Memory Technology

Oct. 17, 1997

Topics

- Memory Hierarchy Basics
- Static RAM
- Dynamic RAM
- Magnetic Disks
- Access Time Gap

Computer System

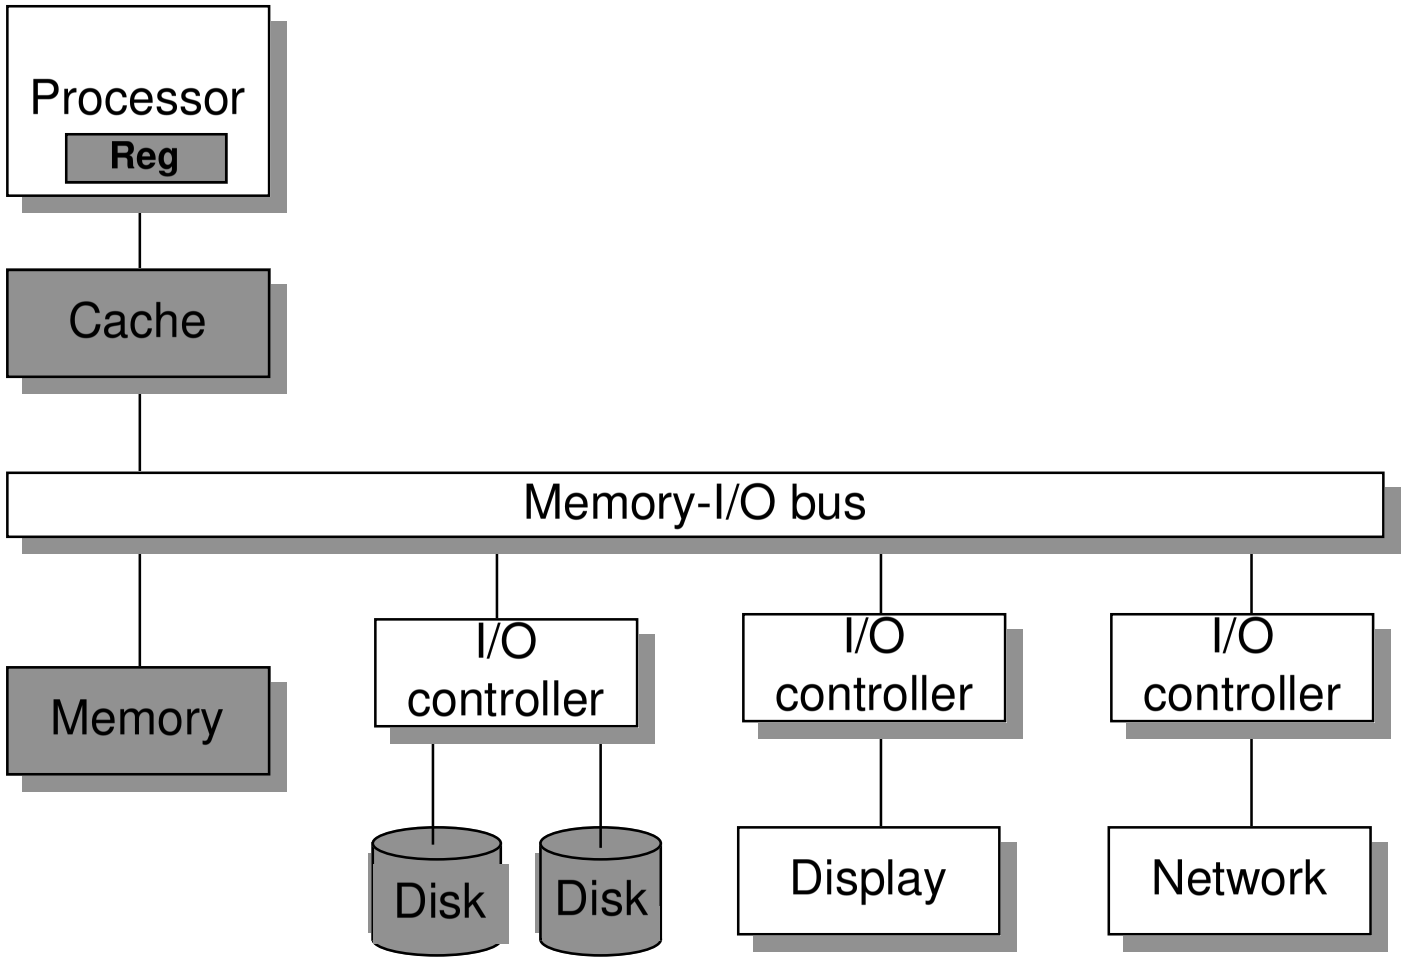

Levels in a typical memory hierarchy

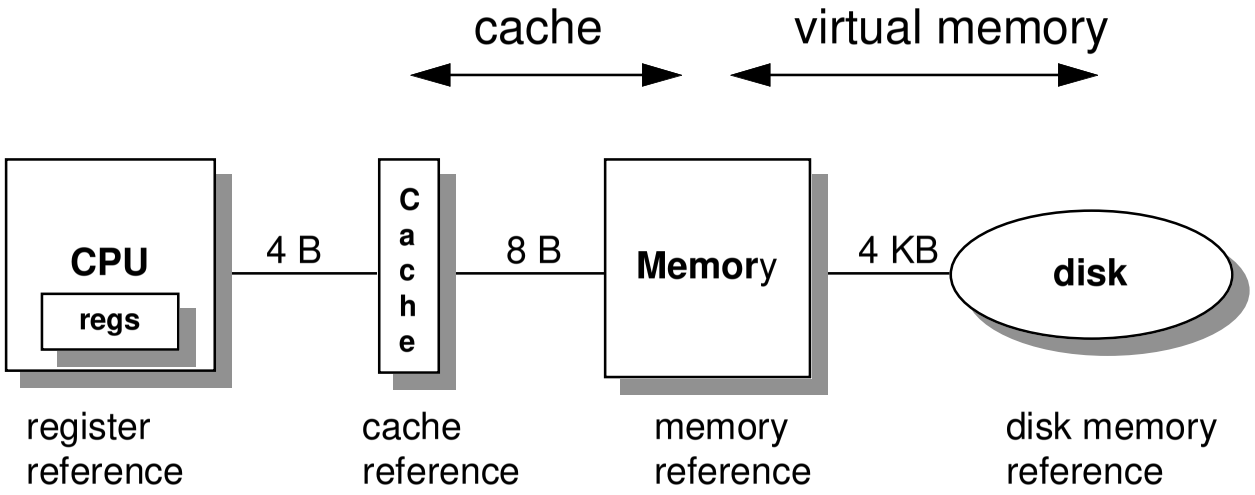

size:	200 B	64 KB	64 MB	2 GB
speed:	5 ns	10 ns	70 ns	10 ms
$/Mbyte:		$200/MB	$4/MB	$0.08/MB
block size:	4 B	8 B	4 KB	

larger, slower, cheaper

Scaling to 0.1μm

- **Semiconductor Industry Association, 1992 Technology Workshop**

Year	1992	1995	1998	2001	2004	2007
Feature size	0.5	0.35	0.25	0.18	0.12	0.10
DRAM cap	16M	64M	256M	1G	4G	16G
Gates/chip	300K	800K	2M	5M	10M	20M
Chip cm^2	2.5	4.0	6.0	8.0	10.0	12.5
I/Os	500	750	1500	2000	3500	5000
off chip MHz	60	100	175	250	350	500
on chip MHz	120	200	350	500	700	1000

Static RAM (SRAM)

Fast

- ~10 ns [1997]

Persistent

- as long as power is supplied
- no refresh required

Expensive

- ~$200/MByte [1997]
- 6 transistors/bit

Stable

- High immunity to noise and environmental disturbances

Technology for caches

Anatomy of an SRAM Cell

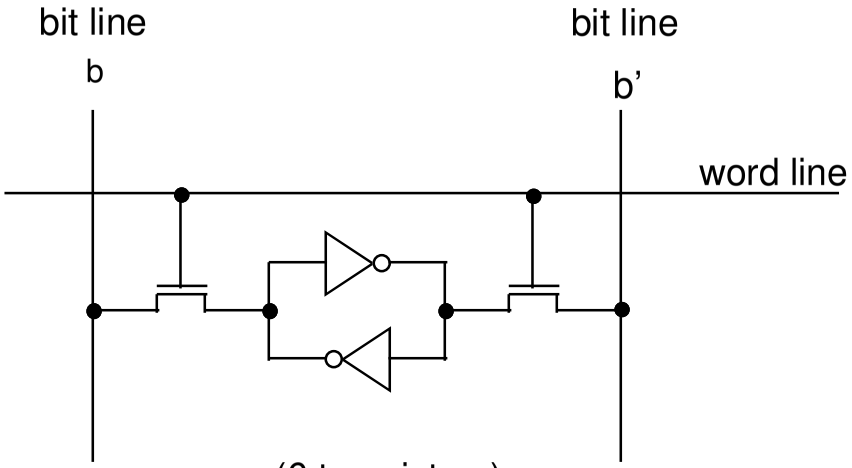

Read:
- set bit lines high
- set word line high
- see which bit line goes low

Write:
- set bit lines to opposite values
- set word line
- Flip cell to new state

SRAM Cell Principle

Inverter Amplifies

- Negative gain
- Slope < –1 in middle
- Saturates at ends

Inverter Pair Amplifies

- Positive gain
- Slope > 1 in middle
- Saturates at ends

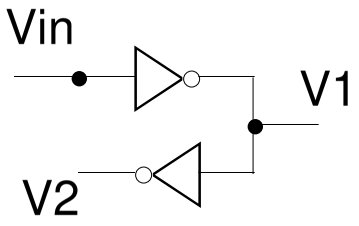

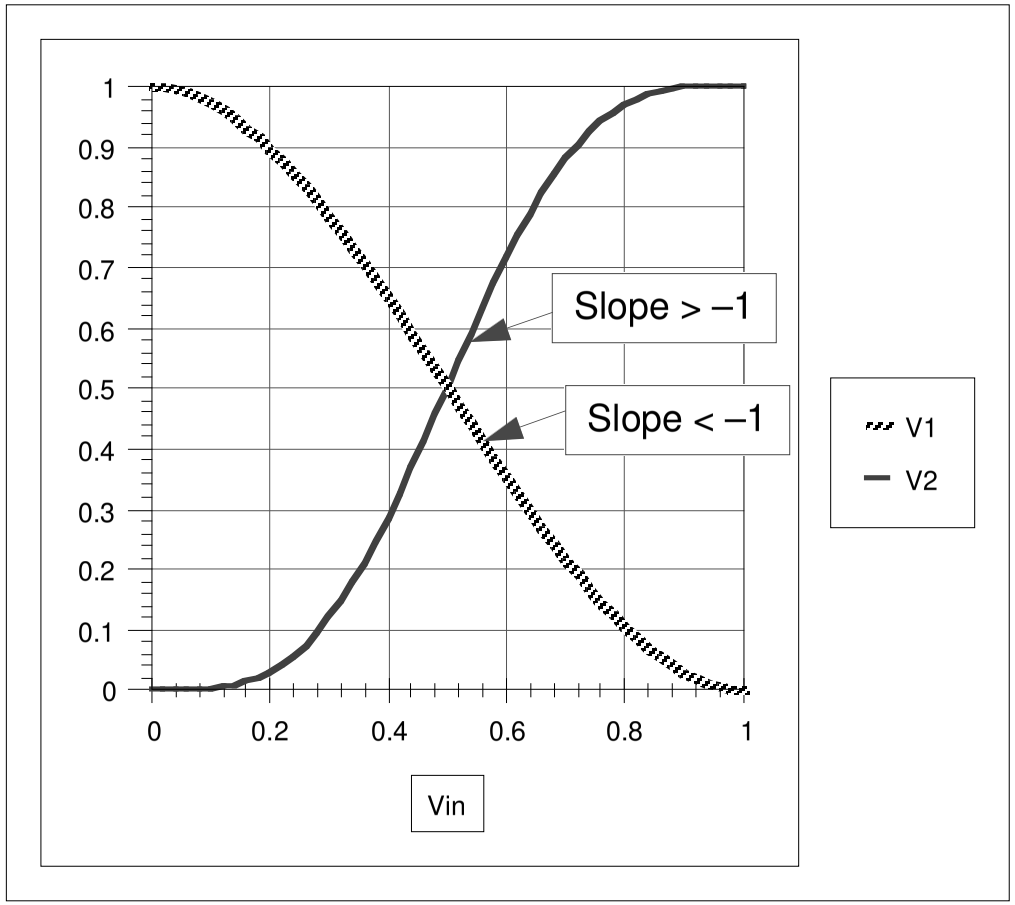

Bistable Element

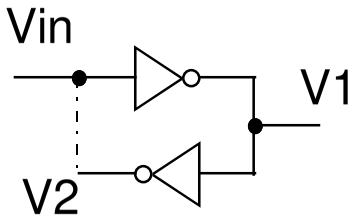

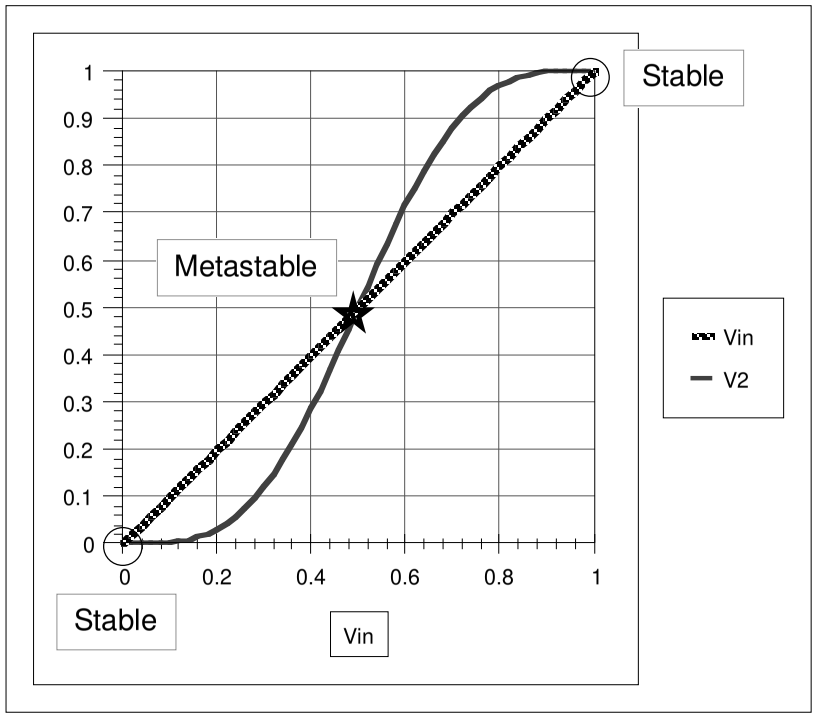

Stability

- **Require Vin = V2**
- **Stable at endpoints**
  - recover from pertubation
- **Metastable in middle**
  - Fall out when perturbed

Ball on Ramp Analogy

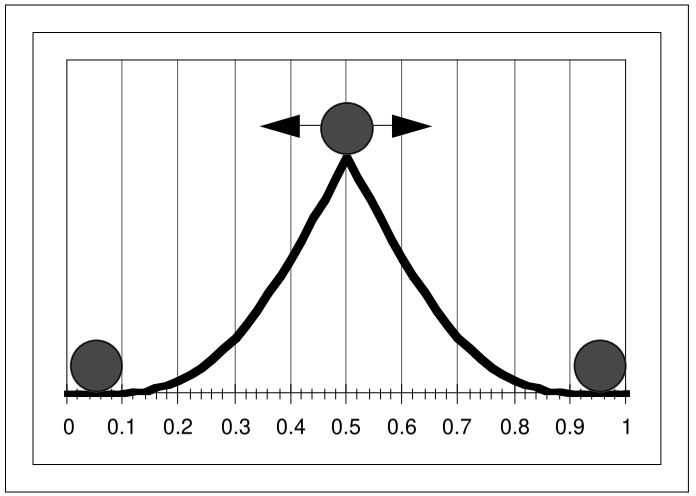

Example 1-level-decode SRAM (16 x 8)

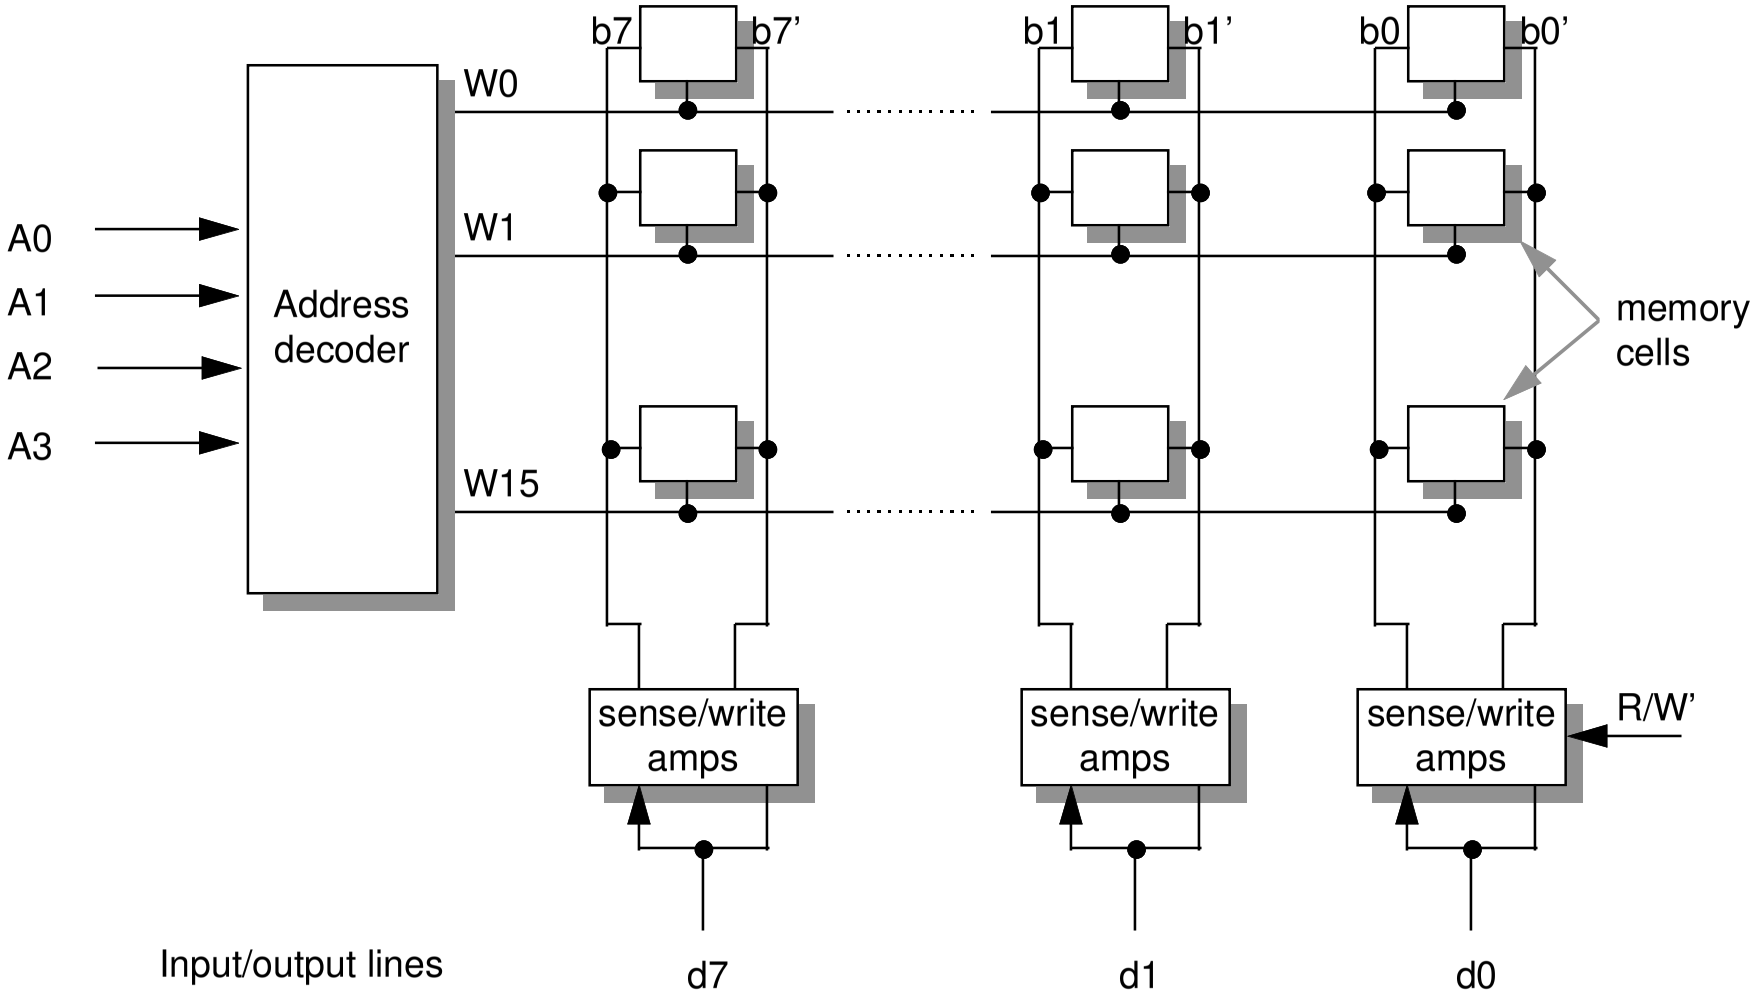

Dynamic RAM (DRAM)

Slower than SRAM

- access time ~70 ns [1997]

Not persistent

- every row must be accessed every ~1 ms (refreshed)

Cheaper than SRAM

- ~$4/MByte [1997]
- 1 transistor/bit

Fragile

- electrical noise, light, radiation

Workhorse memory technology

Anatomy of a DRAM Cell

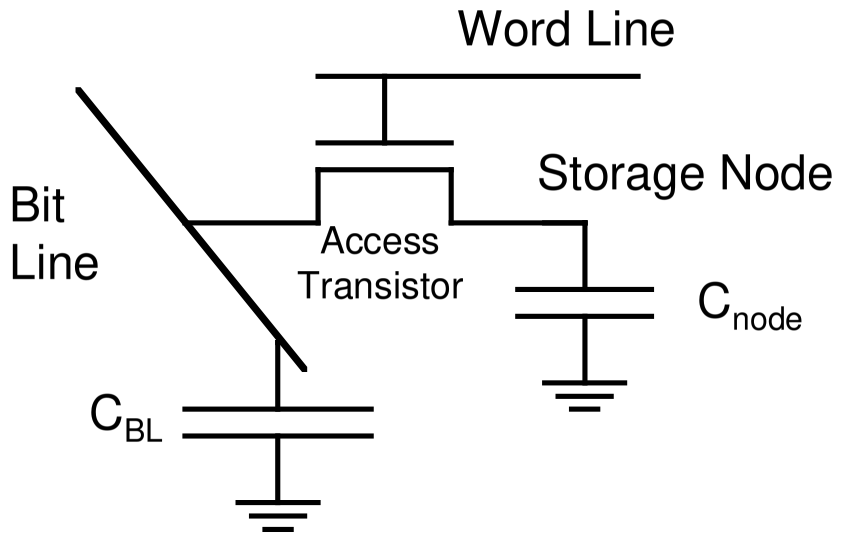

Writing

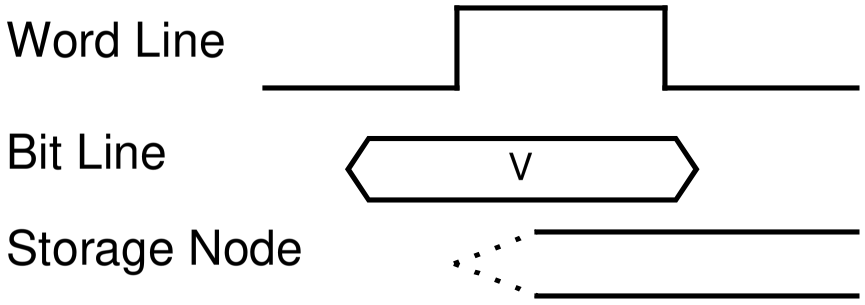

Reading

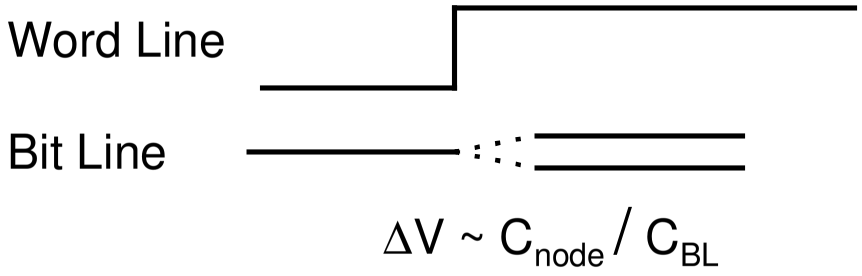

Addressing arrays with bits

Consider an R x C array of addresses, where R = 2^r and C = 2^c.
Then for each address,
row(address) = address / C = leftmost r bits of address
col(address) = address % C = righmost c bits of address

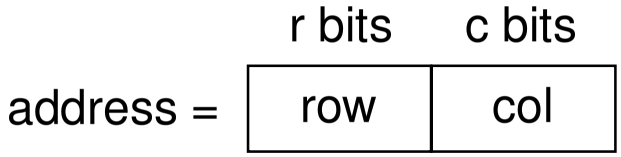

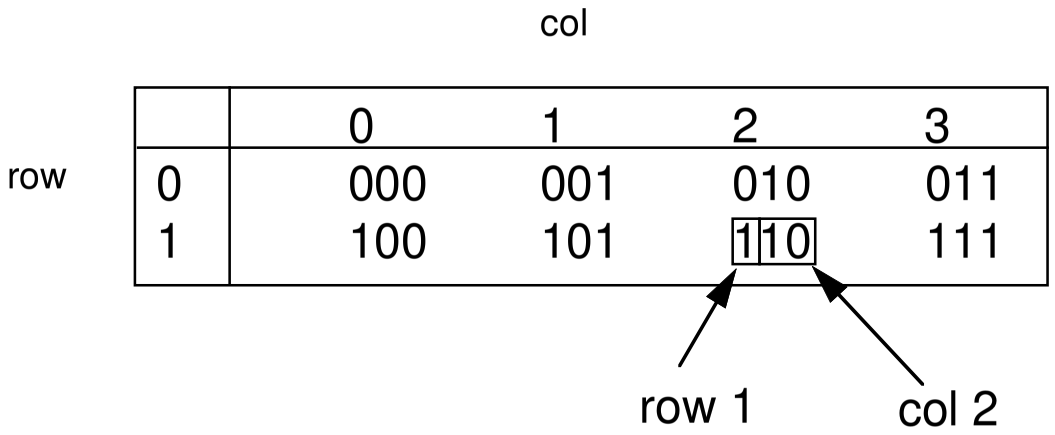

Example 2-level decode DRAM (64Kx1)

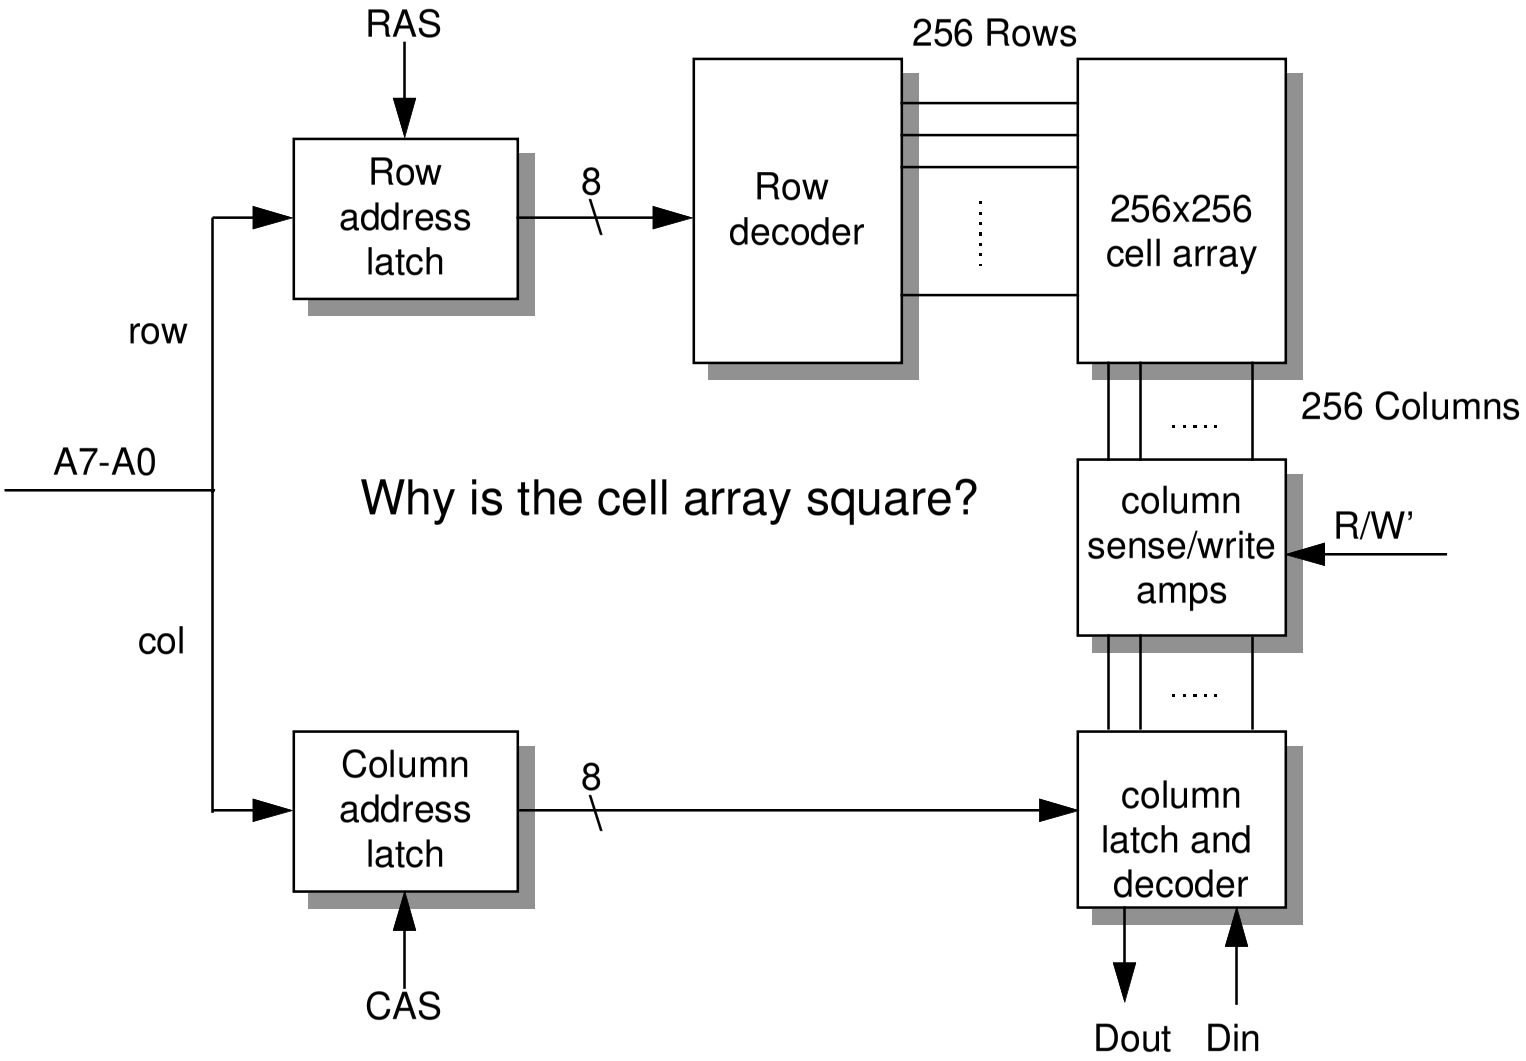

DRAM Operation

Row Address (~50ns)

- **Set Row address on address lines & strobe RAS**
- **Entire row read & stored in column latches**
- **Contents of row of memory cells destroyed**

Column Address (~10ns)

- **Set Column address on address lines & strobe CAS**
- **Access selected bit**
  - READ: transfer from selected column latch to Dout
  - WRITE: Set selected column latch to Din

Rewrite (~30ns)

- **Write back entire row**

Observations About DRAMs

Timing

- Access time = 60ns < cycle time = 90ns
- Need to rewrite row

Must Refresh Periodically

- Perform complete memory cycle for each row
- Approx. every 1ms
- Sqrt(n) cycles
- Handled in background by memory controller

Inefficient Way to Get Single Bit

- Effectively read entire row of Sqrt(n) bits

Enhanced Performance DRAMs

Conventional Access

- **Row + Col**
- **RAS CAS RAS CAS ...**

Page Mode

- **Row + Series of columns**
- **RAS CAS CAS CAS ...**
- **Gives successive bits**

Video RAM

- **Shift out entire row sequentially**
- **At video rate**

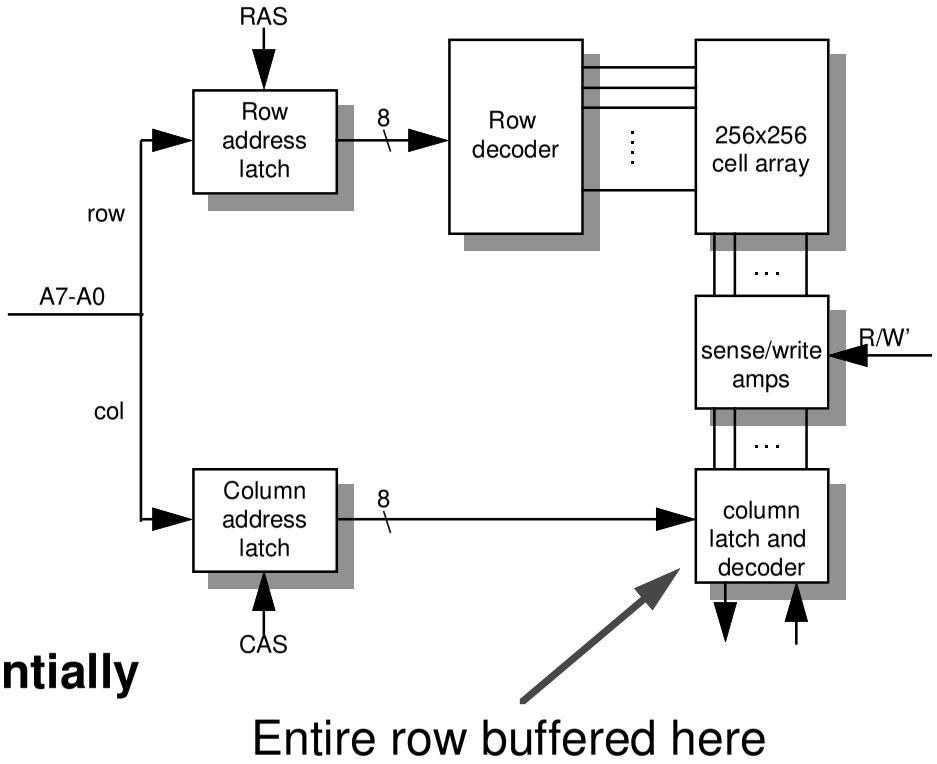

Typical Performance

row access time	col access time	cycle time	page mode cycle time
50ns	10ns	90ns	25ns

DRAM Driving Forces

Capacity

- **4X per generation**
  - Square array of cells
- **Typical scaling**
  - Lithography dimensions 0.7X
    - » Areal density 2X
  - Cell function packing 1.5X
  - Chip area 1.33X
- **Scaling challenge**
  - Typically $C_{node} / C_{BL} = 0.1–0.2$
  - Must keep C_{node} high as shrink cell size

Retention Time

- **Typically 16–256 ms**
- **Want higher for low-power applications**

DRAM Storage Capacitor

Planar Capacitor

- Up to 1Mb
- C decreases linearly with feature size

Trench Capacitor

- 4–256 Mb
- Lining of hole in substrate

Stacked Cell

- > 1Gb
- On top of substrate
- Use high ε dielectric

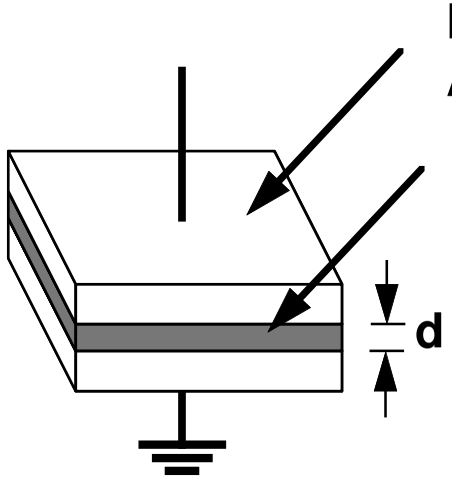

$$C = \varepsilon A/d$$

Trench Capacitor

Process

- **Etch deep hole in substrate**
  - Becomes reference plate
- **Grow oxide on walls**
  - Dielectric
- **Fill with polysilicon plug**
  - Tied to storage node

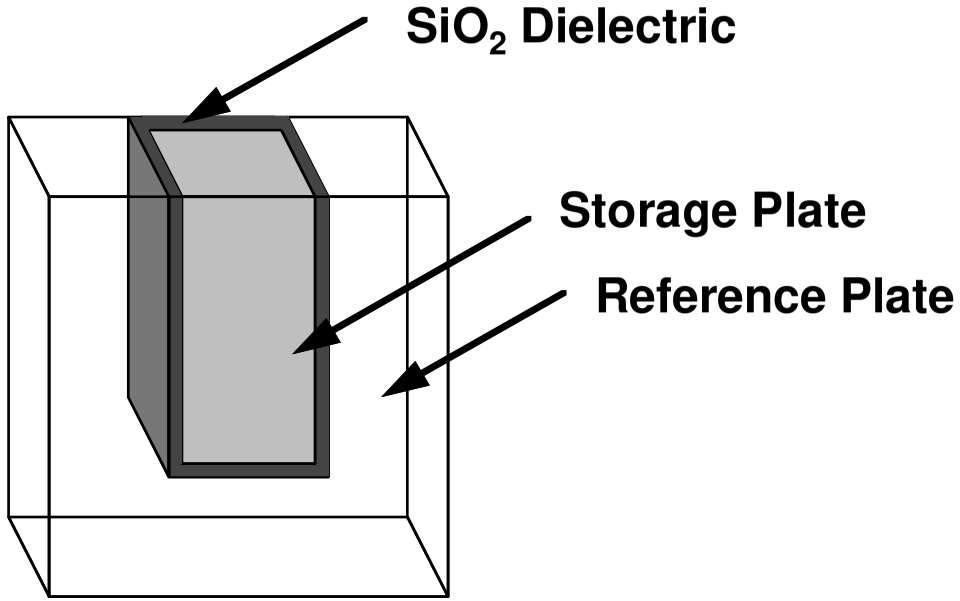

IBM DRAM Evolution

- IBM J. R&D, Jan/Mar '95
- Evolution from 4 – 256 Mb
- 256 Mb uses cell with area 0.6 μm^2

4 Mb Cell Structure

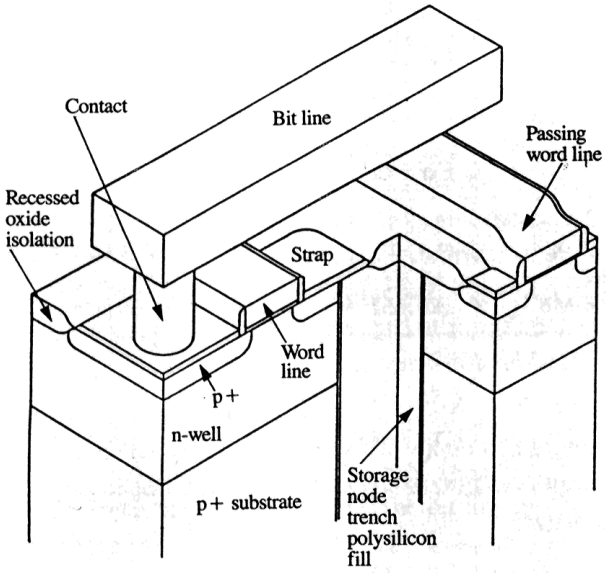

Cell Layouts

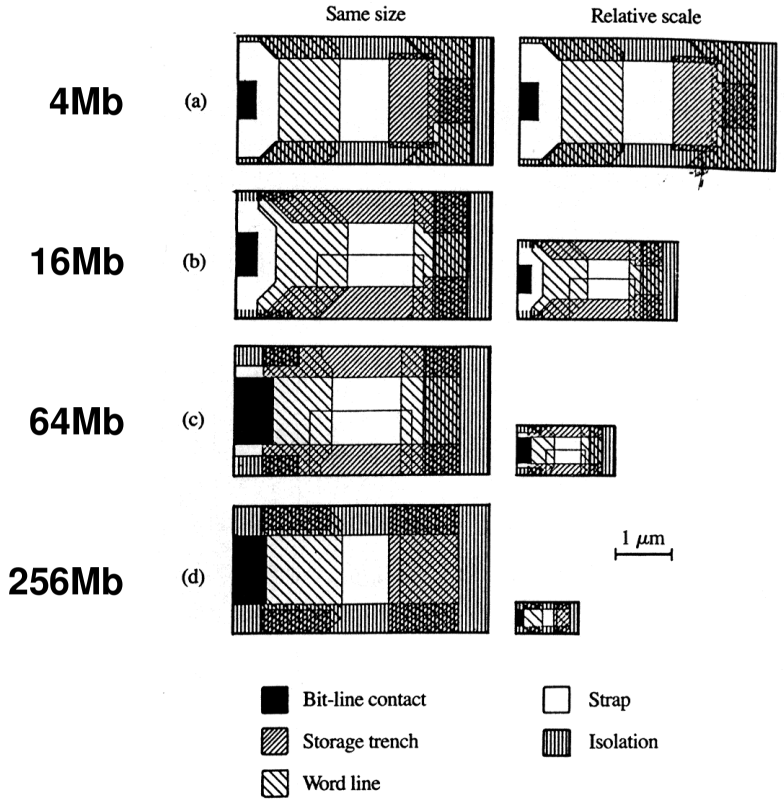

Mitsubishi Stacked Cell DRAM

- **IEDM '95**
- **Claim suitable for 1 – 4 Gb**

Technology

- **0.14 μm process**
  - Synchrotron X-ray source
- **8 nm gate oxide**
- **0.29 μm^2 cell**

Storage Capacitor

- **Fabricated on top of everything else**
- **Rubidium electrodes**
- **High dielectric insulator**
  - 50X higher than SiO_2
  - 25 nm thick
- **Cell capacitance 25 femtofarads**

Cross Section of 2 Cells

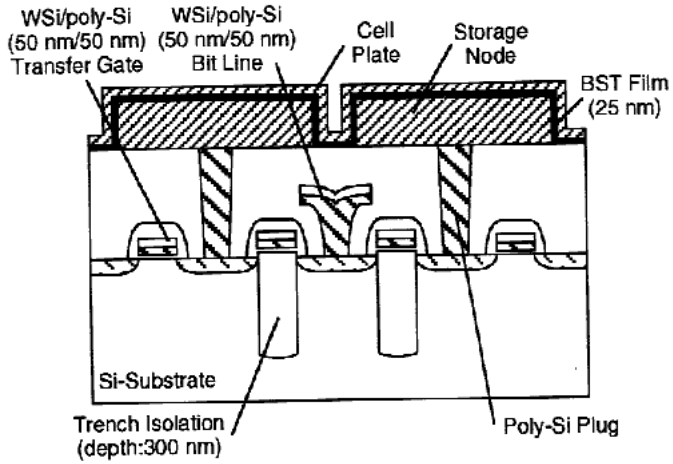

Fig. 2 Schematic cross-sectional view of DRAM memory cells with Ru/BST/Ru stacked capacitors.

Mitsubishi DRAM Pictures

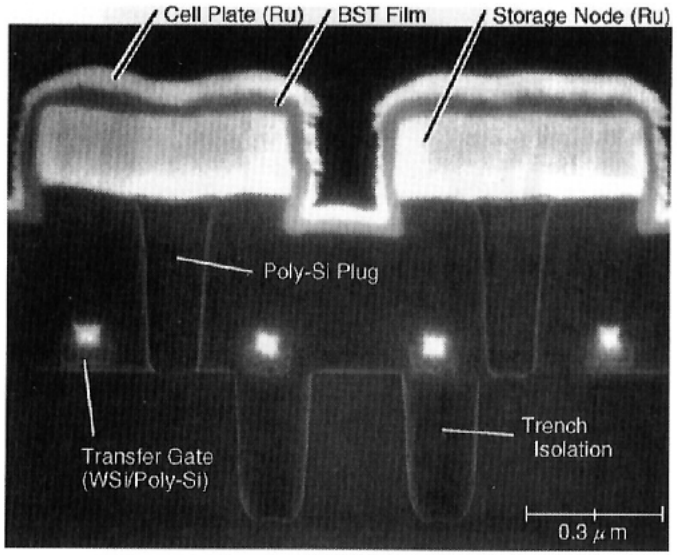

Fig. 3 SEM cross-sectional photograph of the fabricated 0.29-μm² memory cell with Ru/BST/Ru stacked capacitor. The facet was fabricated by focused ion beam etching.

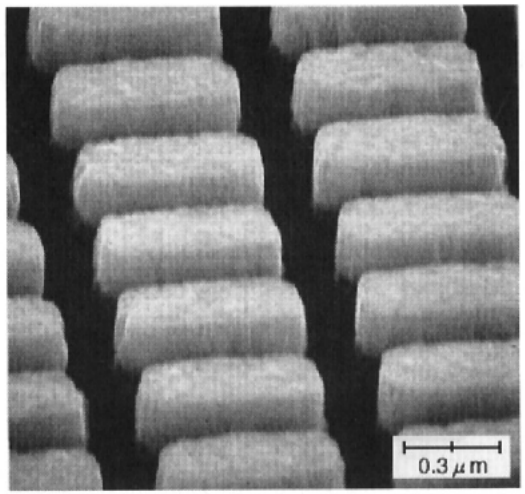

Fig. 8 SEM photograph of a Ru-metal storage node array with a projection a height of 0.2 μm.

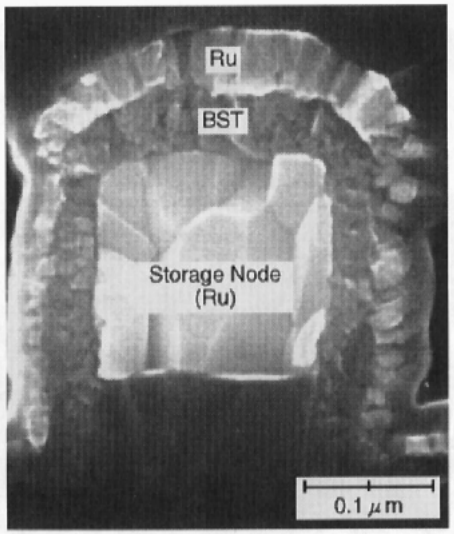

Fig. 10 SEM cross-sectional view of a Ru/BST/Ru capacitor cell. The facet shown is a cleaved facet.

Magnetic Disks

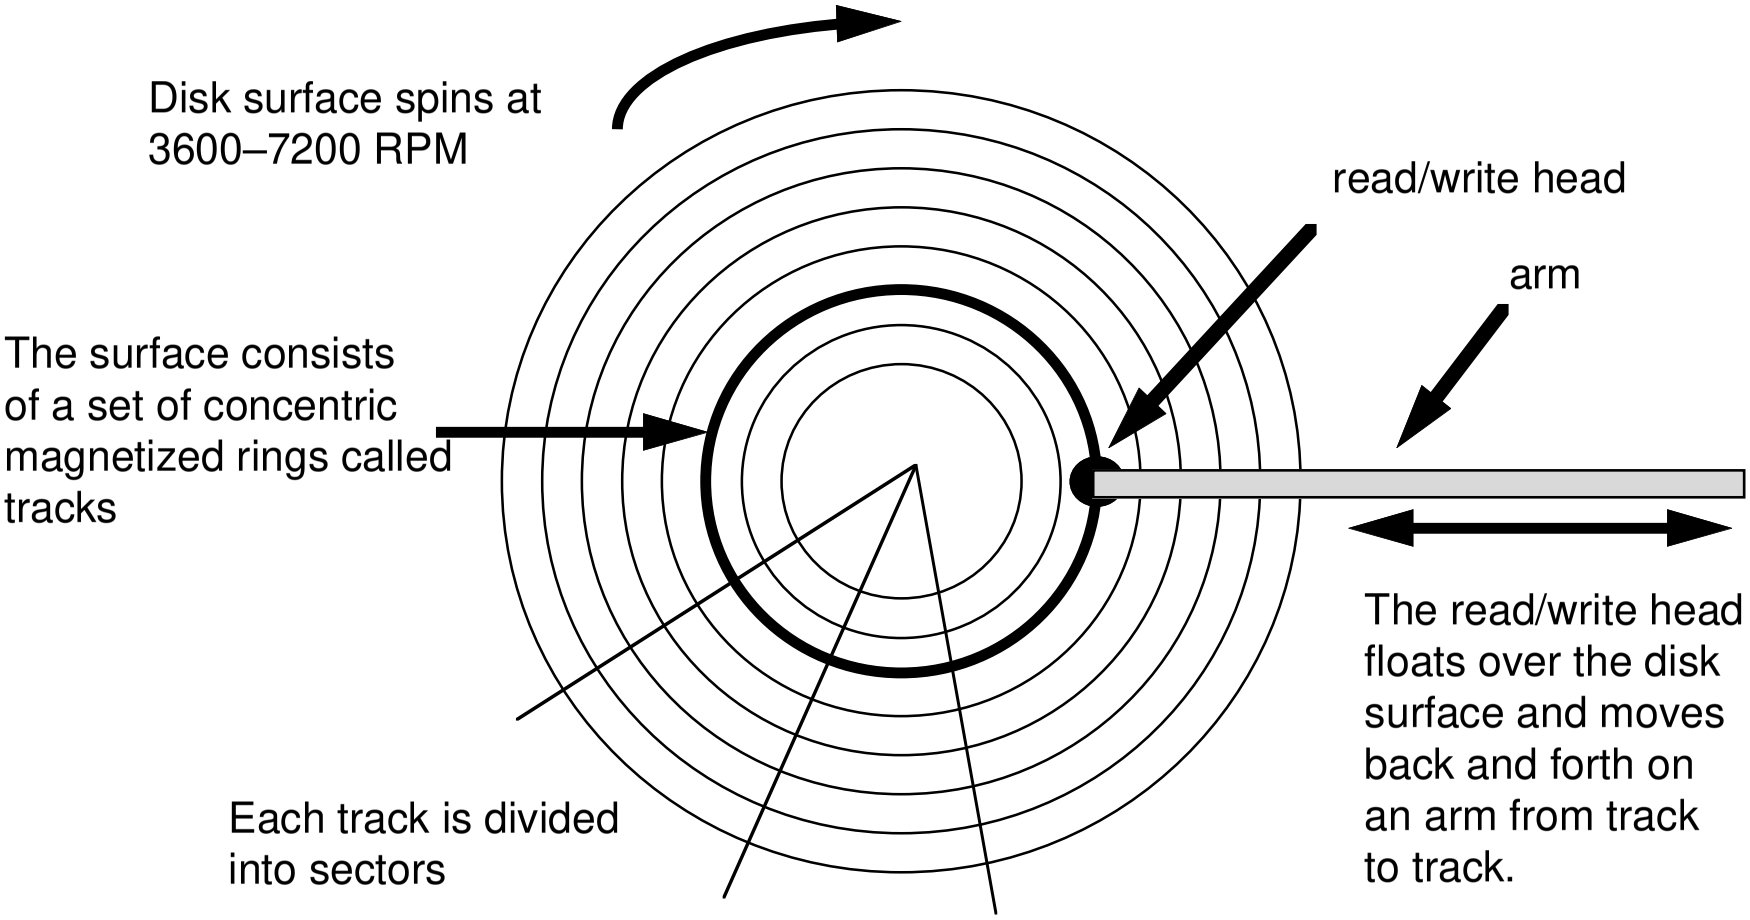

Disk Capacity

Parameter	540MB Example
• Number Platters	8
• Surfaces / Platter	2
• Number of tracks	1046
• Number sectors / track	63
• Bytes / sector	512
Total Bytes	**539,836,416**

Disk Operation

Operation

- **Read or write complete sector**

Seek

- **Position head over proper track**
- **Typically 10ms**

Rotational Latency

- **Wait until desired sector passes under head**
- **Worst case: complete rotation**
  - 3600RPM: 16.7 ms

Read or Write Bits

- **Transfer rate depends on # bits per track and rotational speed**
- **E.g., 63 * 512 bytes @3600RPM = 1.9 MB/sec.**

Disk Performance

Getting First Byte

- **Seek + Rotational latency 10,000 – 27,000 microseconds**

Getting Successive Bytes

- **~ 0.5 microseconds each**

Optimizing

- **Large block transfers more efficient**
- **Try to do other things while waiting for first byte**
  - Switch context to other computing task
  - Disk controller buffers sector
  - Interrupts processor when transfer completed

Disk Technology

Seagate ST-12550N Barracuda 2 Disk

• **Linear density**	**52,187.**	**bits per inch (BPI)**
– Bit spacing	0.5	microns
• **Track density**	**3,047.**	**tracks per inch (TPI)**
– Track spacing	8.3	microns
• **Total tracks**	**2,707.**	**tracks**
• **Rotational Speed**	**7200.**	**RPM**
• **Avg Linear Speed**	**86.4**	**kilometers / hour**
• **Head Floating Height**	**0.13**	**microns**

Analogy

- **Put Sears Tower on side**
- **Fly around world 2.5 cm off ground**
- **8 seconds per orbit**

Storage trends (memory)

SRAM

metric	1980	1985	1990	1995	1995:1980
$/MB	19,200	2,900	320	256	75
access (ns)	300	150	35	15	20

DRAM

metric	1980	1985	1990	1995	1995:1980
$/MB	8,000	880	100	30	266
access (ns)	375	200	100	70	5
typical size(MB)	0.064	0.256	4	16	250

culled from back issues of Byte and PC Magazine

Storage trends (disk)

Disks

metric	1980	1985	1990	1995	1995:1980
$/MB	500	100	8	0.30	1,600
access (ms)	87	75	28	10	9
typical size(MB)	1	10	160	1,000	1,000

culled from back issues of Byte and PC Magazine

Storage price/MByte

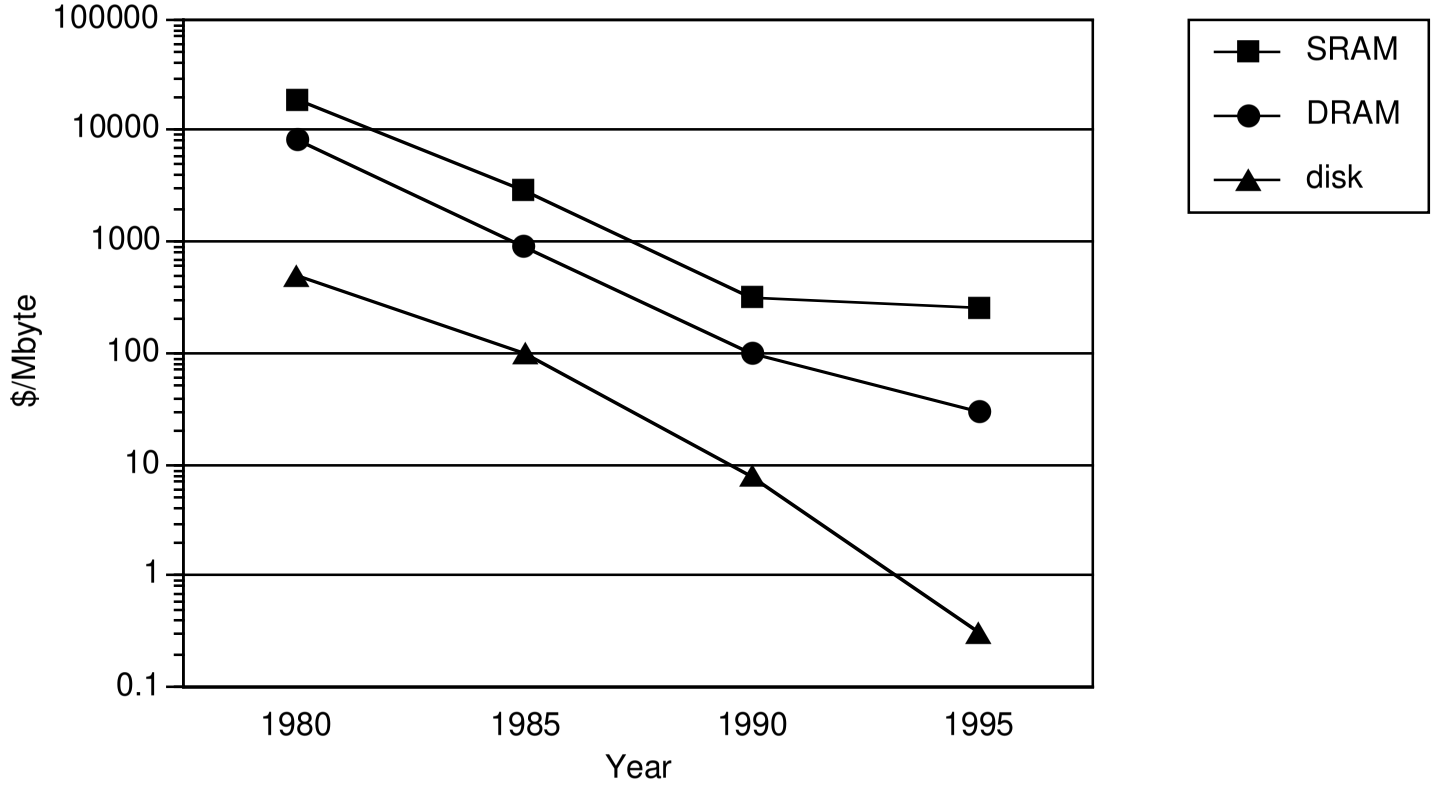

culled from back issues of Byte and PC Magazine

Storage access times

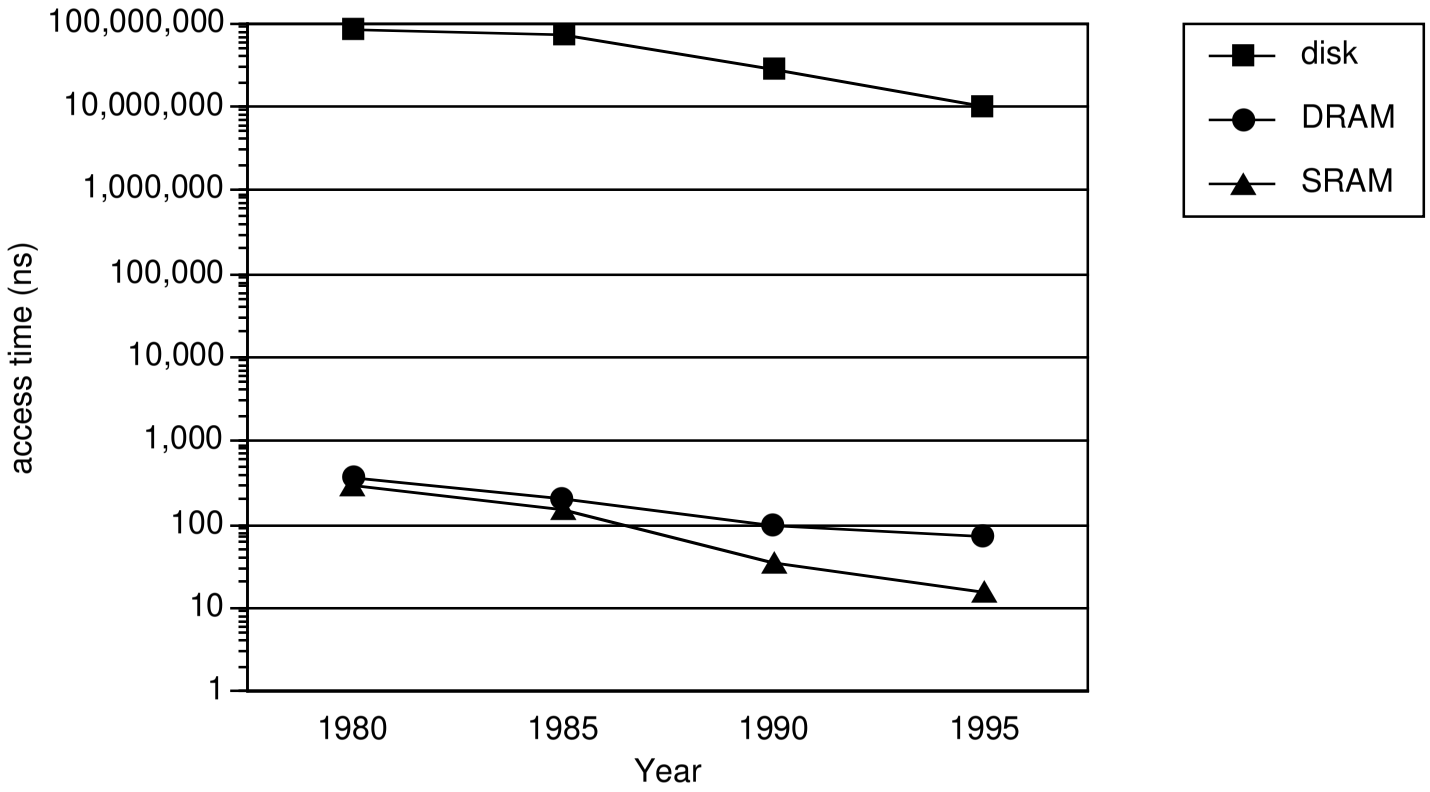

culled from back issues of Byte and PC Magazine

Processor clock rates

Processors

metric	1980	1985	1990	1995	1995:1980
typical clock(MHz)	1	6	20	150	150
processor	8080	286	386	pentium	

culled from back issues of Byte and PC Magazine

The widening processor/memory gap

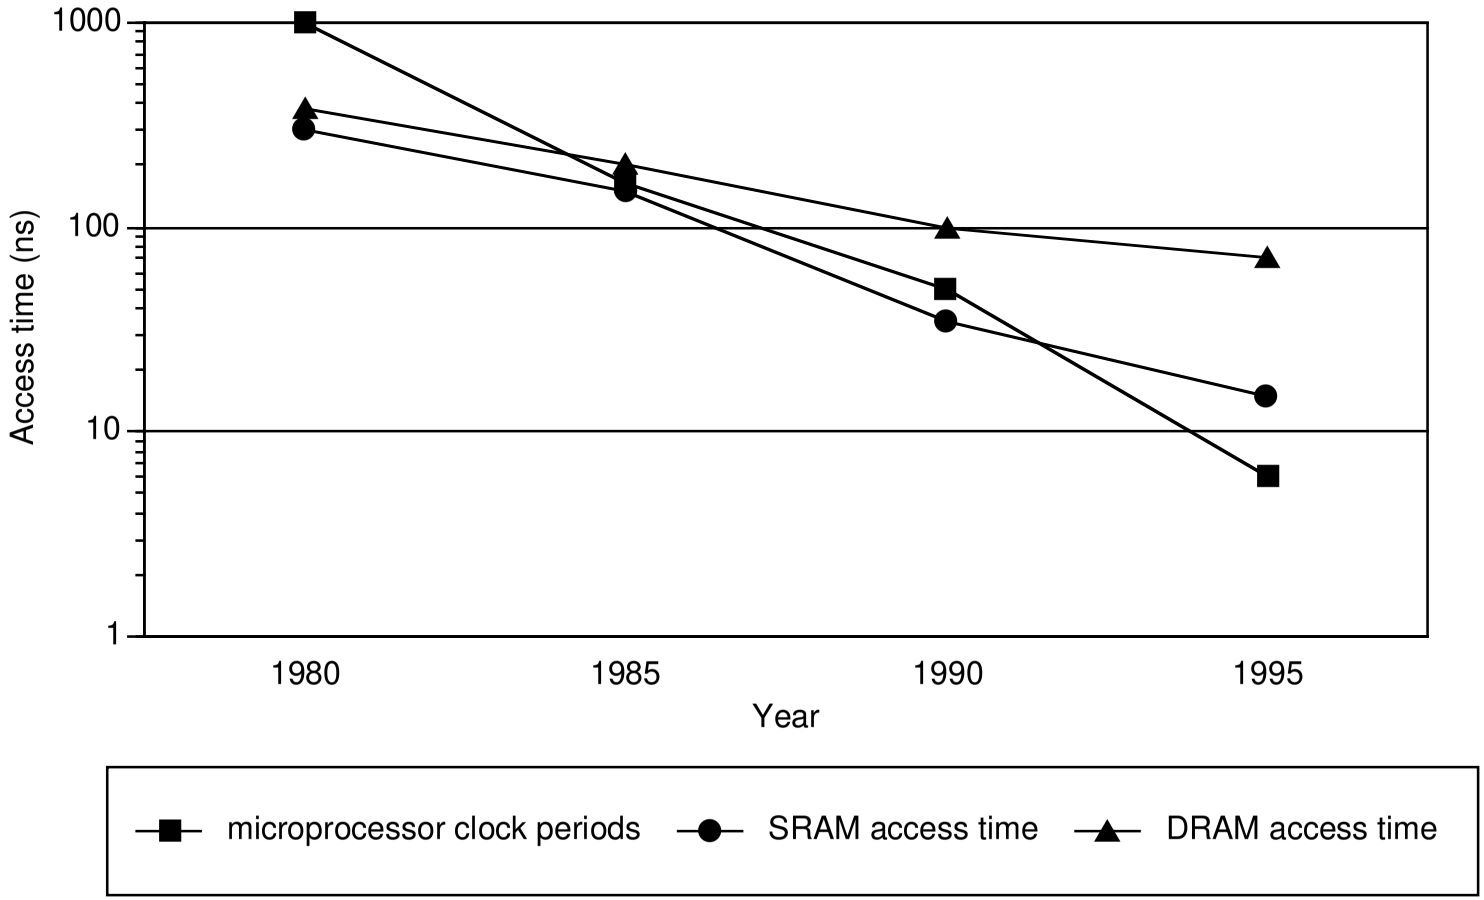

culled from back issues of Byte and PC Magazine

Memory technology summary

Cost and density improving at enormous rates.

Speed lagging processor performance

Memory hierarchies help narrow the gap:

- **small fast SRAMS (cache) at upper levels**
- **large slow DRAMS (main memory) at lower levels**
- **Incredibly large & slow disks to back it all up**

Locality of reference makes it all work

- **Keep most frequently accessed data in fastest memory**