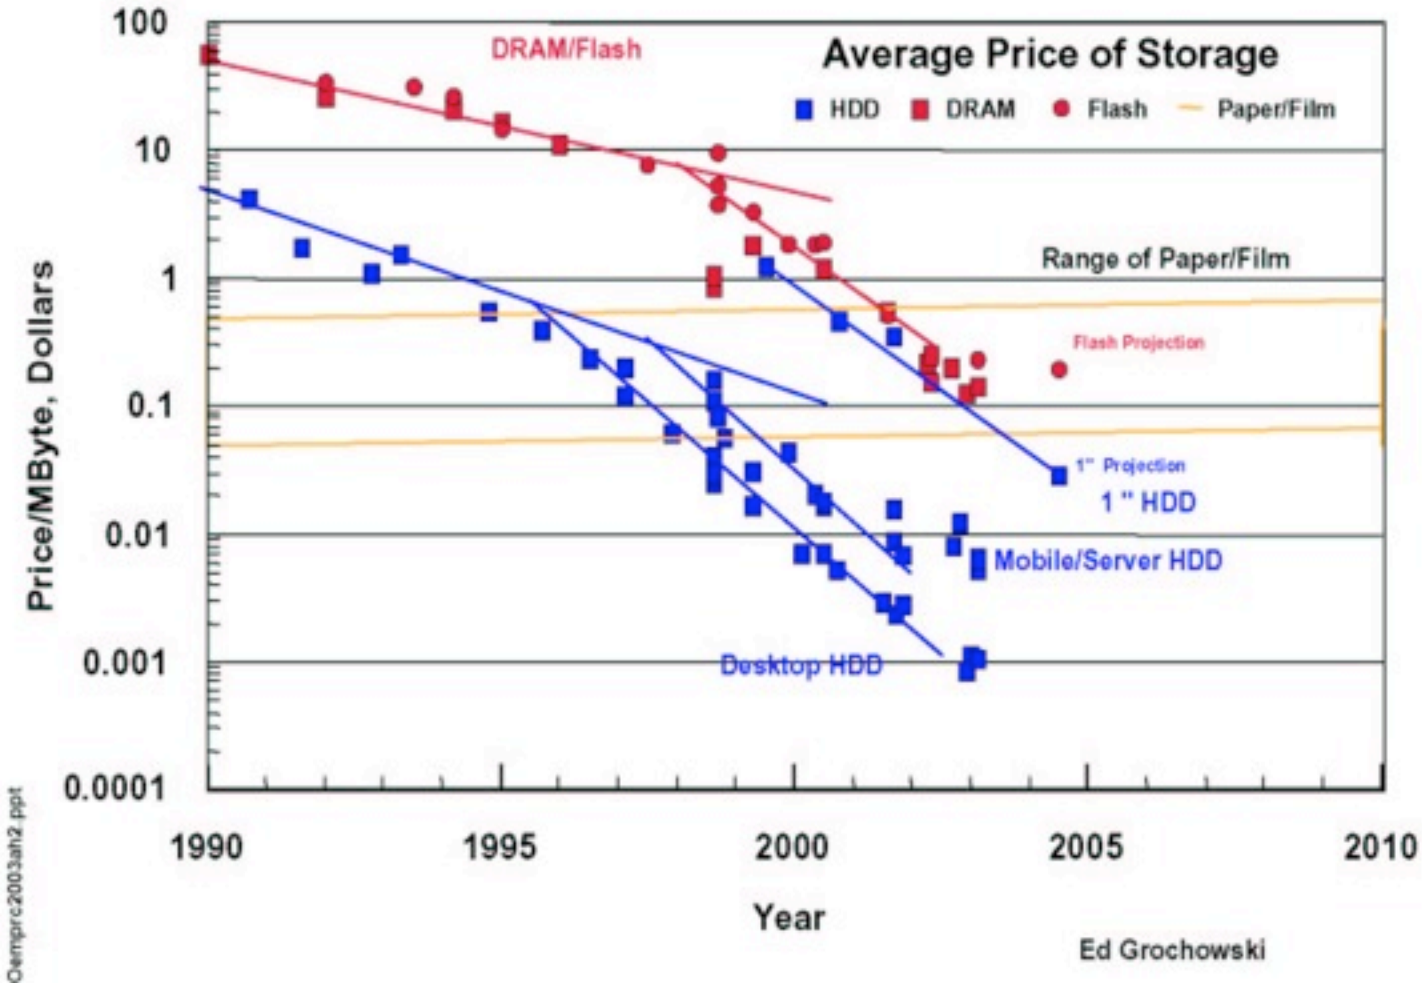

flash is beating the 1" HDD in some apps

e.g., mp3, camera

power consumption may be the larger issue

power consumption is not an advantage

latency ...

fundamental limits of magnetism?

power consumption is not an advantage

latency ...

fundamental limits of magnetism?

more importantly ...

power consumption is not an advantage

latency ...

fundamental limits of magnetism?

more importantly ...
will we get "scooped" like bubbles?

SO!

what about the competition?

how does flash work?

the basics of Flash

NOR Cell

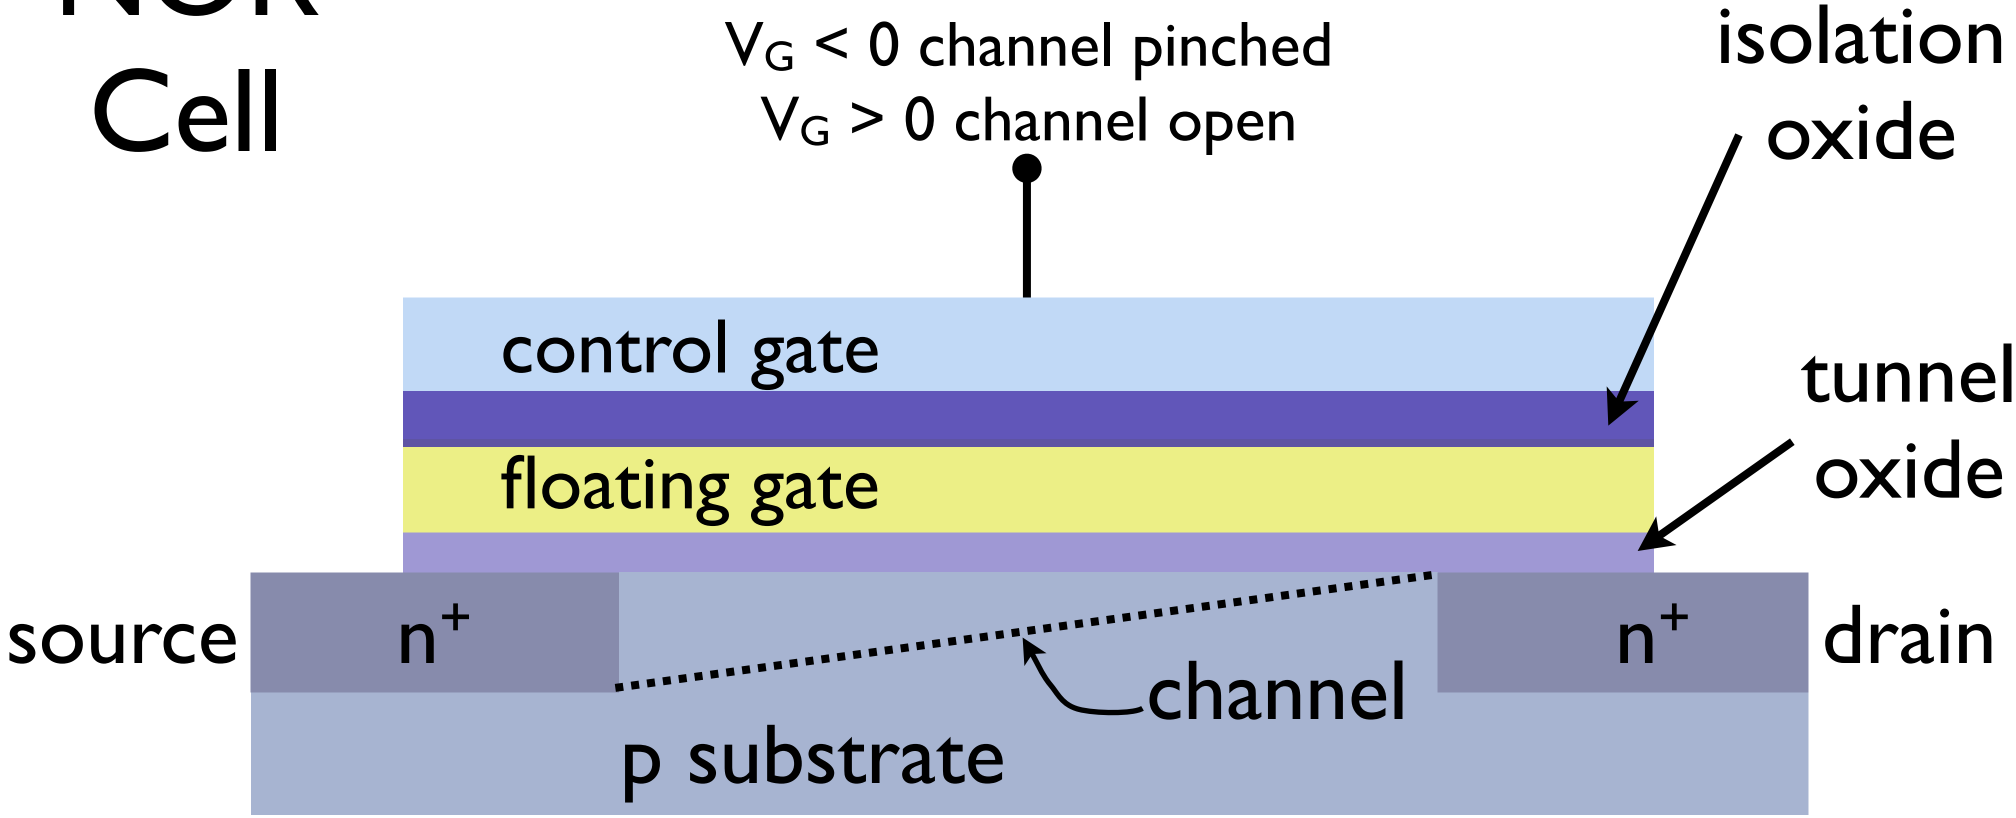

- like a MOSFET
- uses 2 gates

writing

"hot electron injection"

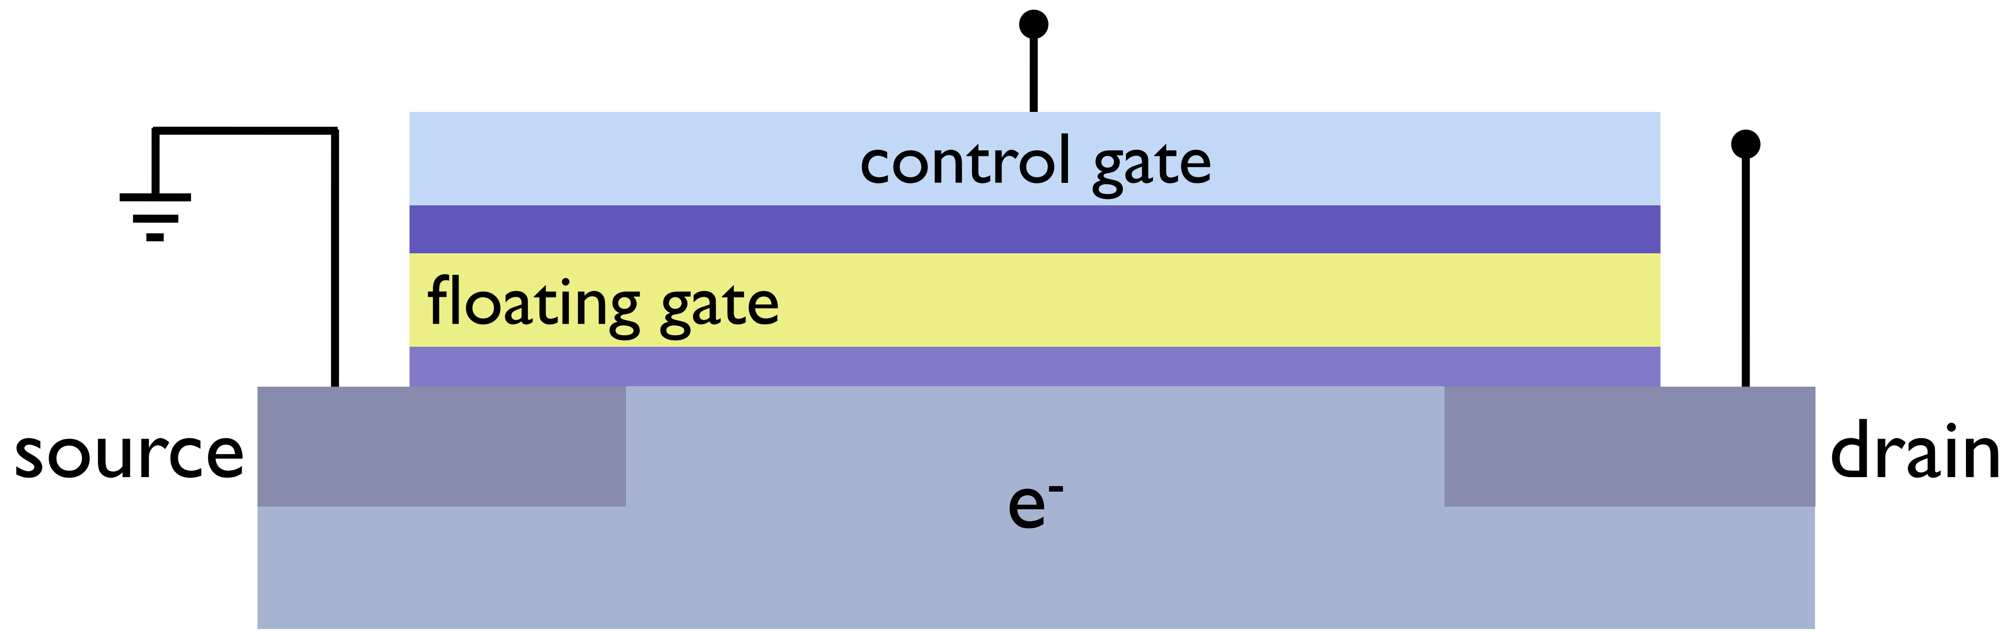

- ~7V to drain
  pull e- through channel
- ~12V to control gate / open channel
  injects e- into floating gate through tunnel oxide
- floating gate now charged

writing

"hot electron injection"

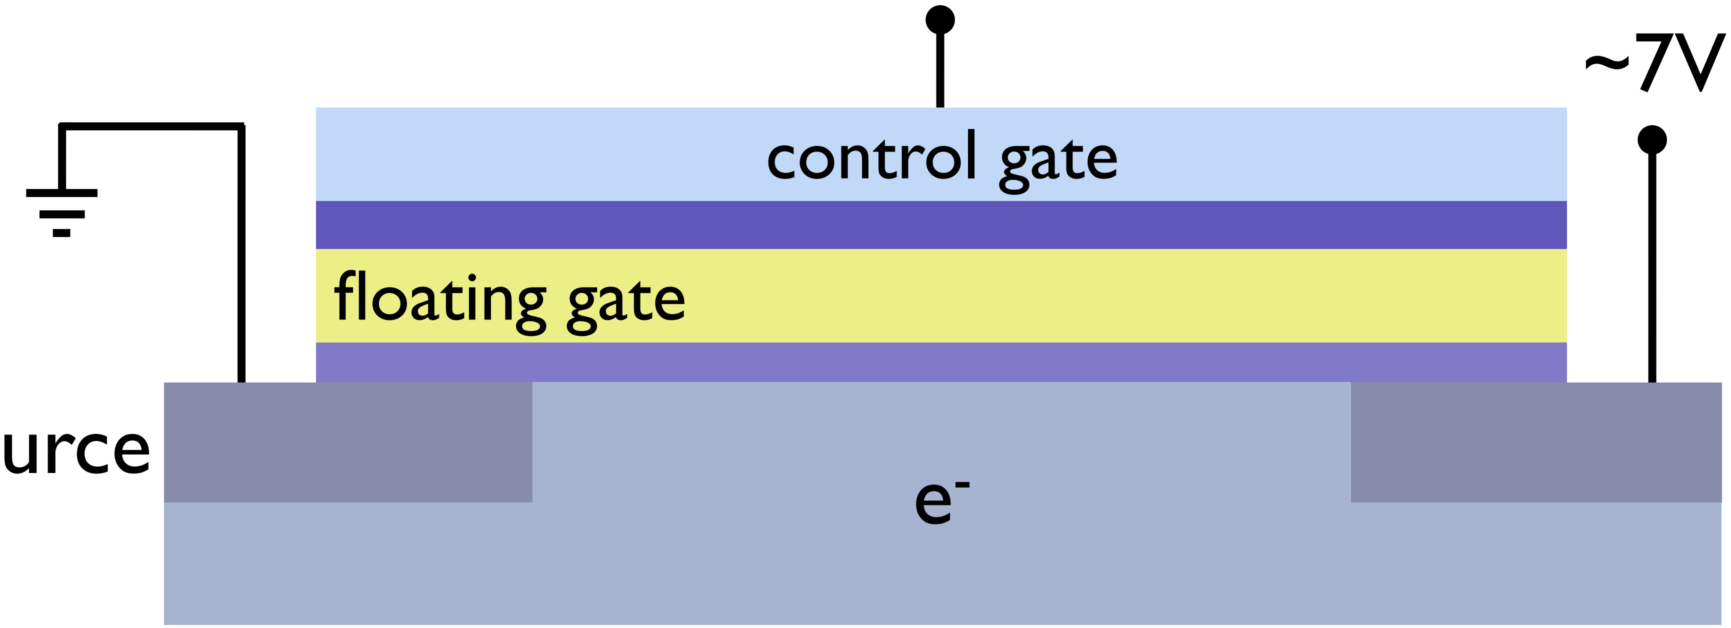

- ~7V to drain
  pull e⁻ through channel
- ~12V to control gate / open channel
  injects e⁻ into floating gate through tunnel oxide
- floating gate now charged

writing

"hot electron injection"

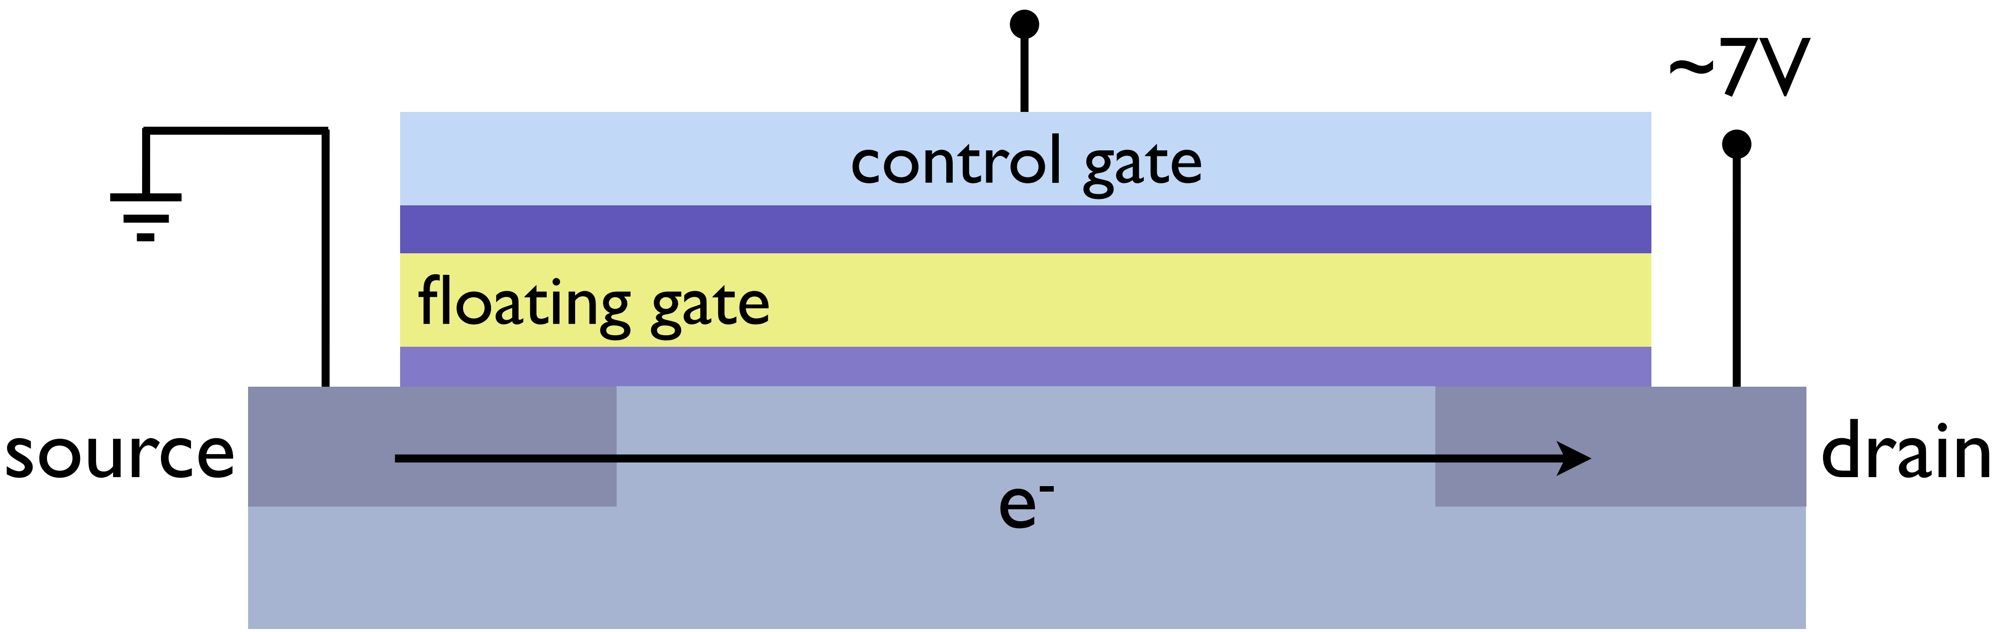

- ~7V to drain
  pull e⁻ through channel
- ~12V to control gate / open channel
  injects e⁻ into floating gate through tunnel oxide
- floating gate now charged

writing

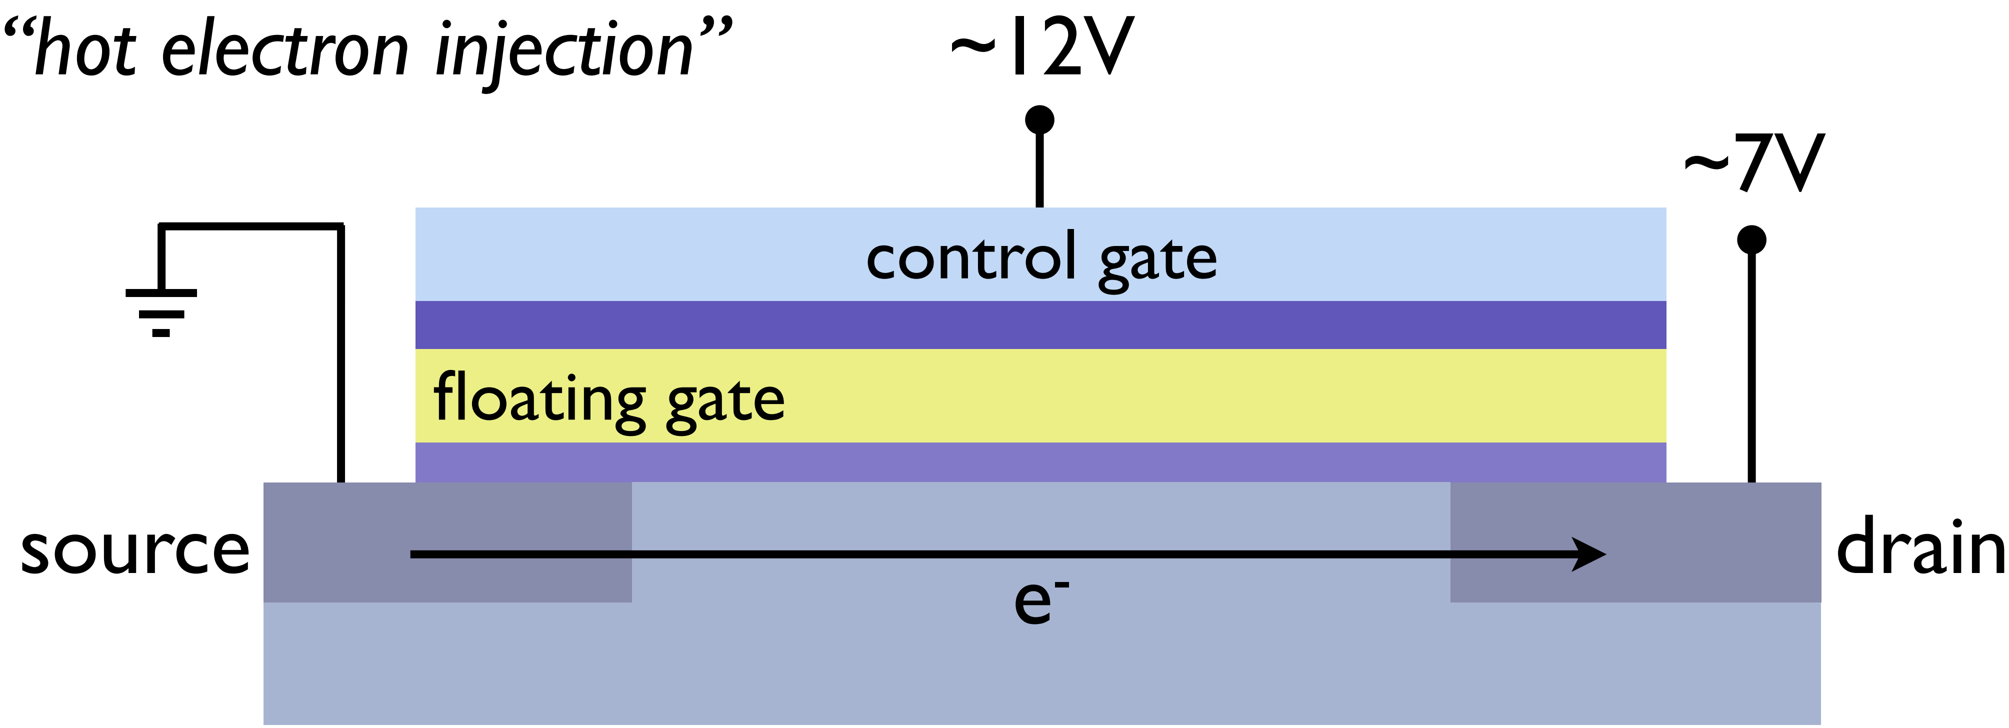

- ~7V to drain
  pull e^- through channel
- ~12V to control gate / open channel
  injects e^- into floating gate through tunnel oxide
- floating gate now charged

writing

"hot electron injection"

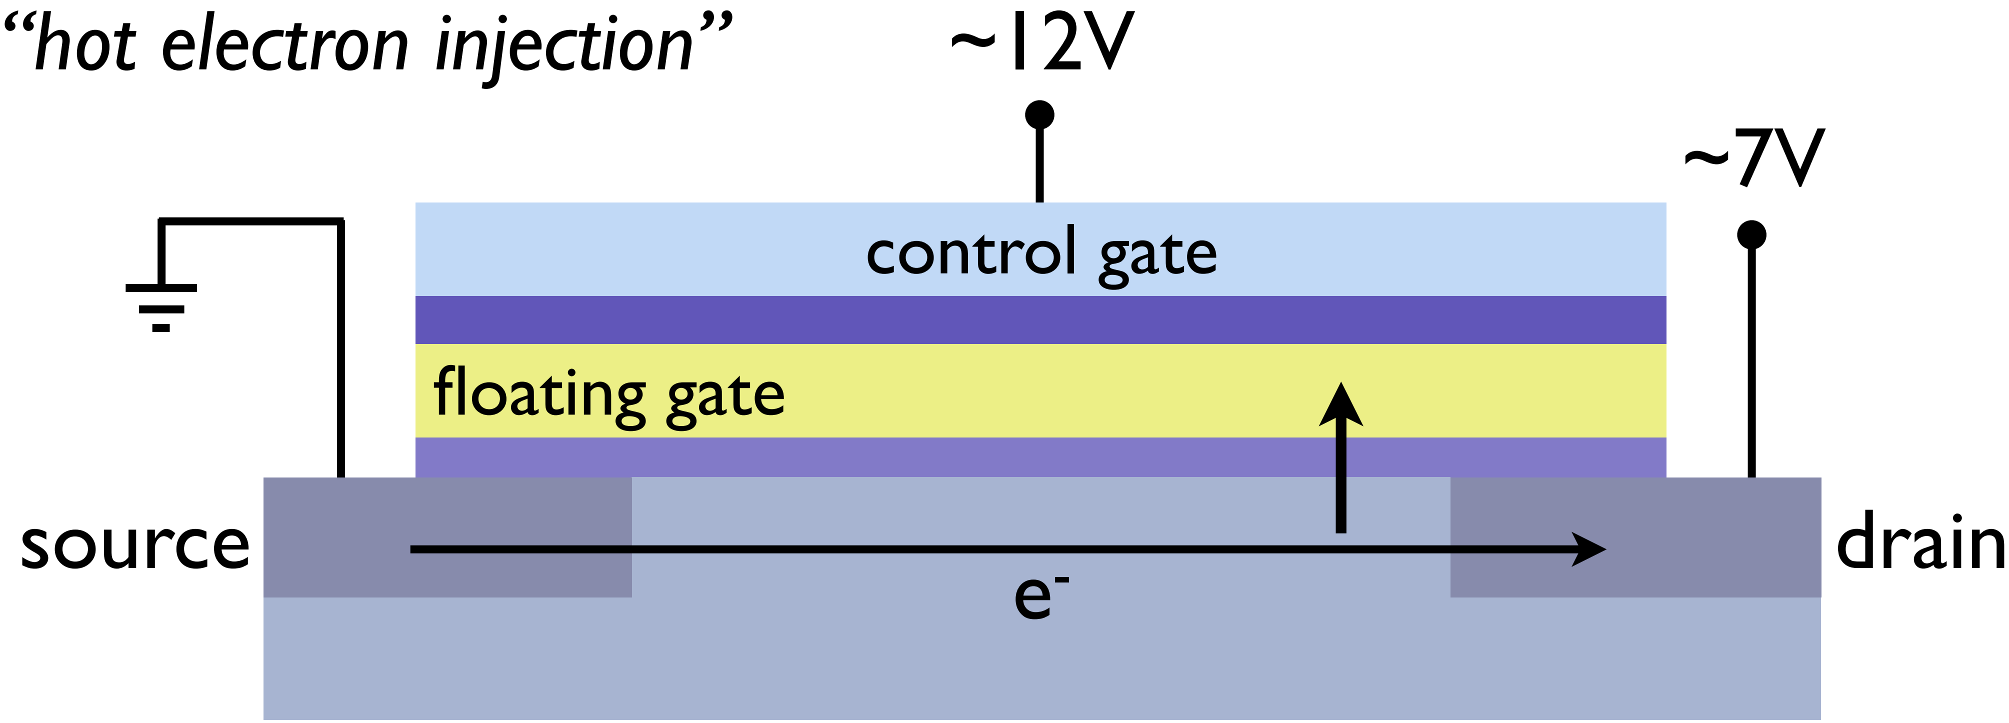

- ~7V to drain
  pull e⁻ through channel
- ~12V to control gate / open channel
  injects e⁻ into floating gate through tunnel oxide
- floating gate now charged

writing

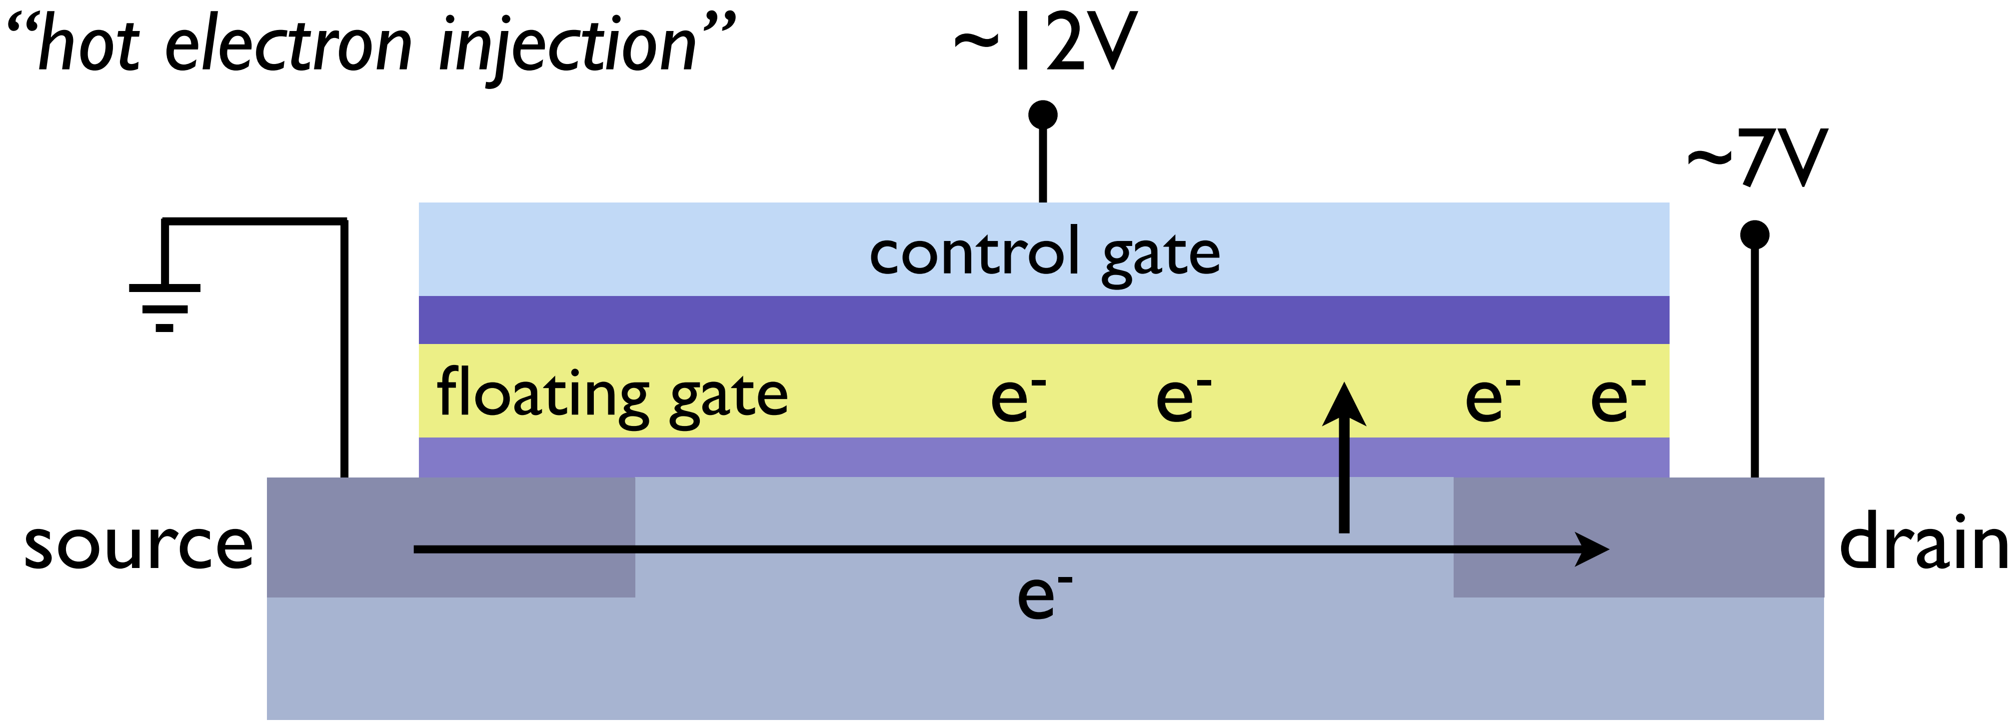

- ~7V to drain
  pull e⁻ through channel
- ~12V to control gate / open channel
  injects e⁻ into floating gate through tunnel oxide
- floating gate now charged

erasing

reset all bits to "1"

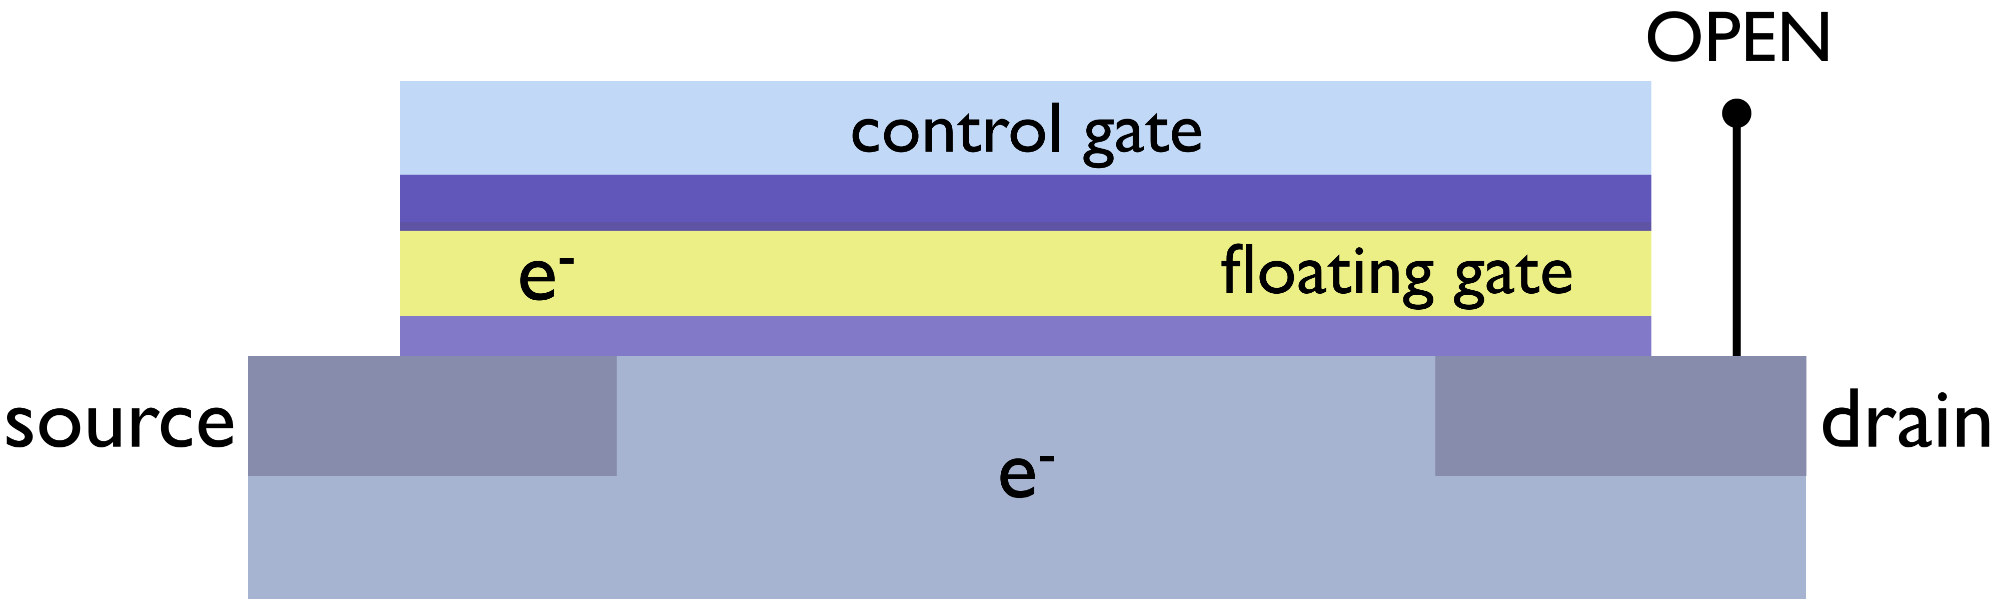

- -9V to control
  - pinch off channel
- ~6V to source
- suck electrons out of floating gate into source
  - Fowler-Nordheim tunneling

erasing

reset all bits to "1"

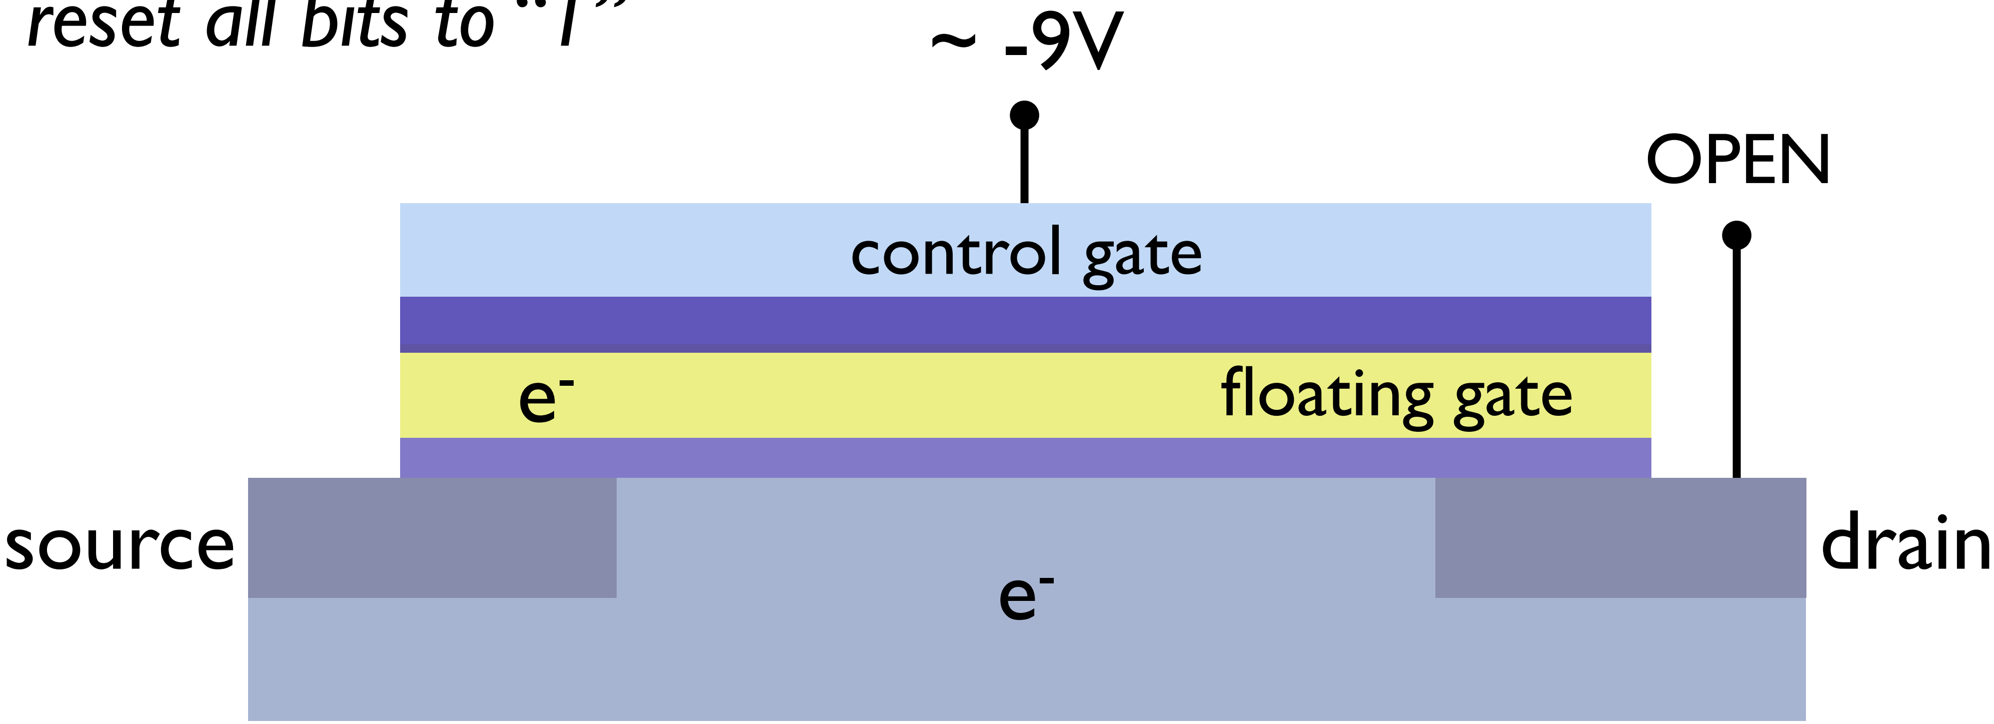

- -9V to control
  pinch off channel
- ~6V to source
- suck electrons out of floating gate into source
  Fowler-Nordheim tunneling

erasing

reset all bits to "1"

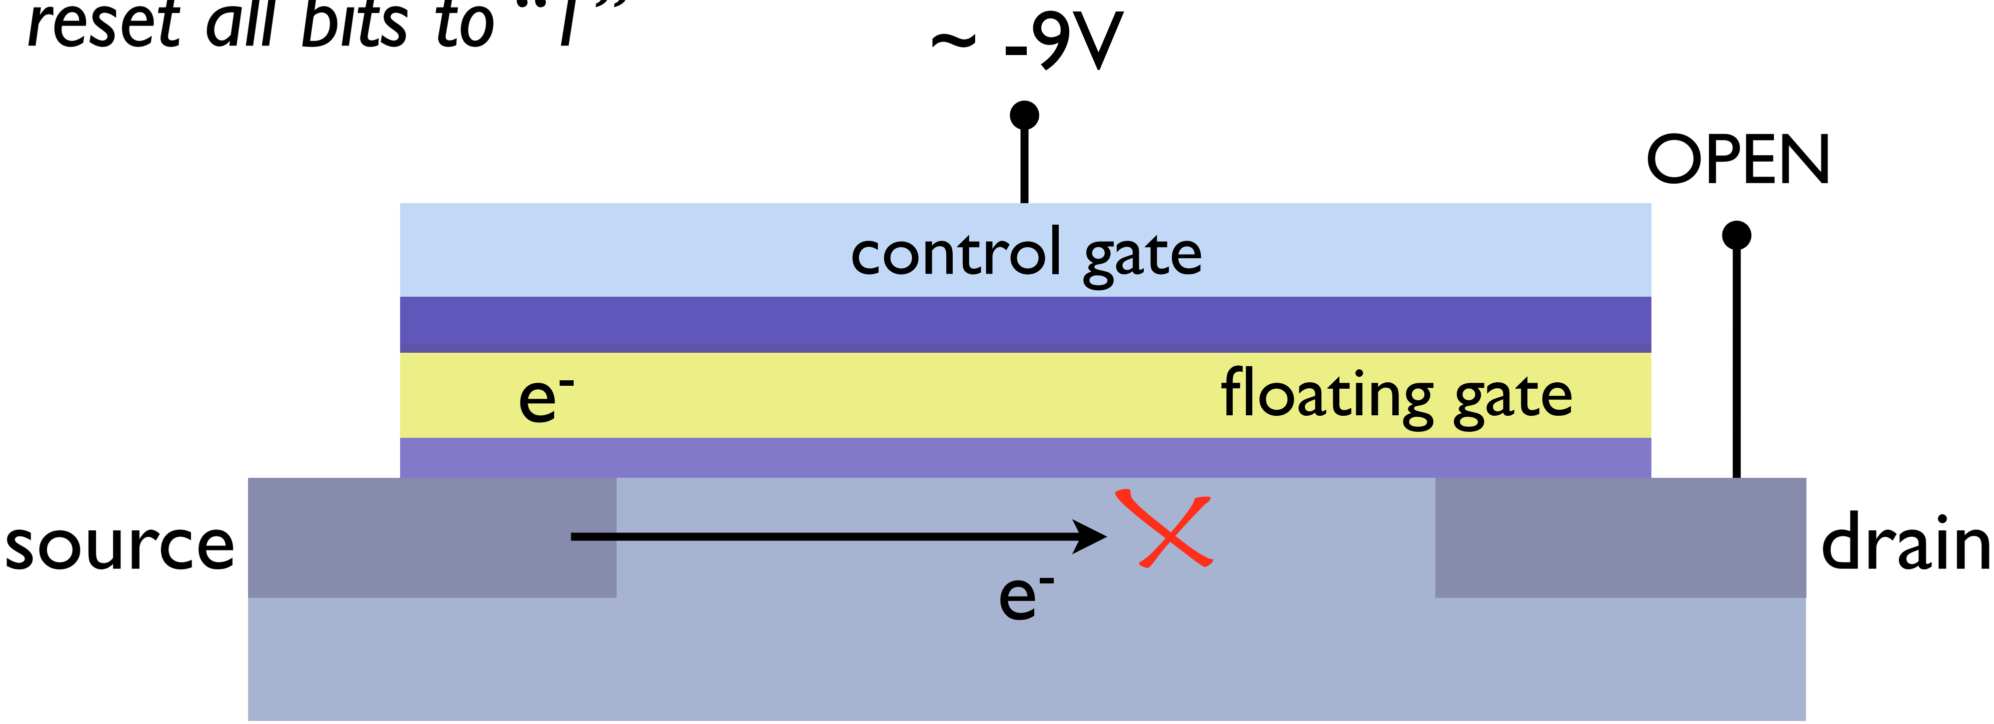

- -9V to control
  pinch off channel
- ~6V to source
- suck electrons out of floating gate into source
  Fowler-Nordheim tunneling

erasing

reset all bits to "1"

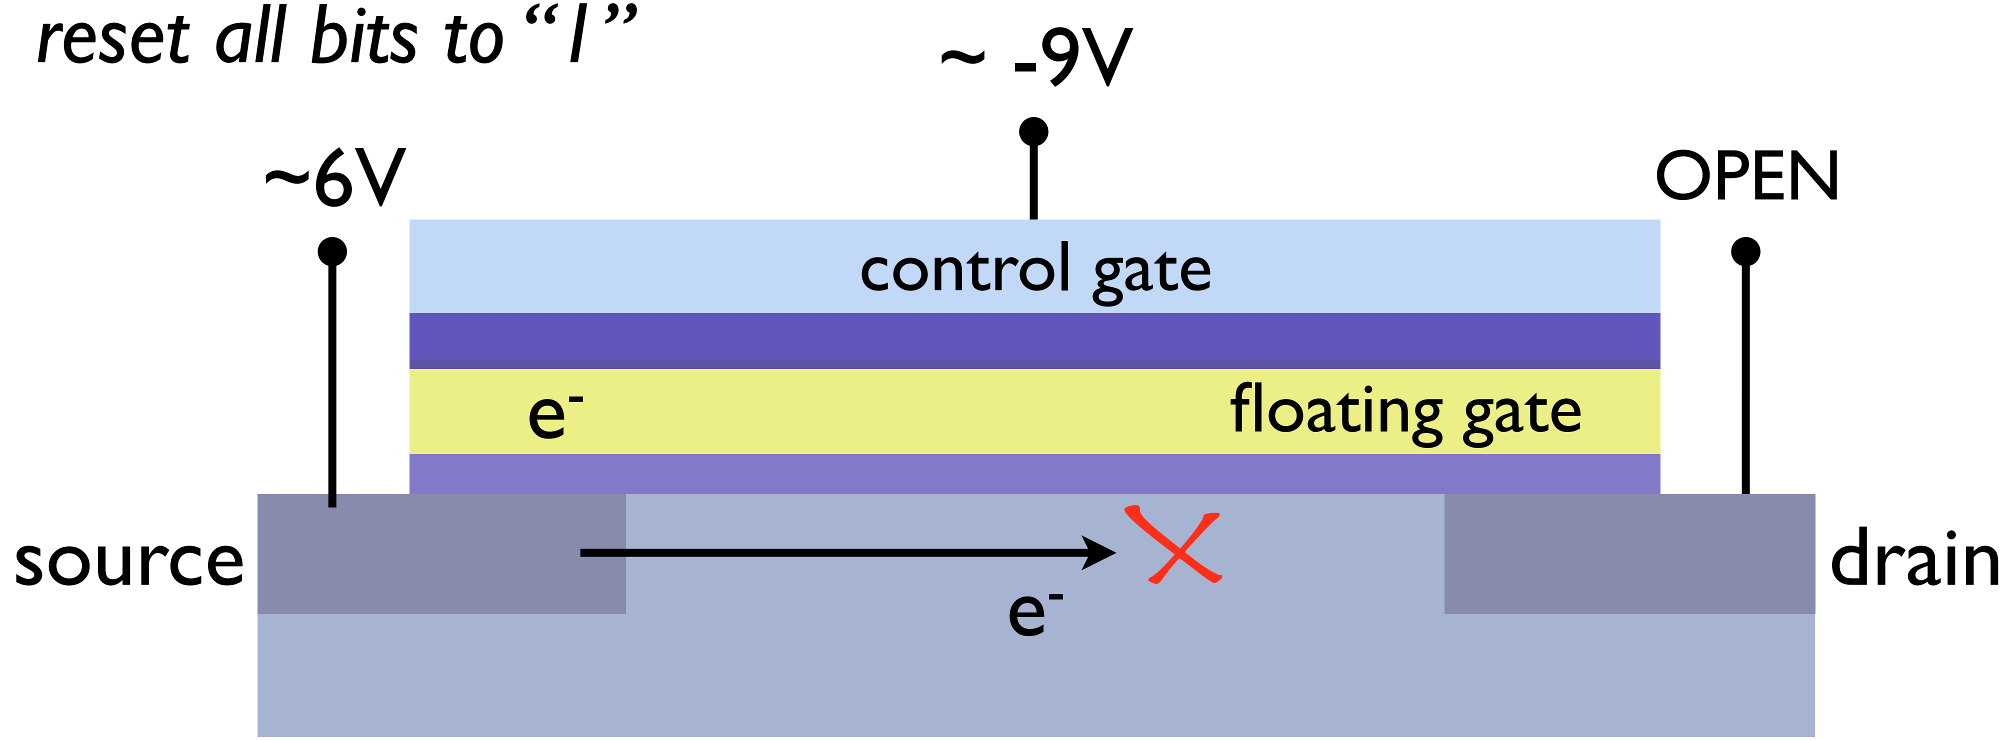

- -9V to control
  pinch off channel
- ~6V to source
- suck electrons out of floating gate into source
  Fowler-Nordheim tunneling

erasing

reset all bits to "1"

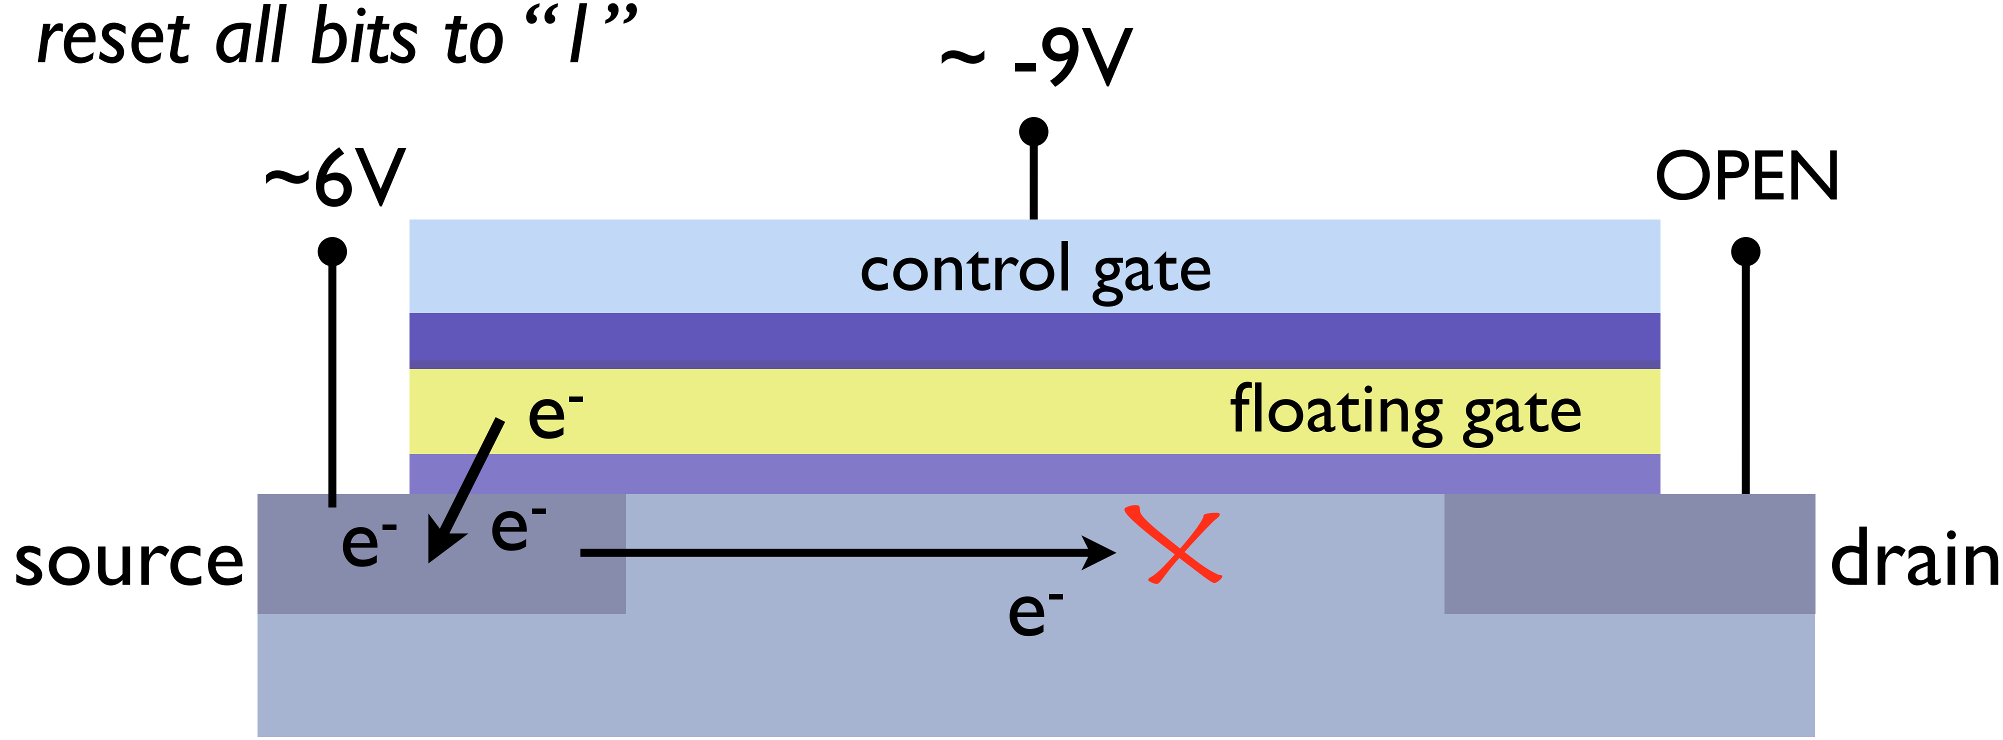

- -9V to control
  pinch off channel
- ~6V to source
- suck electrons out of floating gate into source
  Fowler-Nordheim tunneling

erasing

reset all bits to "1"

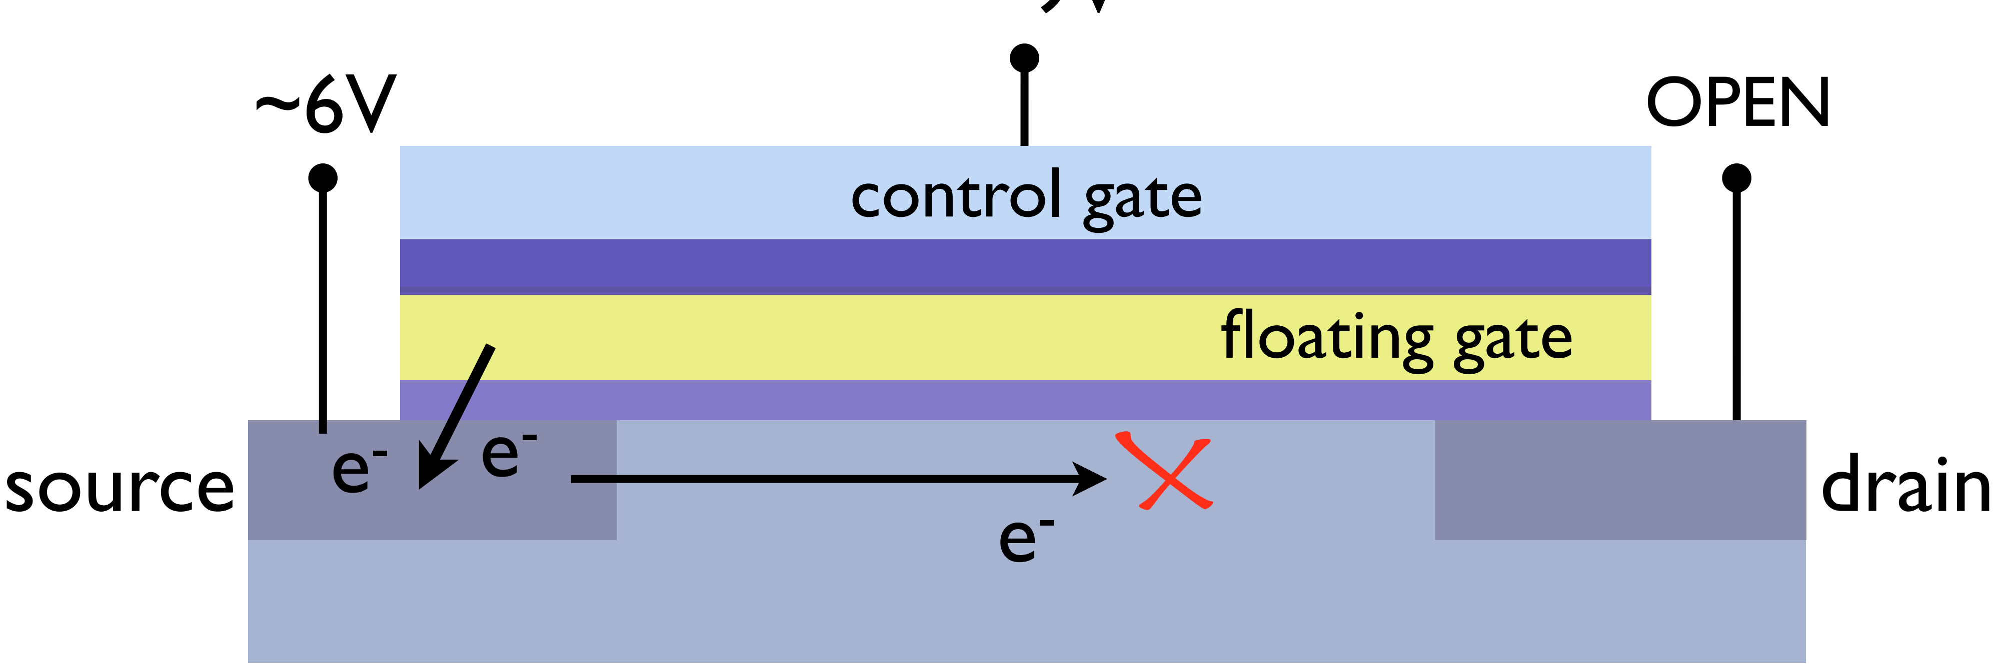

- -9V to control
  - pinch off channel
- ~6V to source
- suck electrons out of floating gate into source
  - Fowler-Nordheim tunneling

reading

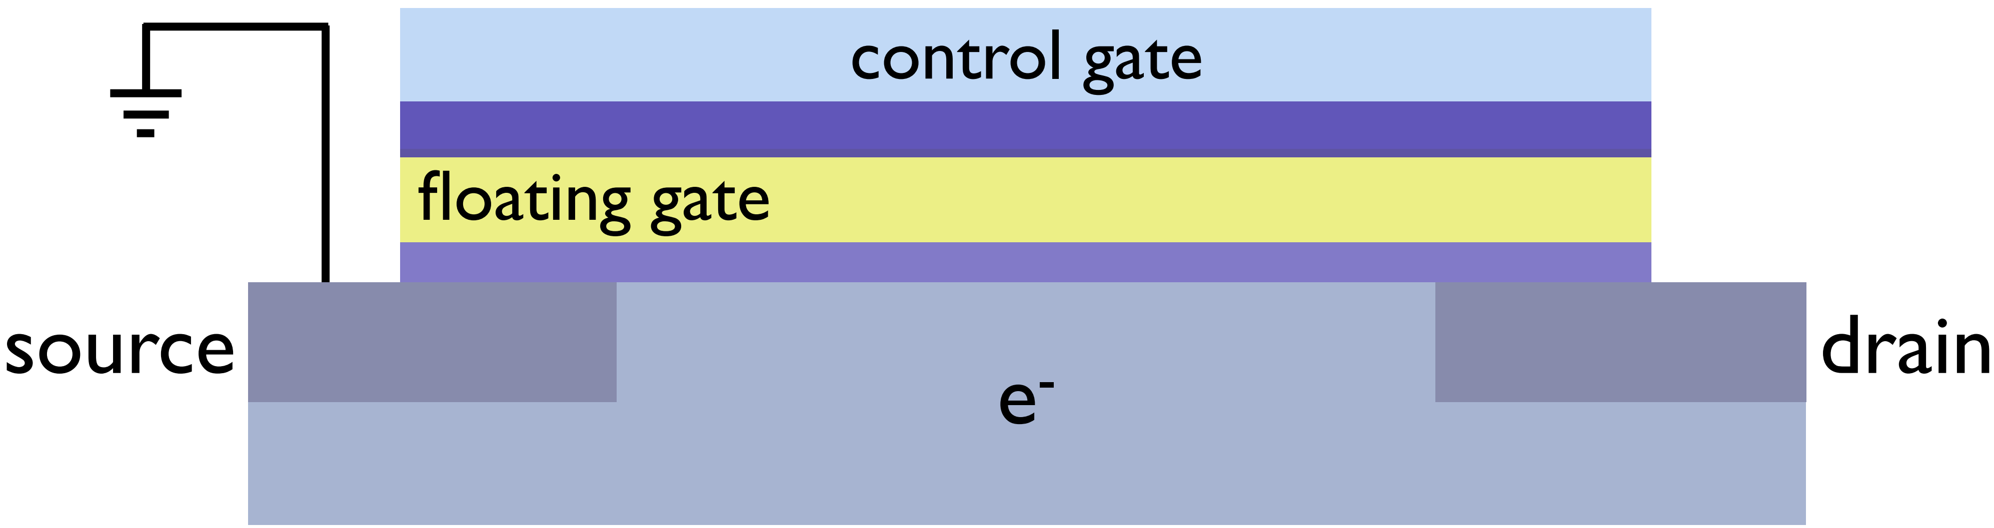

- 5V to control
- 1V to drain
- floating gate charged = channel is pinched off = "0"
- floating gate *discharged* = channel open = "1"

presence of charge modulates I_{SD} !

reading

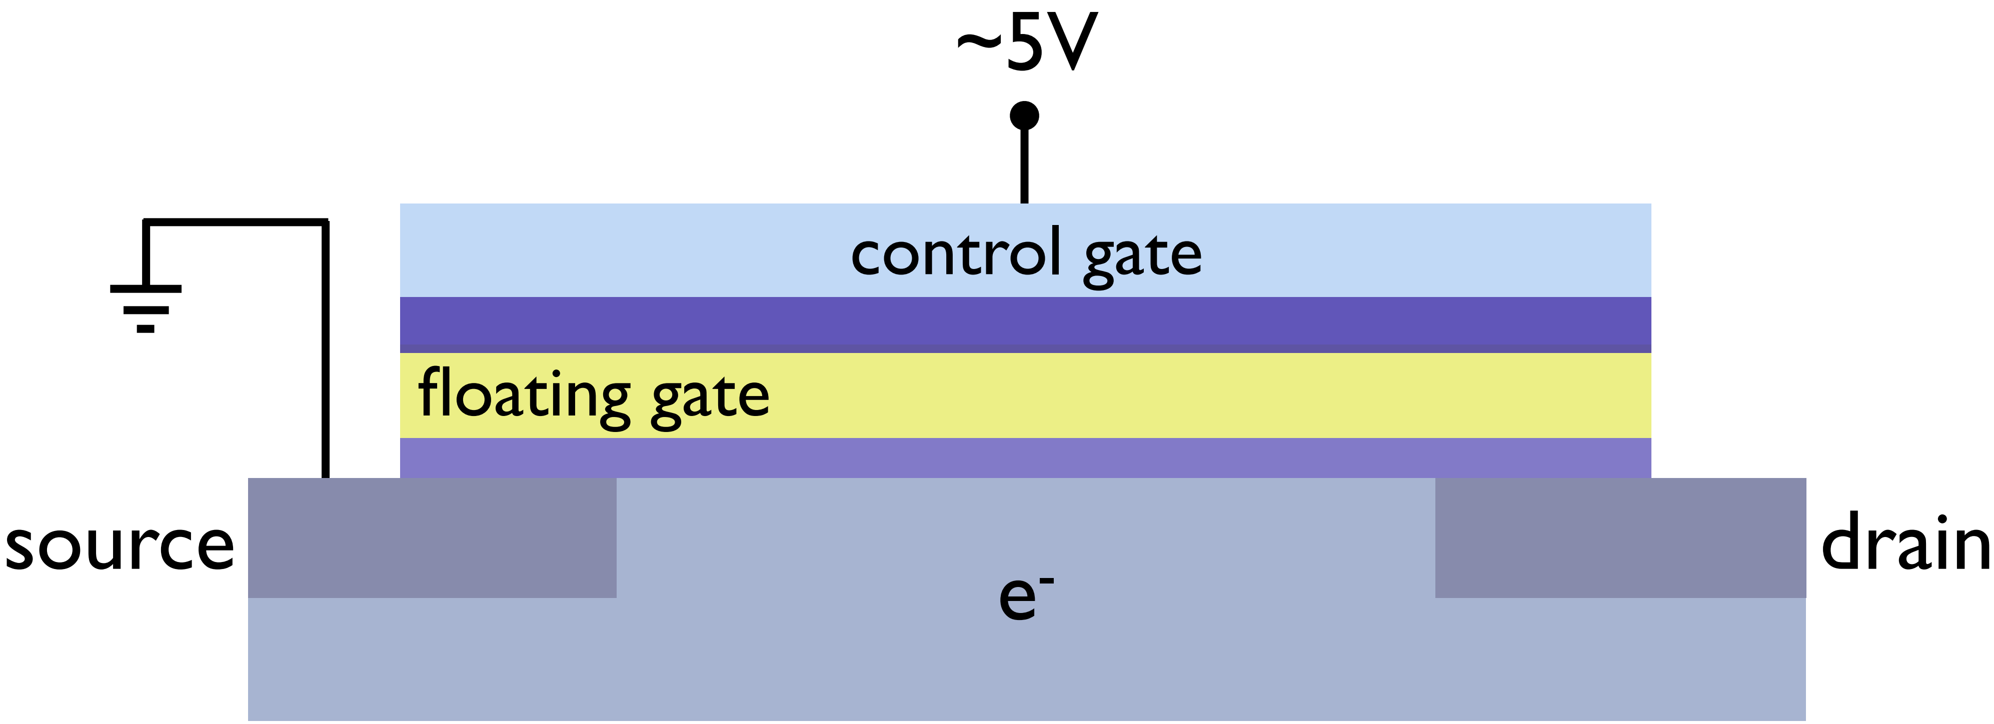

- 5V to control
- 1V to drain
- floating gate charged = channel is pinched off = "0"
- floating gate *discharged* = channel open = "1"

presence of charge modulates I_{SD} !

reading

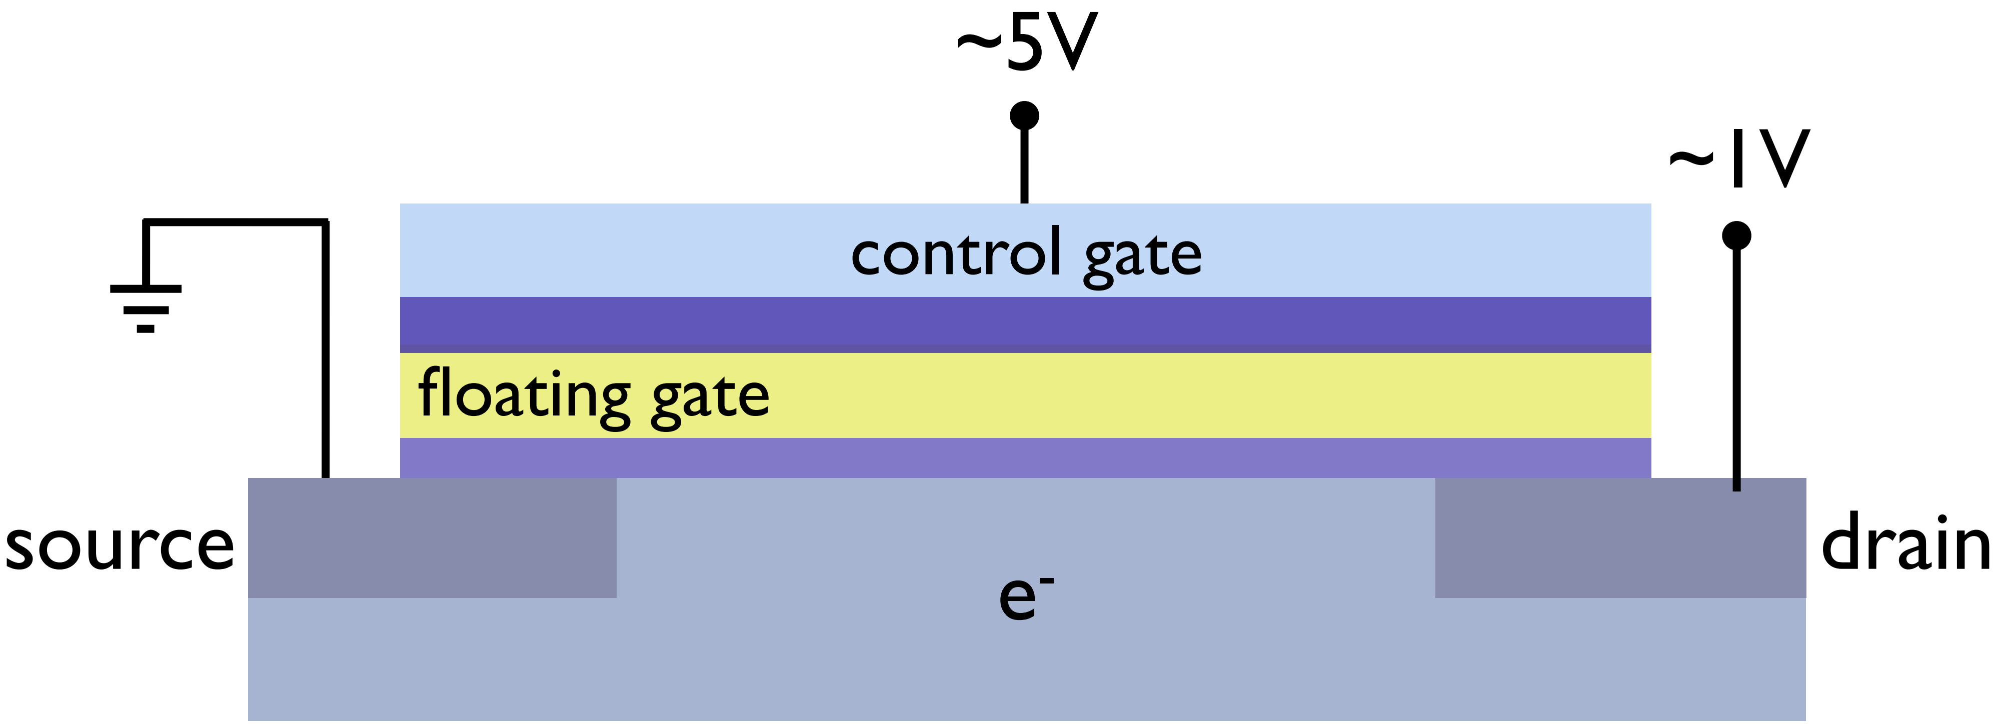

- 5V to control
- 1V to drain
- floating gate charged = channel is pinched off = "0"
- floating gate *discharged* = channel open = "1"

presence of charge modulates I_{SD} !

reading

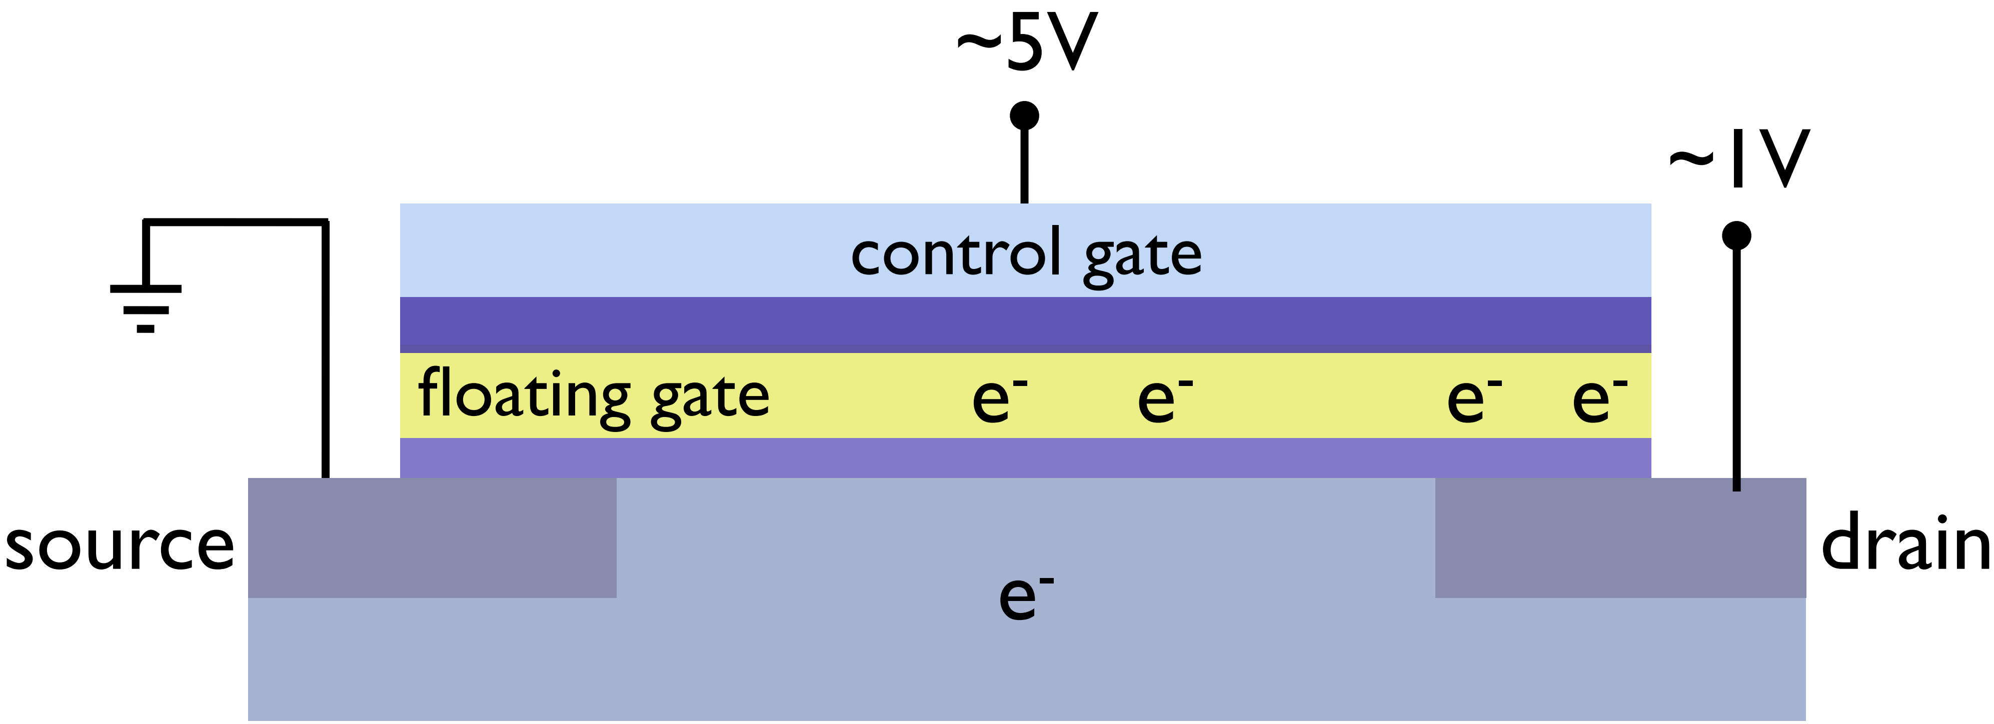

- 5V to control
- 1V to drain
- floating gate charged = channel is pinched off = "0"
- floating gate *discharged* = channel open = "1"

presence of charge modulates I_{SD} !

reading

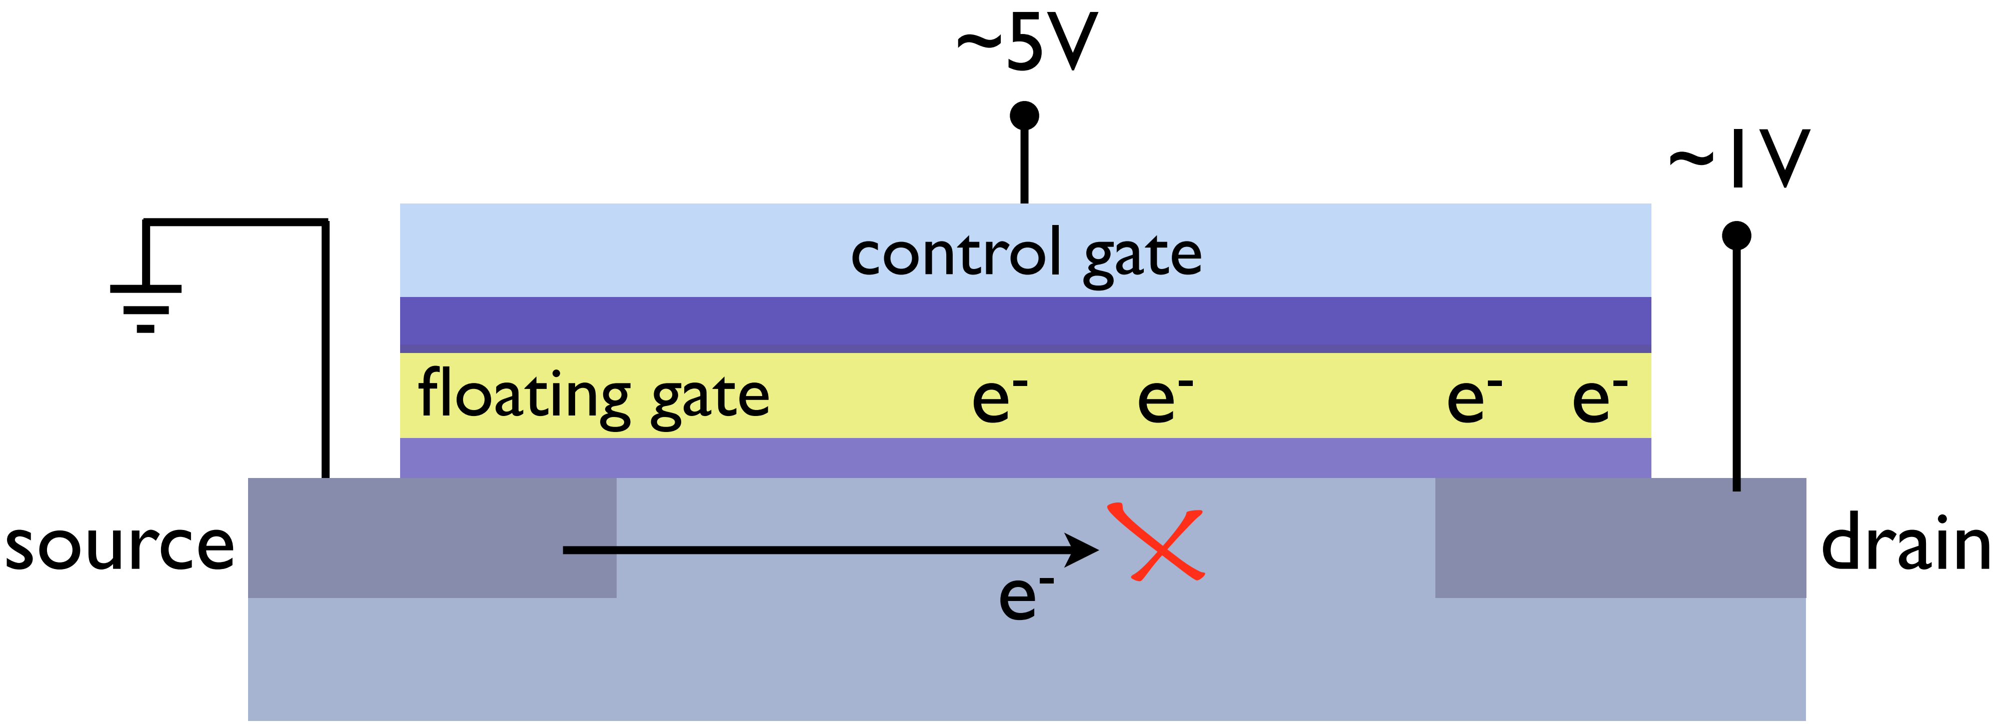

- 5V to control
- 1V to drain
- floating gate charged = channel is pinched off = "0"
- floating gate *discharged* = channel open = "1"

presence of charge modulates I_{SD} !

reading

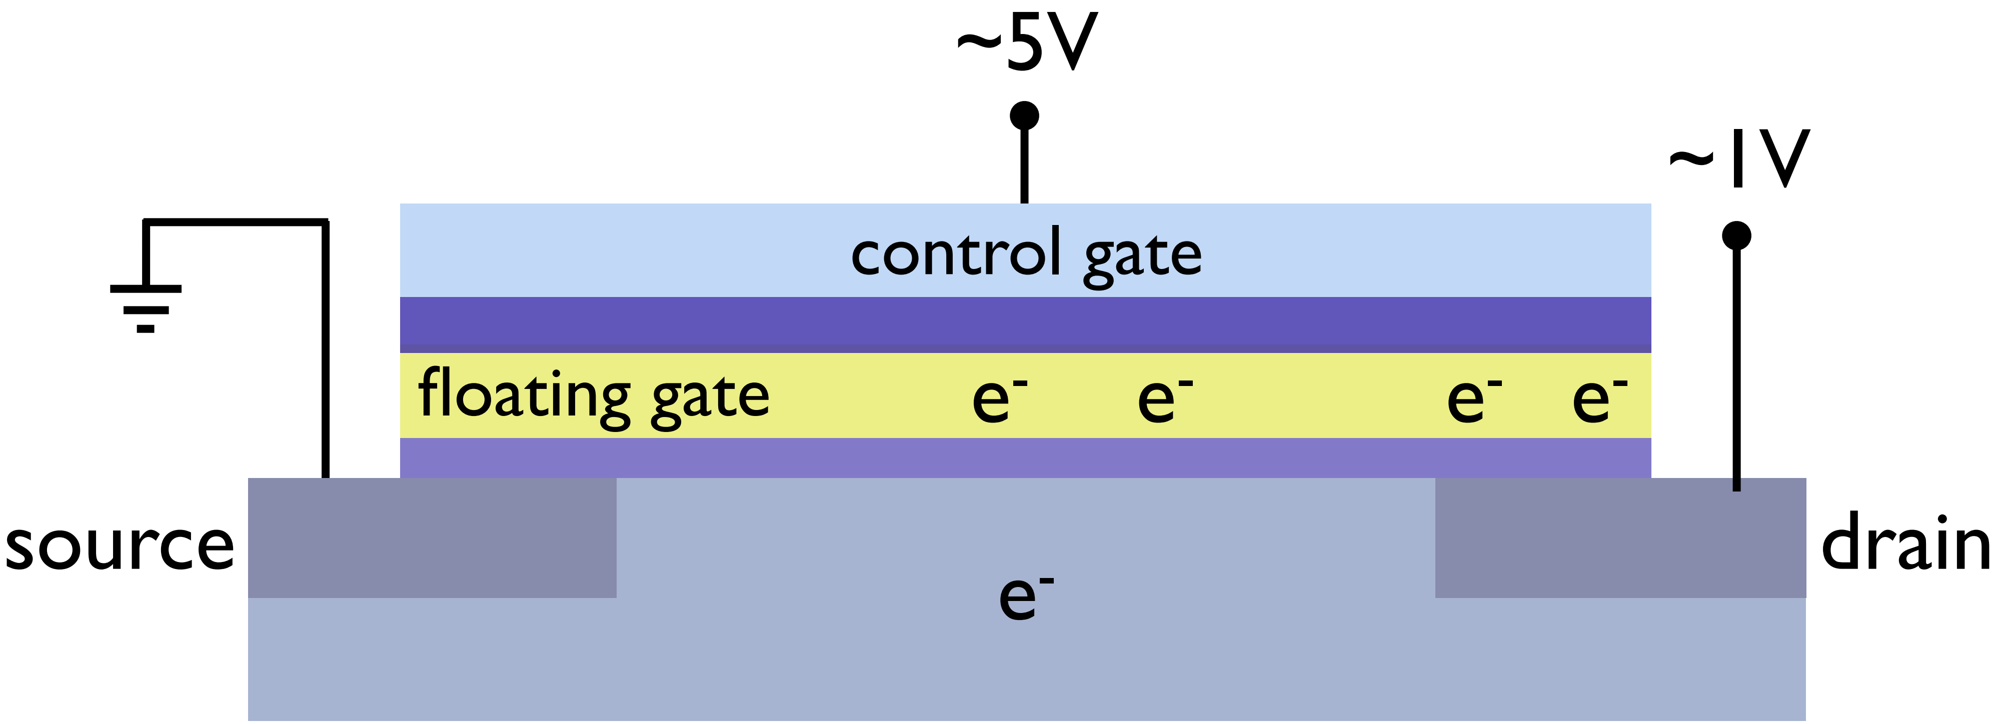

- 5V to control
- 1V to drain
- floating gate charged = channel is pinched off = "0"
- floating gate *discharged* = channel open = "1"

presence of charge modulates I_{SD} !

reading

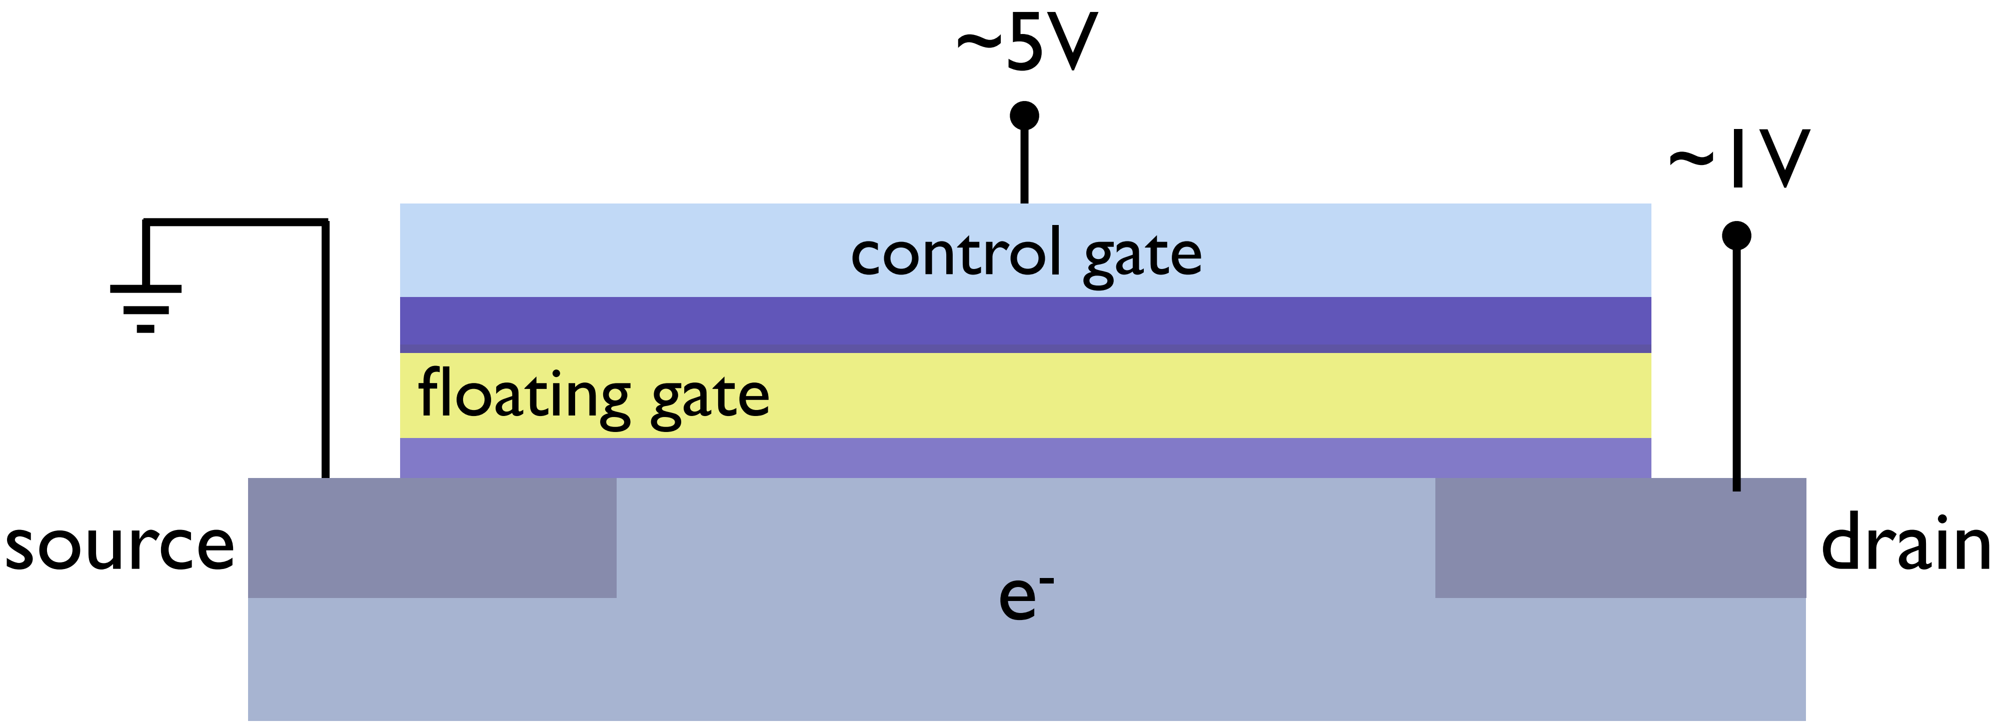

- 5V to control
- 1V to drain
- floating gate charged = channel is pinched off = "0"
- floating gate *discharged* = channel open = "1"

presence of charge modulates I_{SD} !

reading

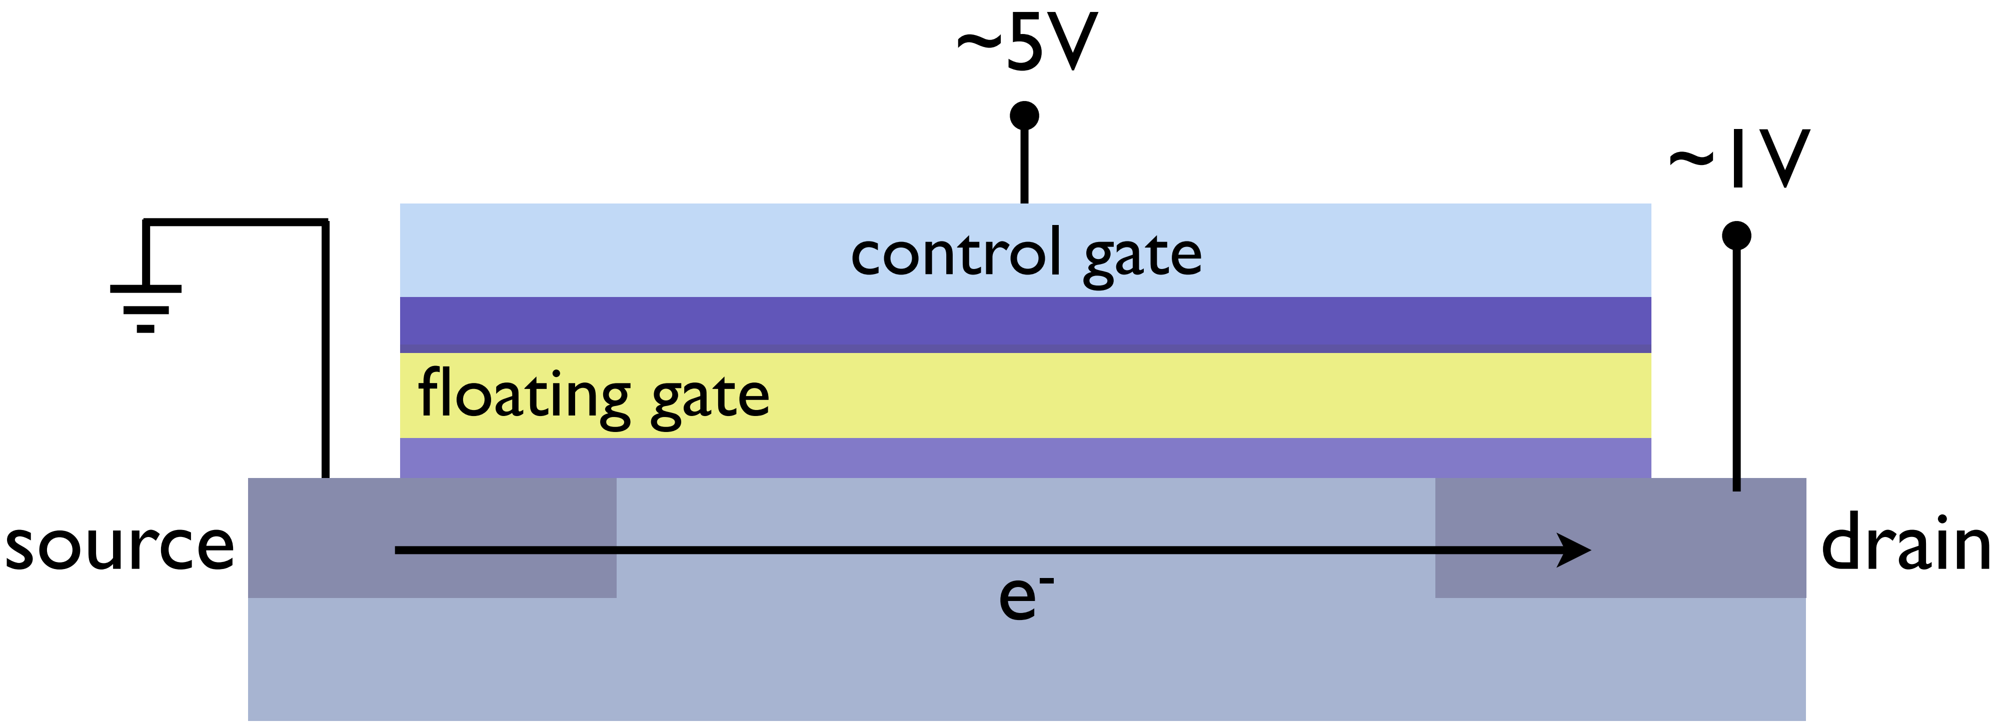

- 5V to control
- 1V to drain
- floating gate charged = channel is pinched off = "0"
- floating gate *discharged* = channel open = "1"

presence of charge modulates I_{SD} !

the basics of Flash

the basics of Flash

+ no mechanical limitations
+ lower latency

  = attractive for speed, noise, power consumption, reliability.

the basics of Flash

+ no mechanical limitations

+ lower latency
  = attractive for speed, noise, power consumption, reliability.

− cost/GB still significantly higher (but decreasing rapidly!)

− finite number of erase/write (typically 10^6 cycles guaranteed)
  unable to support an OS (swap!)
  warranties on flash-based disks trending ≥ HDD

Transcript

http://bama.ua.edu/~pleclair/PH587/

PDF version of these slides