

Memory 3.0 (Three Dot O)

Sangyeun Cho

Memory Solutions Lab, Memory Division
Samsung Electronics Co.

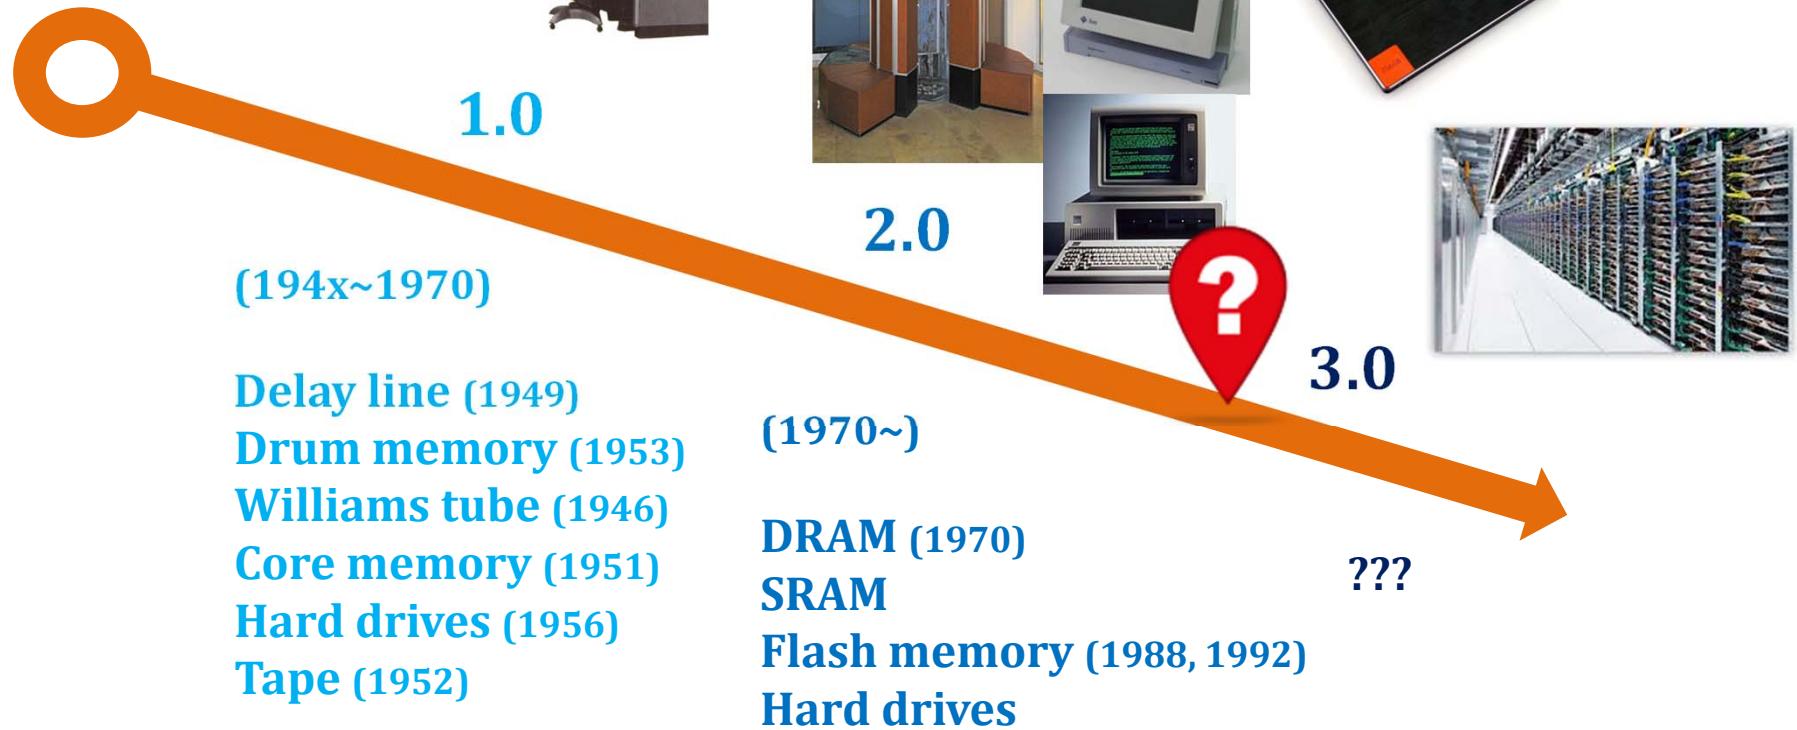
mem·o·ry *noun* \ˈmem-rē, 'me-mə-\

1 a: the power or process of reproducing or recalling what has been learned and retained especially through associative mechanisms

...

4 a: a device (as a chip) or a component of a device in which information especially for a computer can be inserted and stored and from which it may be extracted when wanted

mem·o·ry



mem·o·ry

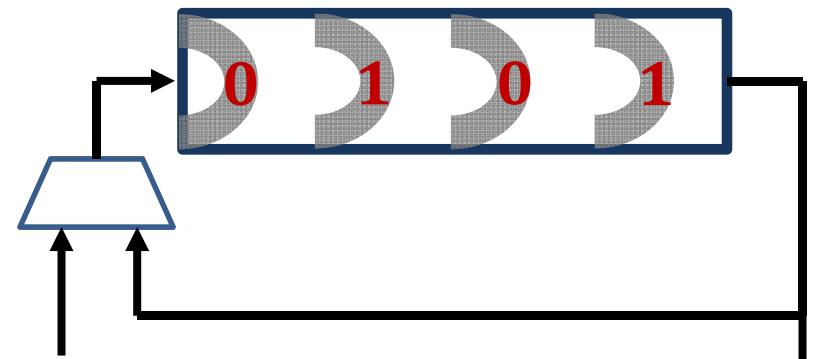
1.0

Delay line (1949)

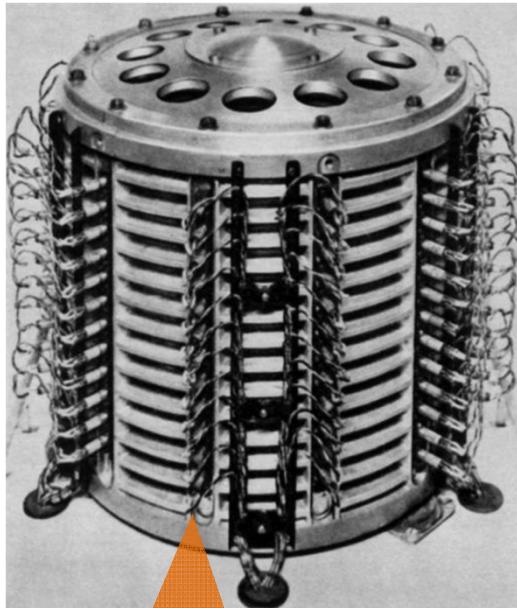
- Sonic waves are injected at one end
 - These waves propagate through the media inside the “line”
- Waves are retrieved at the other end and re-injected
 - States are preserved
 - New values can be injected instead of old values
- 100's of bits
- 100's μ sec access latency
- Address interleaving



[For UNIVAC I, 1951]



Drum memory (1953)

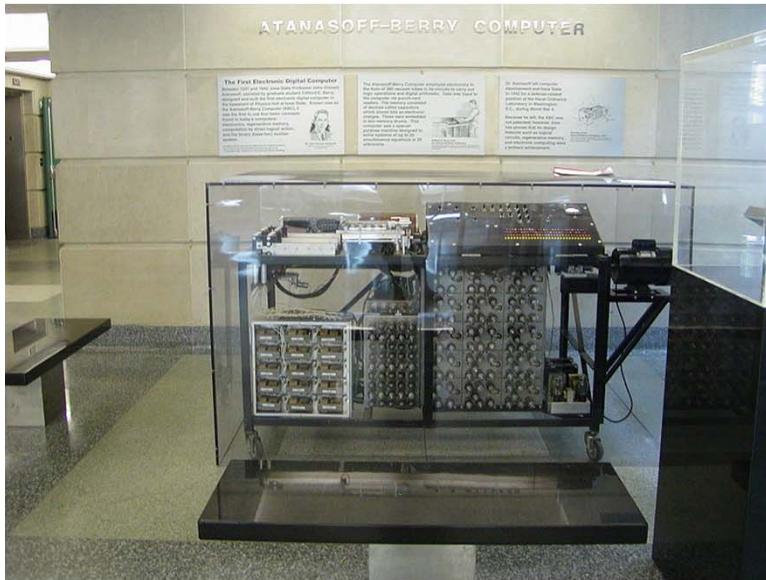


[For ZAM-41, 1961]

In BSD Unix, `/dev/drum` is the name of the default swap device

- Rotating drum (metal cylinder)
 - Many heads (fixed)
 - (Random) access time of milliseconds
 - <100KiB
 - Non-volatile
 - Rotational speed determines performance
 - Address interleaving
- Similar to the soon available hard drive technology

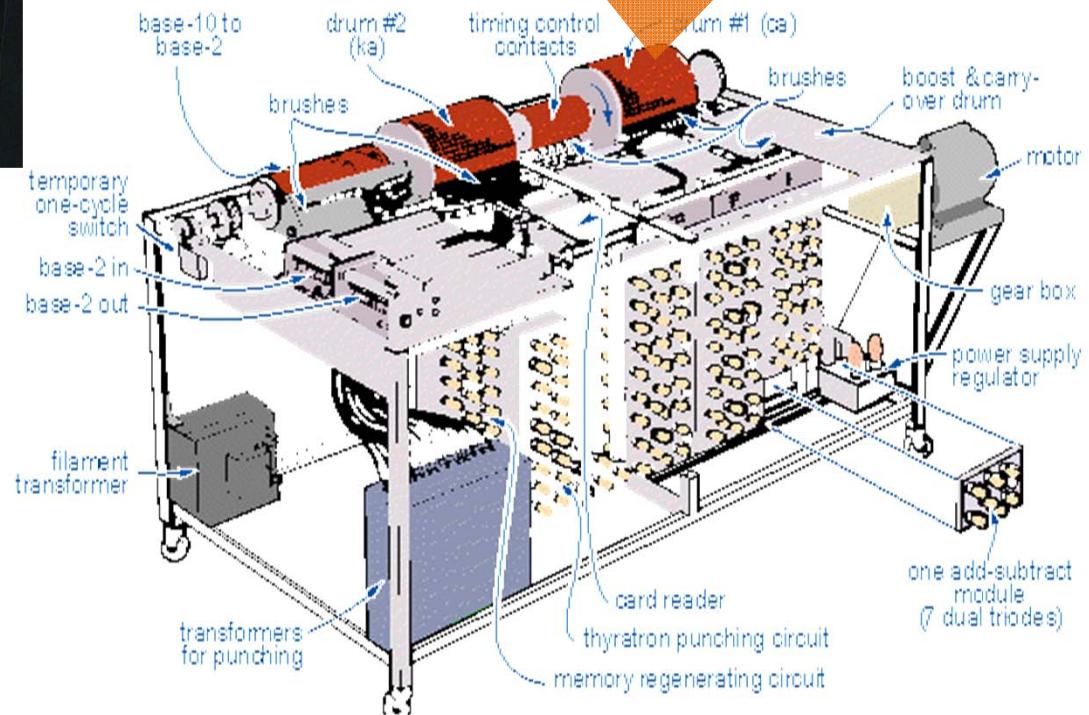
Atanasoff-Berry Computer (1942)



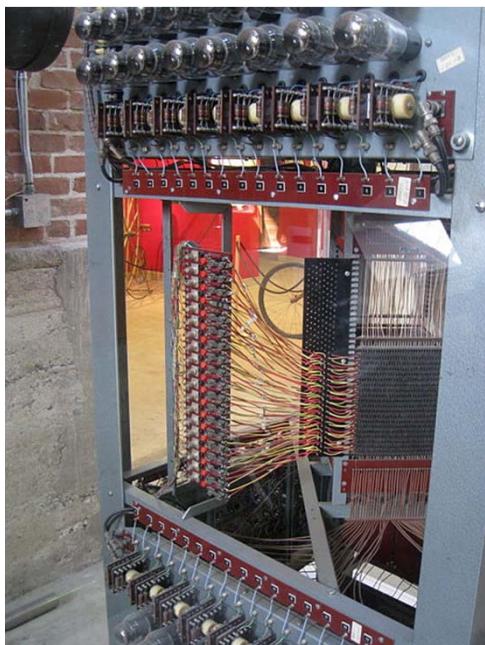
["ABC" @Iowa State University]

Each rotating drum has
1,600 capacitors, refreshed
or updated every second

The Atanasoff-Berry Computer



Core memory (1951)



[For Whirlwind, 1951]

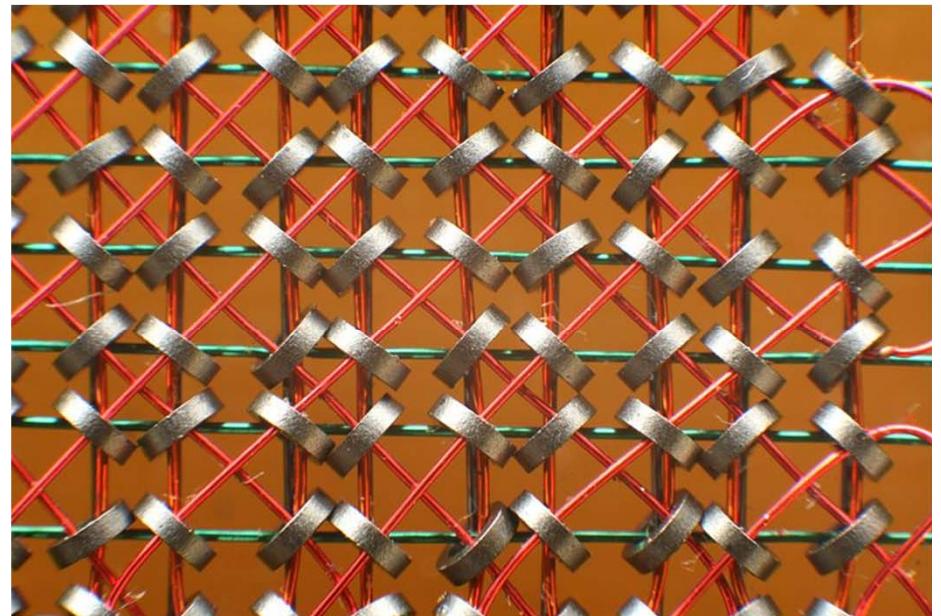
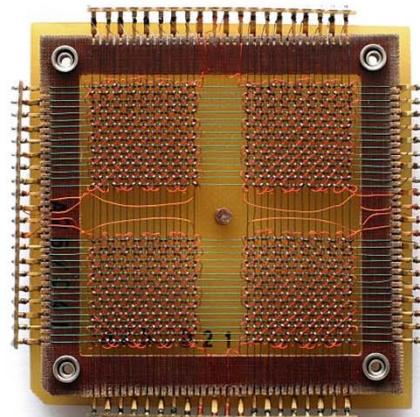


Figure 2d: Sensing the state of a core

The state of a core
no change to core field,
 $-1/2I_s/I_s$ no pulse on sense wire

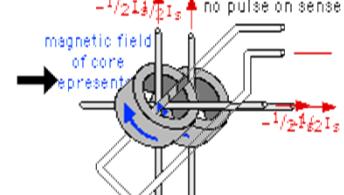
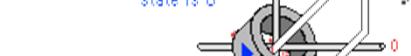


Figure 2b:
Coincident current addressing

If stored state is 0
current I_s flows in one direction



1/2I_s →
If stored state is 1
current I_s flows in the other direction

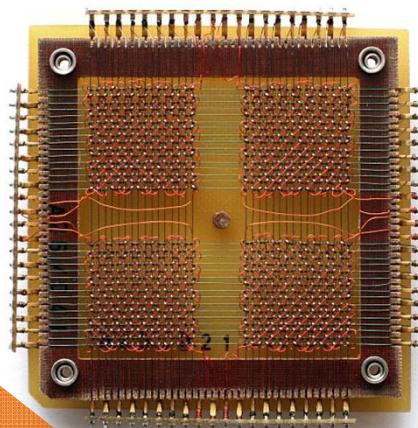


Read is destructive...
Need to reprogram after
each read



Memory 3.0

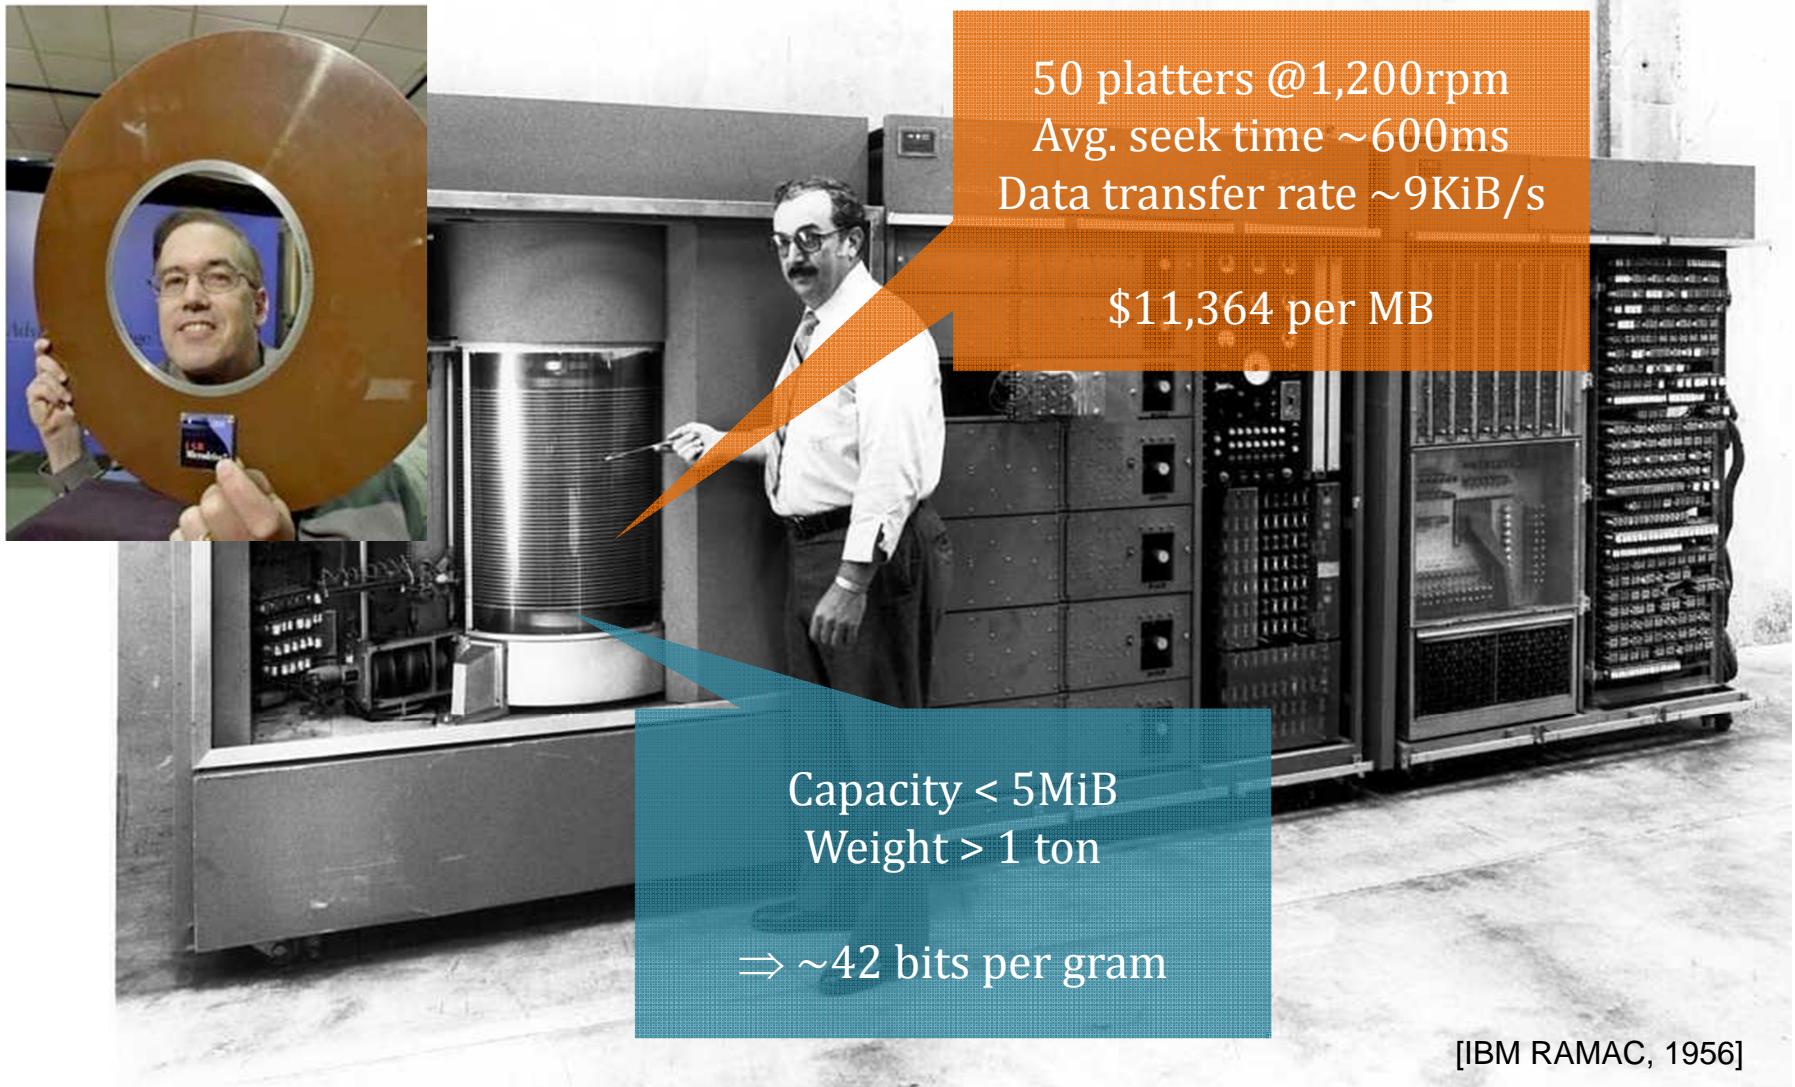
Core memory (1951)



- High density
 - \$1 per bit \Rightarrow \$0.01 per bit
- High performance
 - 1MHz clock rate
- Non-volatile
 - This property was utilized in some systems

In many systems, a dump of memory contents (after system crash) is called “core dump”

First hard drive (1956)



Summary of mem·o·ry 1.0

- Introduction of familiar concepts like:
 - Sequential access vs. random access
 - Address interleaving
 - Retention vs. refreshing
 - Destructive reading
- Births and deployment of lasting (or recurring) memory technologies like:
 - Hard drives
 - Tapes
 - Magnetic RAM
 - Capacitive storage (\Rightarrow DRAM)

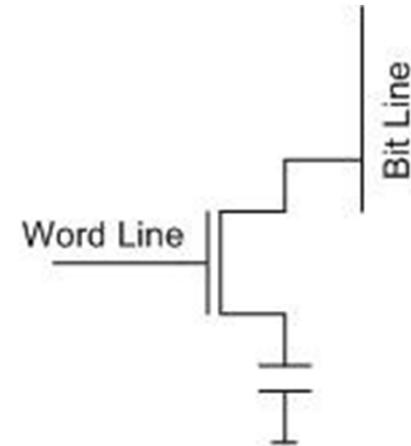
mem·o·ry 2.0

DRAM

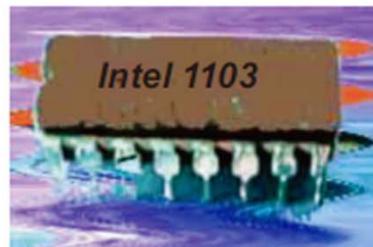


Dr. Robert H. Dennard, a Fellow at the IBM created the one-transistor DRAM in 1966

1966



World 1'st Available DRAM Chip



- .. 1 Kb
- .. 3 Transistors

1970

1'st DRAM Chip with 1-Tr. & 1-Cap.



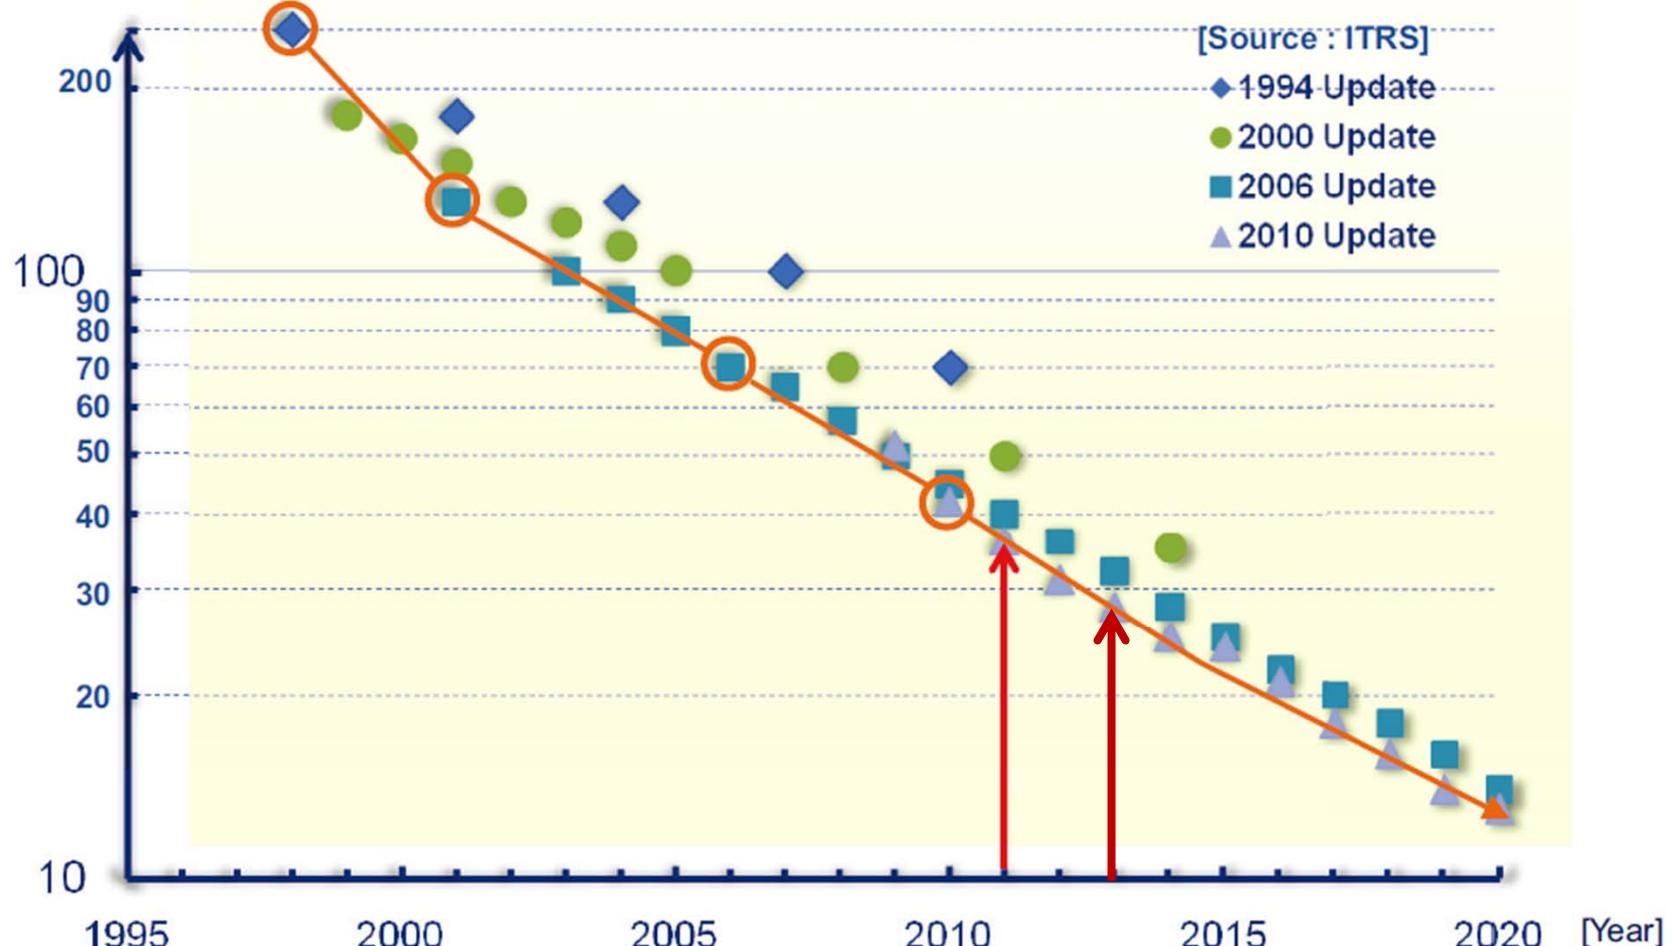
- .. 4 Kb
- .. 1-Tr. & 1-Cap.
- .. Address multiplexing
- .. VDD : 11.4~12.6V
- .. Access time of 300ns
- .. Refresh time of 2ms

1973

[Cha, 2011 VLSI Tech. Short Course]

DRAM scaling

Minimum Feature Size [nm]



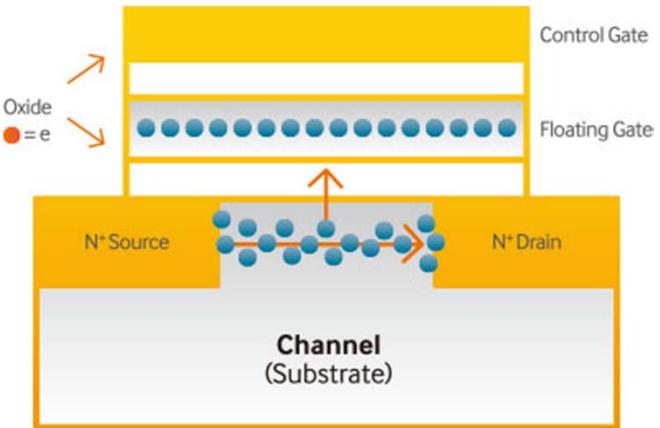
[Cha, 2011 VLSI Tech. Short Course]

NAND flash



Dr. Fujio Masaoka @Toshiba invents flash memory in 1980

1980



Intel produces first NOR flash in 1988

1988

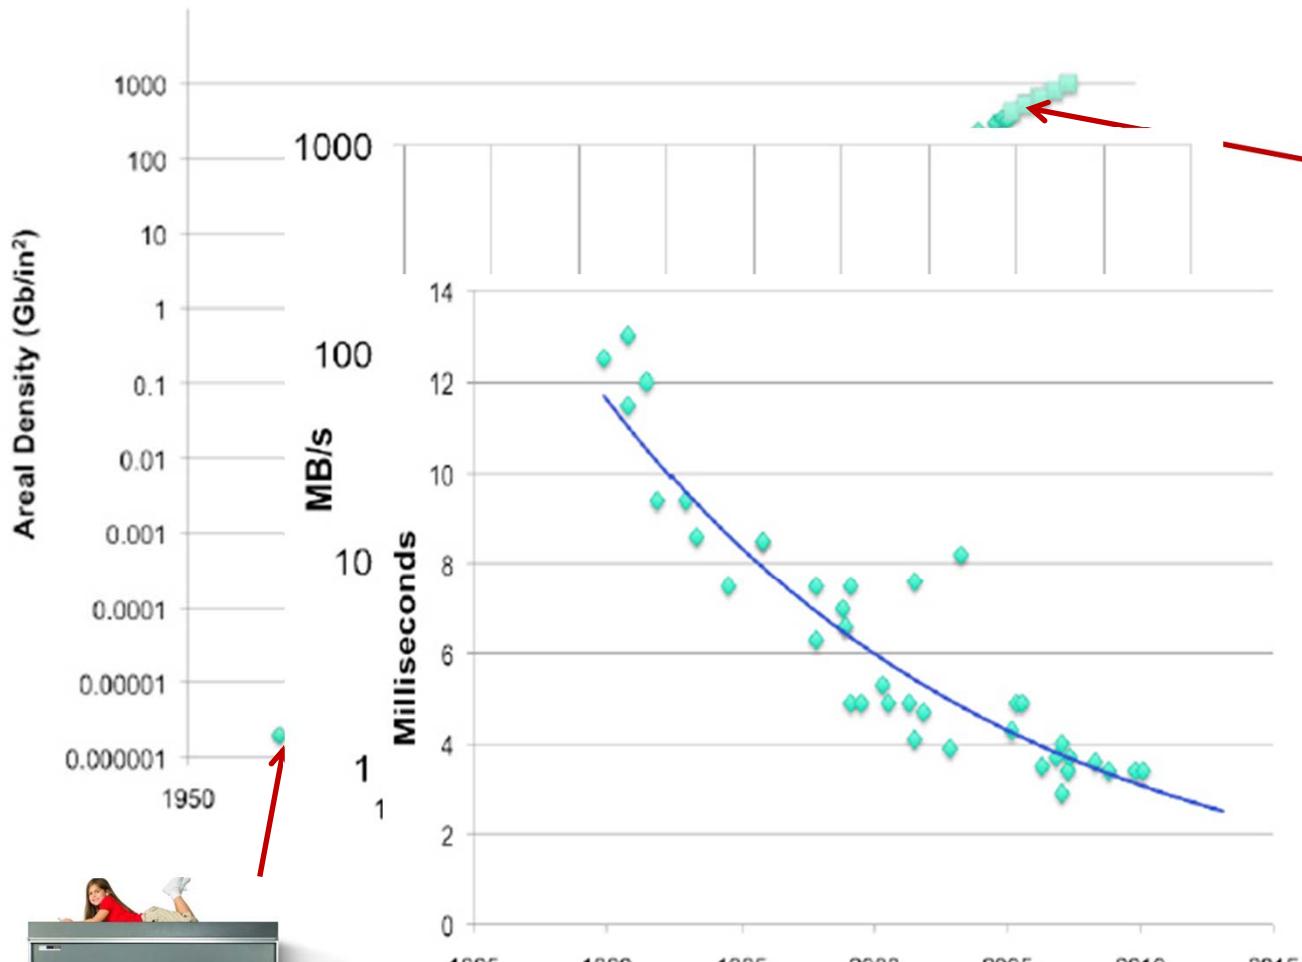
Toshiba introduces 4Mb NAND flash in 1992

1992

Samsung develops 16Mb NAND flash in 1994

1994

Hard drives



WD Se 4TB SATA drive (2013)

7,200 RPM
64MB buffer
Seek (avg.): several ms
4TB
0.75kg

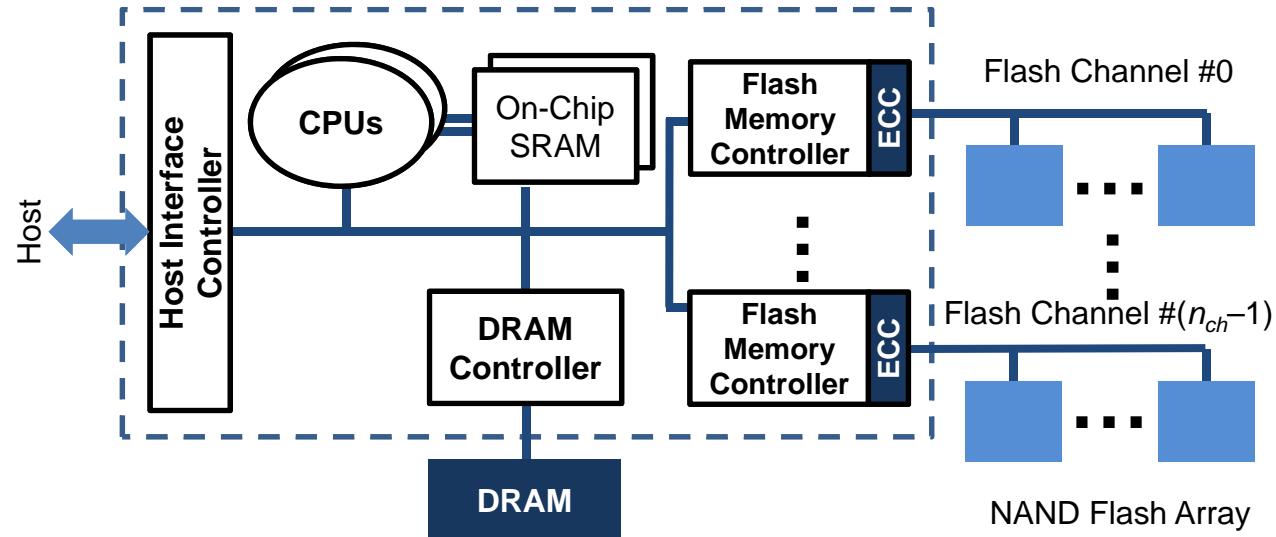


Memory 3.0

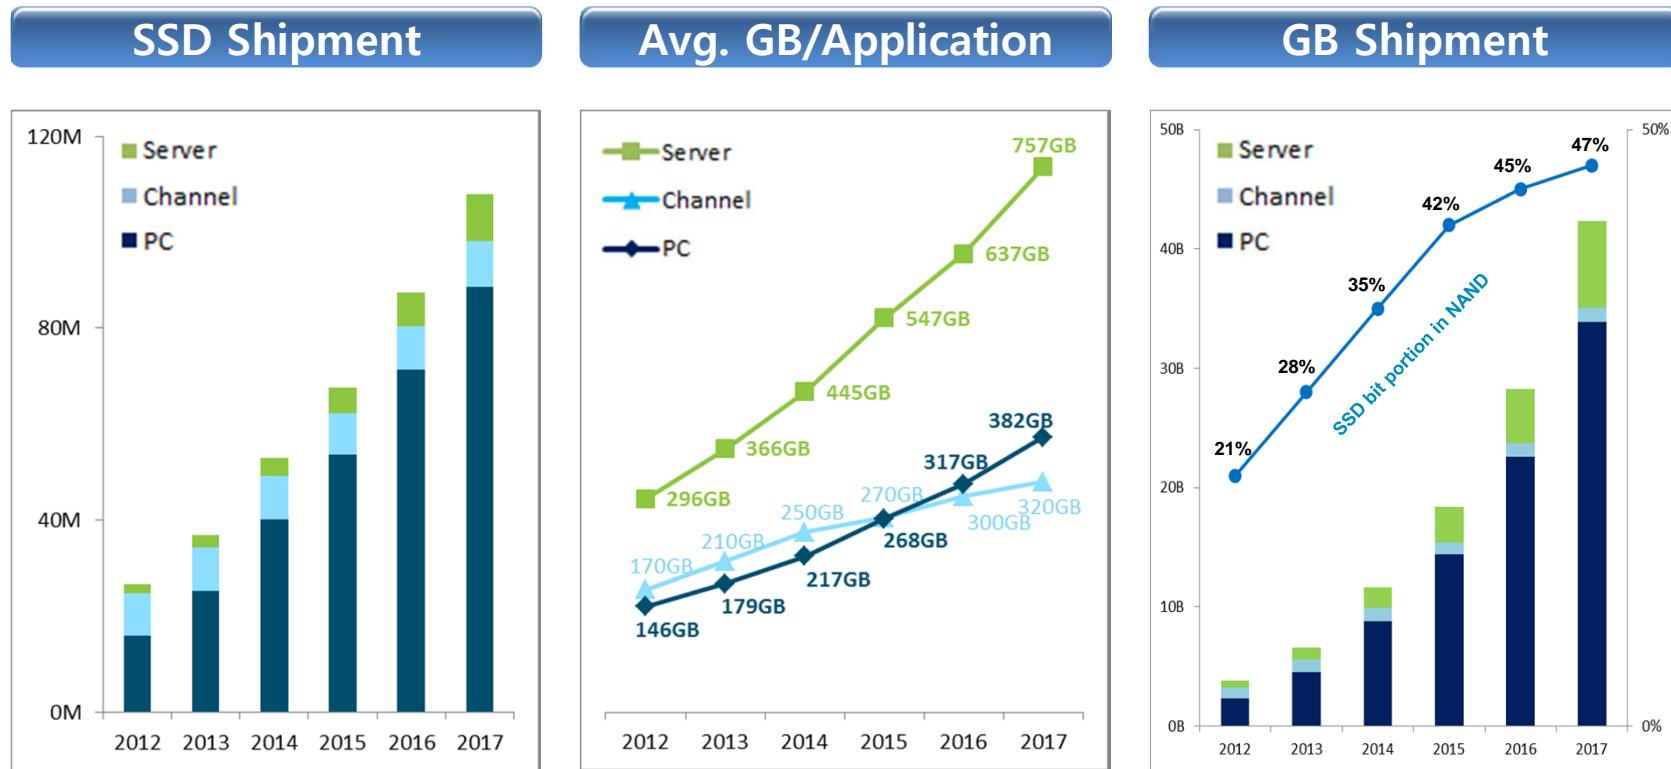
Hard drives, then and now

	RAMAC (1956)	WD Se (2013)	Ratio
Inch	60	2.5	1/24
Capacity	5MiB	4TiB	800k
Weight	>1 ton	0.75 kg	1/1,333
Rotation speed	1,200 rpm	7,200 rpm	6
Avg. seek	600ms	<5ms	1/120
Bits per gram	42	43B	>1B
Bandwidth	~9KiB/s	~100MiB/s	11.1k
Time to read out	9.25 min	667 min	72
Time to read out (4KiB random)	21 min	35 days	2,413

Solid-state drives



SSD market forecast



[Source: IDC May 2013]

Samsung: #1 SSD provider since 2007

Hard drive vs. SSD

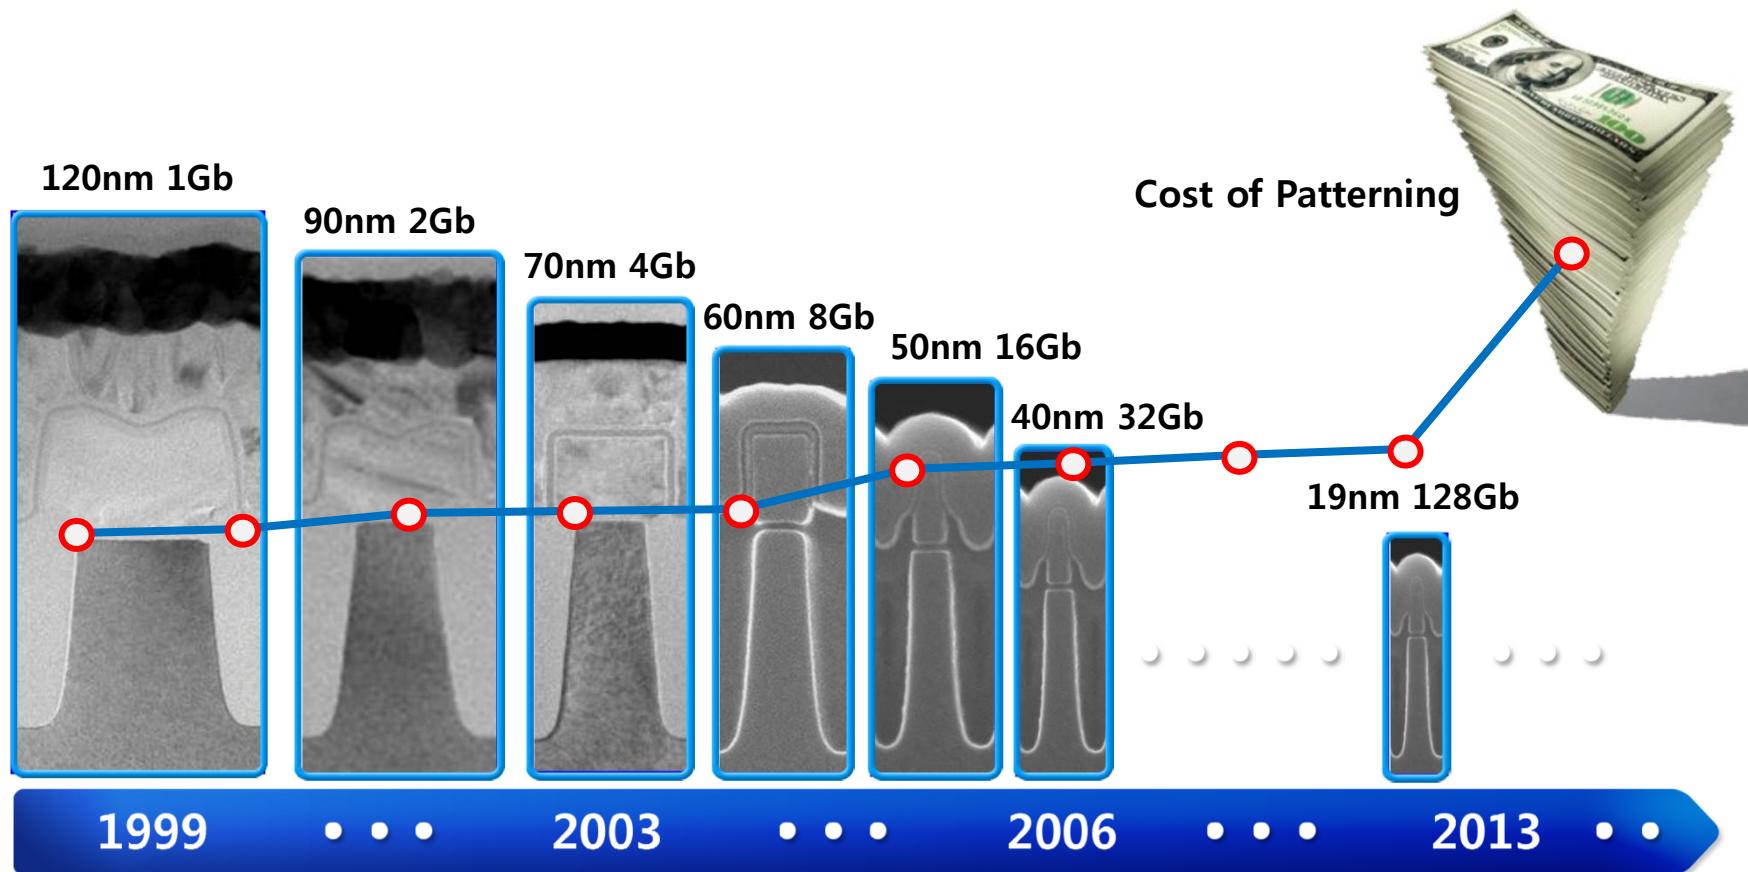
	WD Se 4TB	Samsung 841	Ratio
Inch	2.5	-	-
Capacity	4TiB	512GiB	1/8
Weight	0.75 kg	0.01 kG	1/75
Rotation speed	7,200 rpm	-	-
Avg. seek	<5ms	(negligible)	-
Bits per gram	43B	410B	9.5
Bandwidth	~100MiB/s	~540MiB/s	5.4
Time to read out	667 min	16 min	1/42
Time to read out (4KiB random)	35 days	22 min	1/2,291

Summary of memory 2.0

- Scaling rules!
 - DRAM has the crown in main memory (DDR_x)
 - Hard drive capacity follows exponential growth curve
 - But... the performance of hard drives is stagnant
 - NAND flash memory starts to replace (high-end) hard drives and enable mobile revolution!
 - Flash is new hard drive, hard drive is new tape
- However, ...
 - Further, economic (planar) scaling is seriously questioned
 - Physical limitations (e.g., cell interference) are becoming (seemingly) harder to overcome

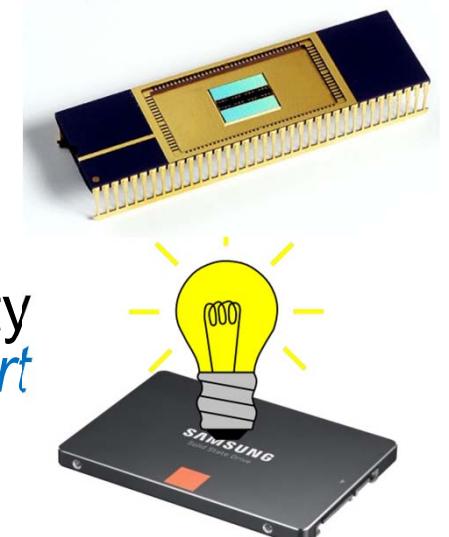
mem·o·ry 3.0

NAND flash scaling trend

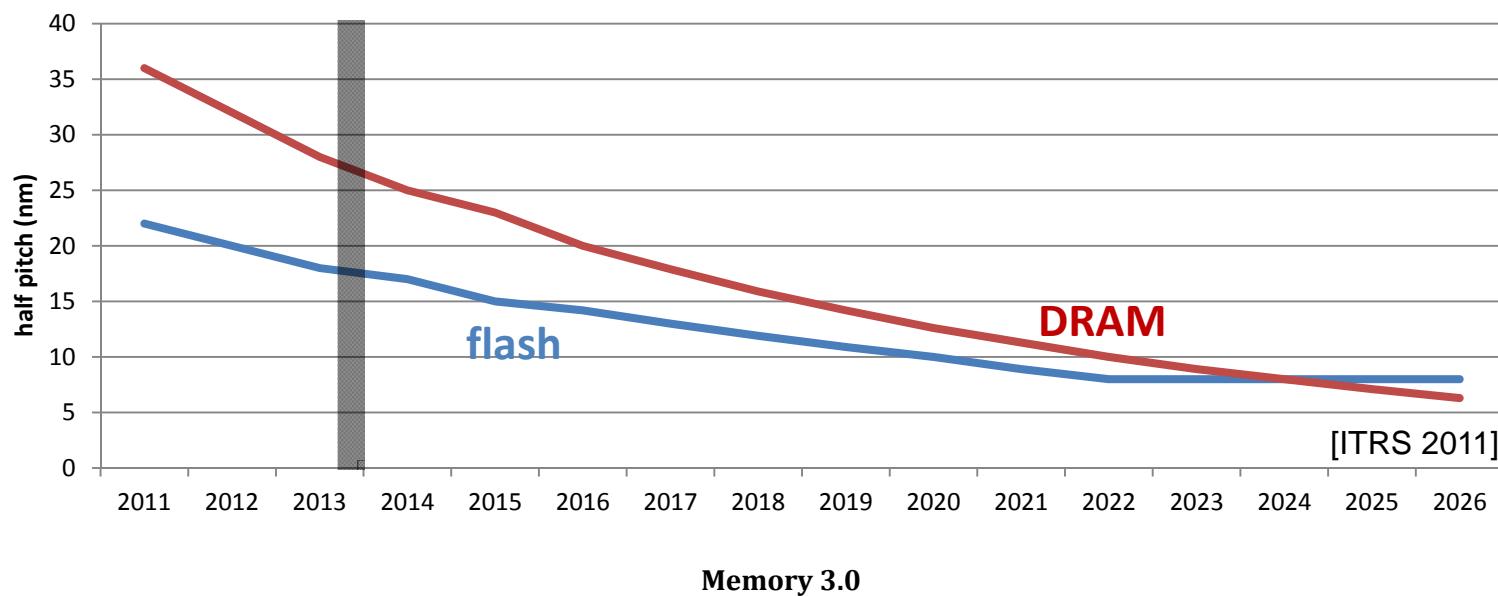


The era of mem·o·ry 3.0

- Economic planar scaling is *very* hard
 - *It's time to start planning for the end of Moore's Law*, August 2013, Bob Colwell (DARPA)
 - *The end of Moore's Law may ultimately be as much about economics as physics*
- We need creative approaches to scaling and adding value to memory solutions
- Consider potentially more scalable memory technologies, e.g., resistive memories
- As data-intensive applications and data locality become increasingly important, *active* or *smart* memory subsystems make more sense



1. Device and technology innovations will continue (but for how long?)

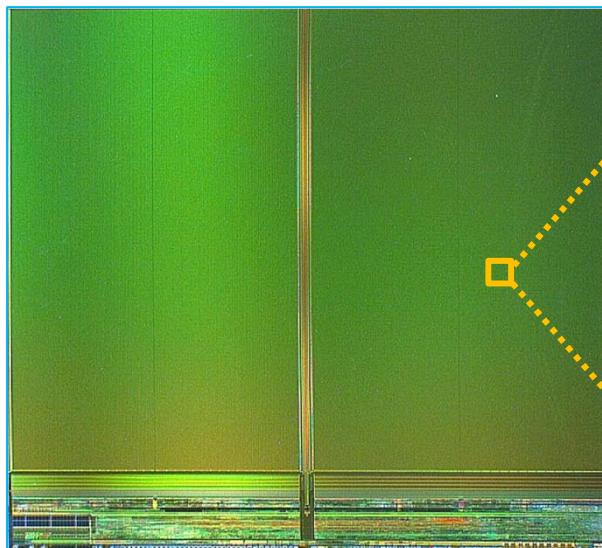


128Gb V-NAND [Elliott and Jung, Flash Memory Summit 2013]

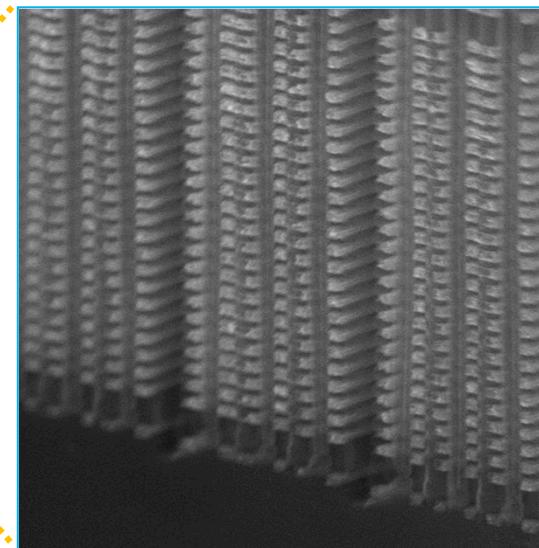
"The World's 1st 3D V-NAND Flash Mass Production"

Comparing with 20nm planar NAND Flash

- 2X Density and Write Speed
- $\frac{1}{2}$ Power Consumption
- 10X Endurance



128Gb V-NAND Flash

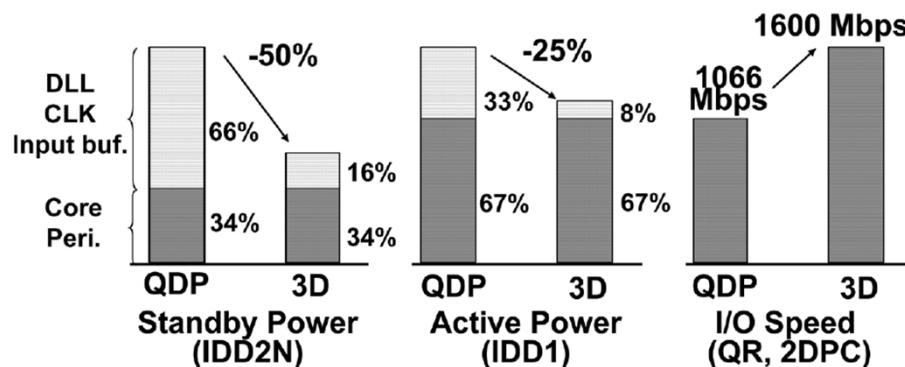
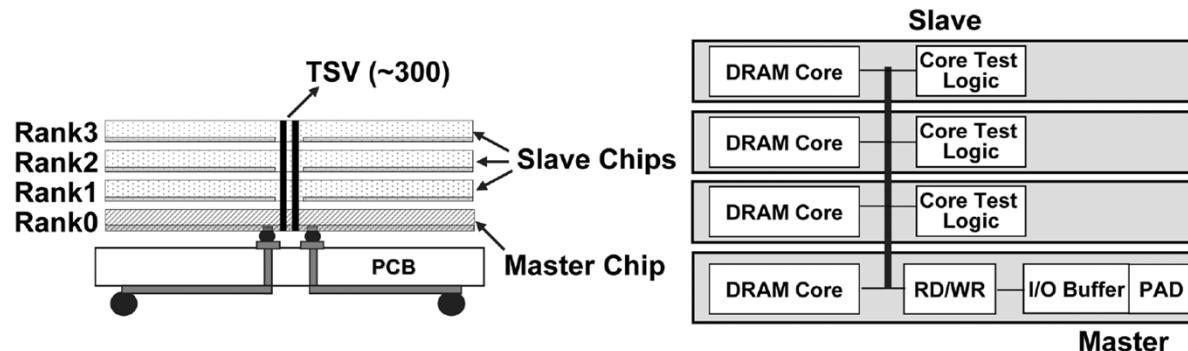


24 Layer Cell Structure

8 Gb 3-D DDR3 DRAM Using Through-Silicon-Via Technology

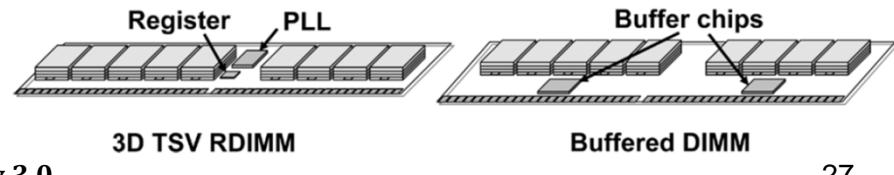
[JSSC 2010]

Uksong Kang, Hoe-Ju Chung, Seongmoo Heo, Duk-Ha Park, Hoon Lee, Jin Ho Kim, Soon-Hong Ahn, Soo-Ho Cha, Jaesung Ahn, DukMin Kwon, Jae-Wook Lee, Han-Sung Joo, Woo-Seop Kim, Member, IEEE, Dong Hyeon Jang, Nam Seog Kim, Jung-Hwan Choi, Tae-Gyeong Chung, Jei-Hwan Yoo, Joo Sun Choi, Changhyun Kim, Senior Member, IEEE, and Young-Hyun Jun



	3D TSV module	Buffered DIMM
*Standby-power (IDD2N)	48%	100%
*Active power (IDD1)	65%	100%
Latency @1333Mbps	1 cycle	3~4 cycles
Cost adder	TSV process cost	2 buffer cost
Density	Max 4 rank	
Speed@2DPC	~ 1600Mbps (QR X4)	

*3D module PLL+register: ~2W, Buffered-DIMM 2 buffer chips: ~4.3W

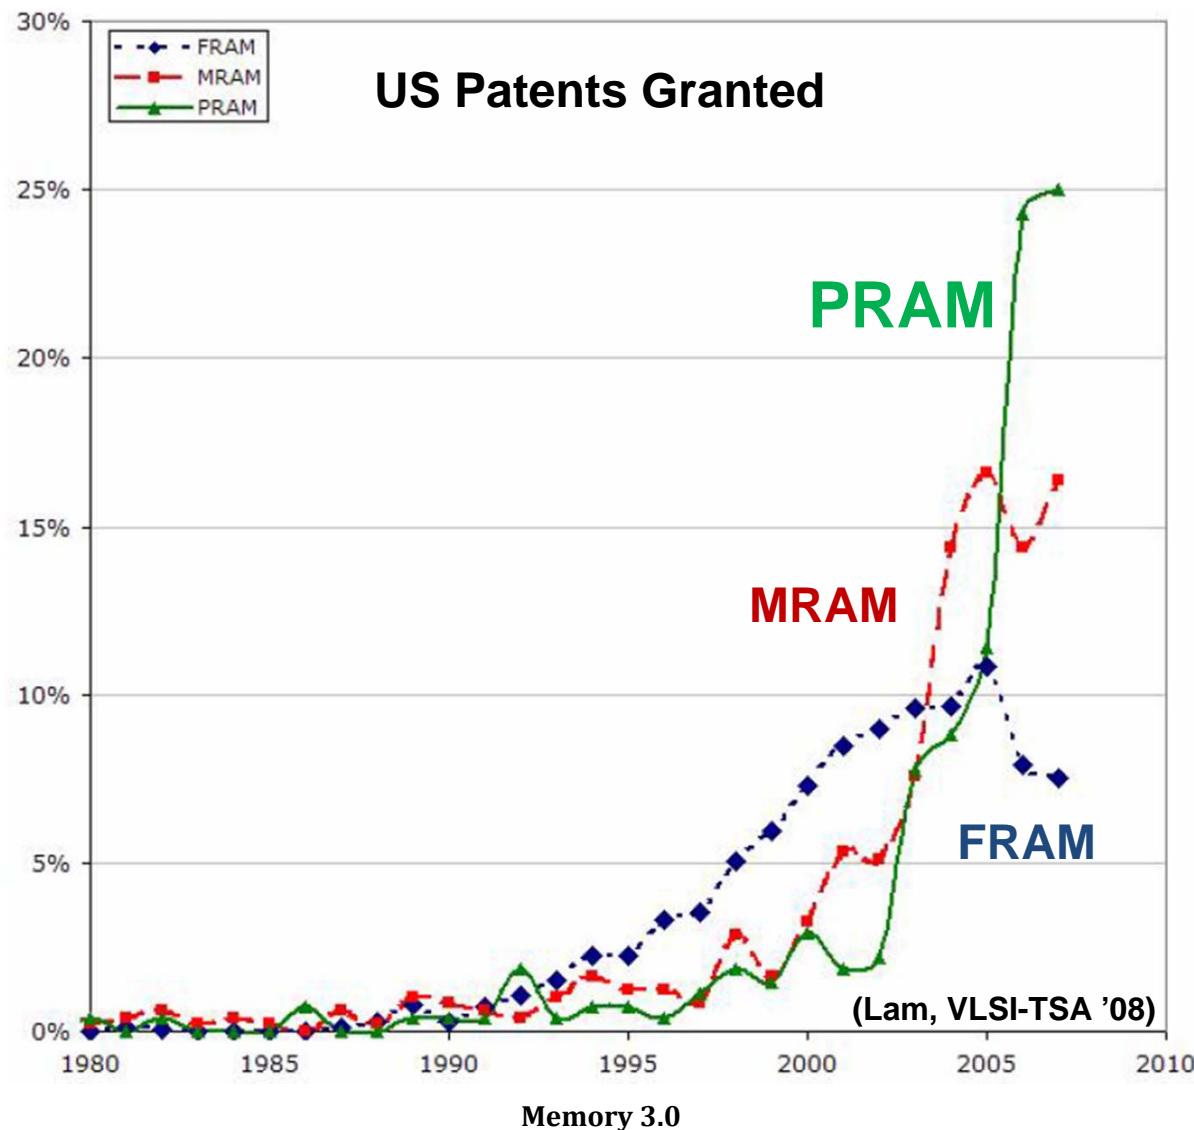


2. New memories are coming?

	Latency			Program energy	Allowable access unit	Retention on power-off	Write endurance	Cell density*
	read	write	erase					
PCM	20ns	100ns	N/A	100 pJ	byte	Yes	$10^8 \sim 10^9$	$5F^2$
STT-MRAM	10ns	10ns	N/A	0.02 pJ	byte	Yes	10^{15}	$4F^2$
ReRAM	10ns	20ns	N/A	2 pJ	byte	Yes	10^6	$6F^2$
DRAM	10ns	10ns	N/A	2 pJ	byte	No	10^{16}	$(2/3)F^2$
NAND flash	$25\mu s$	$200\mu s$	1.5ms	10 nJ	page/block	Yes	$10^4 \sim 10^6$	$4 \sim 5F^2$
HDD	8.5ms	9.5ms	N/A	N/A	sector	Yes	N/A	$2 \sim 3F^2$

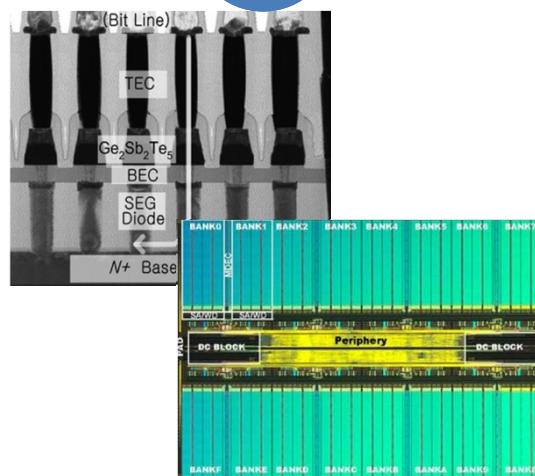
[Cryder and Kim, Trans. Magnetics 2009]

Interest grows

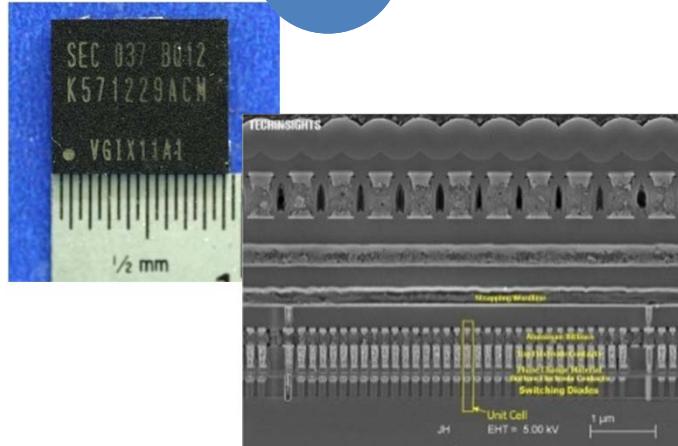


Samsung

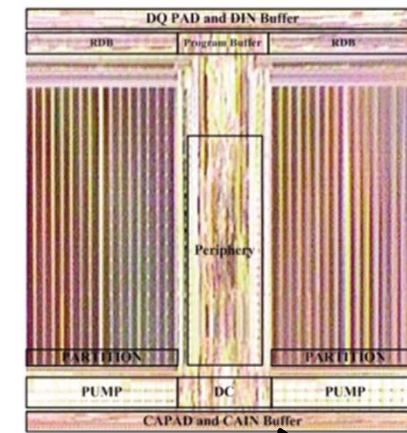
Lee et al. ISSCC '07
Lee et al. JSSC '08



Techinsights decap '10

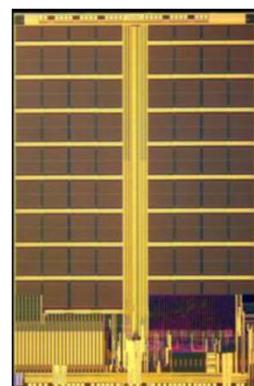


Chung et al. ISSCC '11



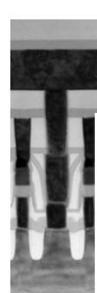
Numonyx (now Micron)

Early access program
(2009)



“Alverstone” (OMNEO)
128Mb @90nm
TR switch design
40MB/s read (?)
<1MB/s write (?)

Numerous press releases
(slated for MP in 2011)



Device/Tecnology	Write Bandwidth (Mb/s)
128Mb Axcell™ P33 130nm NOR Flash	1.3
128Mb P8P 90nm PCM	2.8
1Gb 45nm PCM	27.3



(Servalli, IEDM '09)

(2011~2012?)

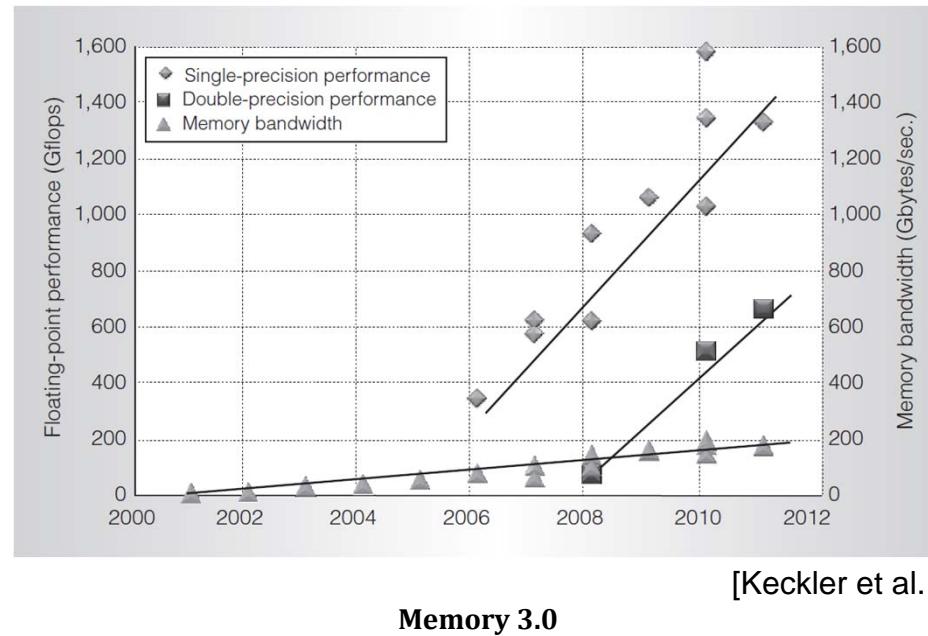
(www.micron.com)

“Bonelli”
1Gb @45nm
1.8V I/O

Memory 3.0

“Imola” and “Mandello”
2Gb & 4Gb @45nm
1.2V & 1.8V I/O
LPDDR2-NVM &
DDR3-NVM

3. Closer and faster, please!



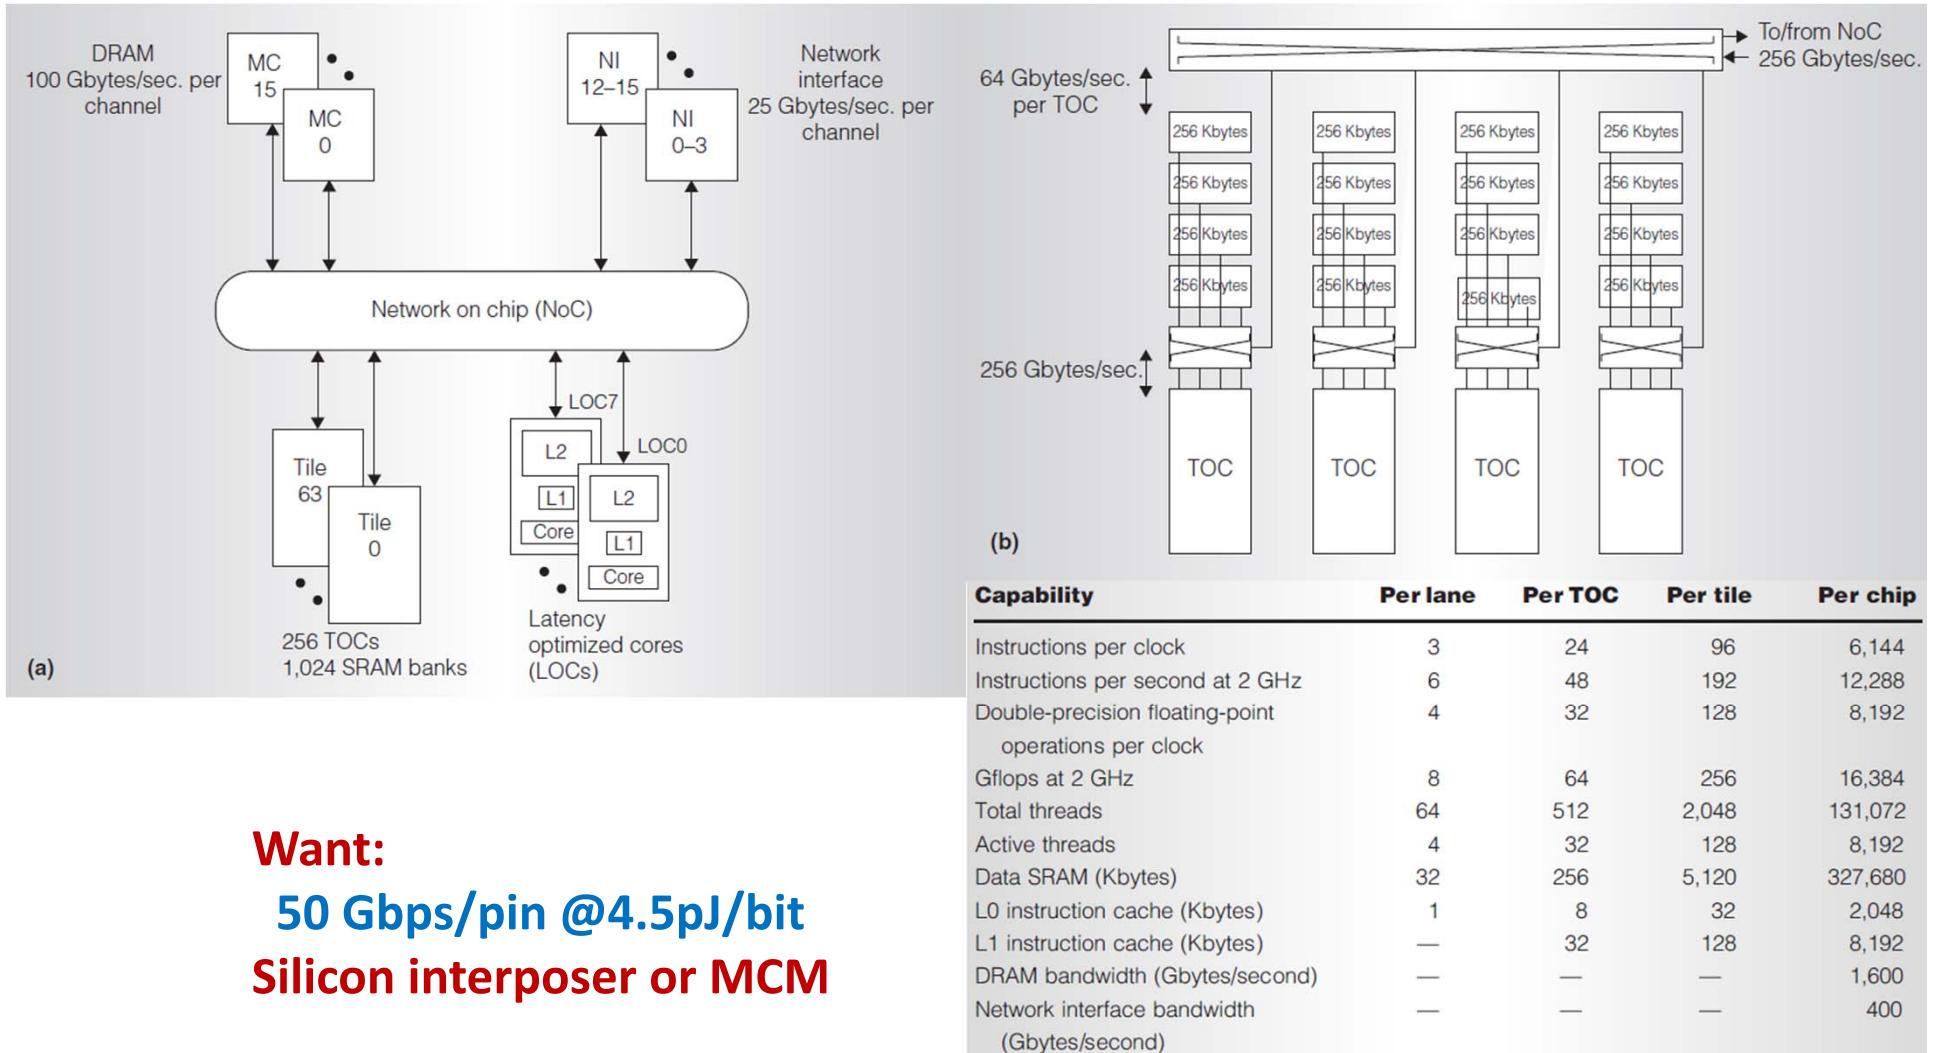
Distance of data sorely felt

[Keckler et al., IEEE Micro 2009]

Process technology	2010	2017	
	40nm	10nm, high freq.	10nm, low volt.
V_{DD} (nominal)	0.9 V	0.75 V	0.65 V
Frequency target	1.6 GHz	2.5 GHz	2 GHz
Double-precision FMA energy	50 pJ	8.7 pJ	6.5 pJ
64-bit read from an 8KiB SRAM	14 pJ	2.4 pJ	1.8 pJ
Wire energy (256 bits, 10mm)	310 pJ	200 pJ	150 pJ
Operand fetch from DRAM	More than 10nJ		

Exascale goal: 20 pJ per floating point operation

nVIDIA Echelon



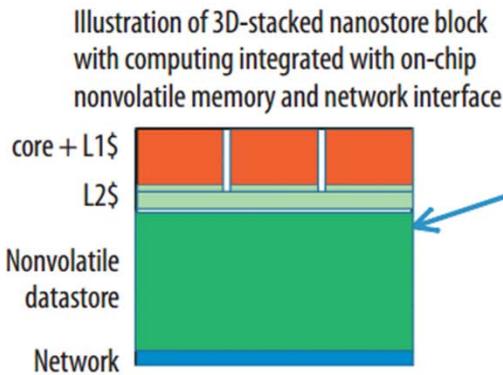
Want:

50 Gbps/pin @4.5pJ/bit

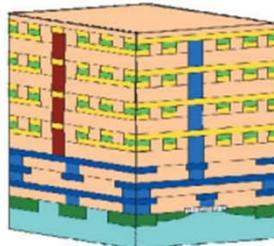
Silicon interposer or MCM

[Keckler et al., IEEE Micro 2009]

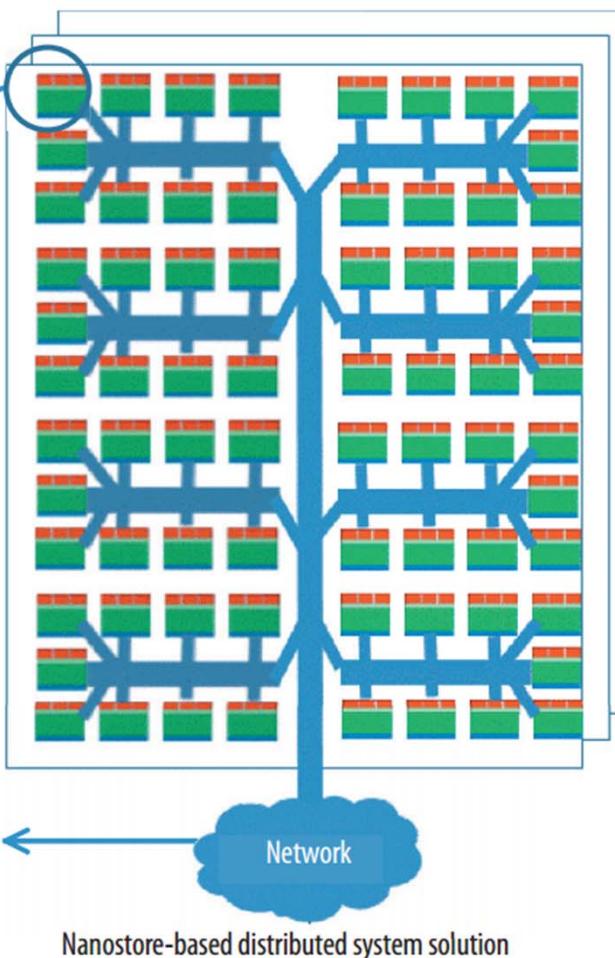
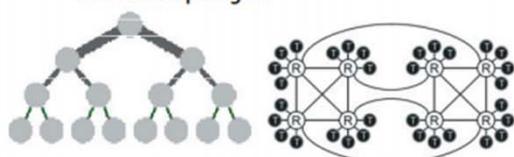
HP Lab Nanostore



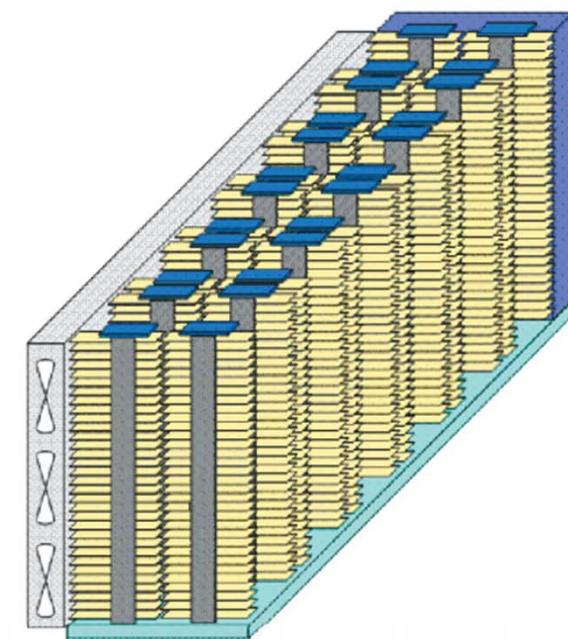
Example 3D-stacked memristor die showing the CMOS layer in the bottom and the cross-bar and wiring layers on top



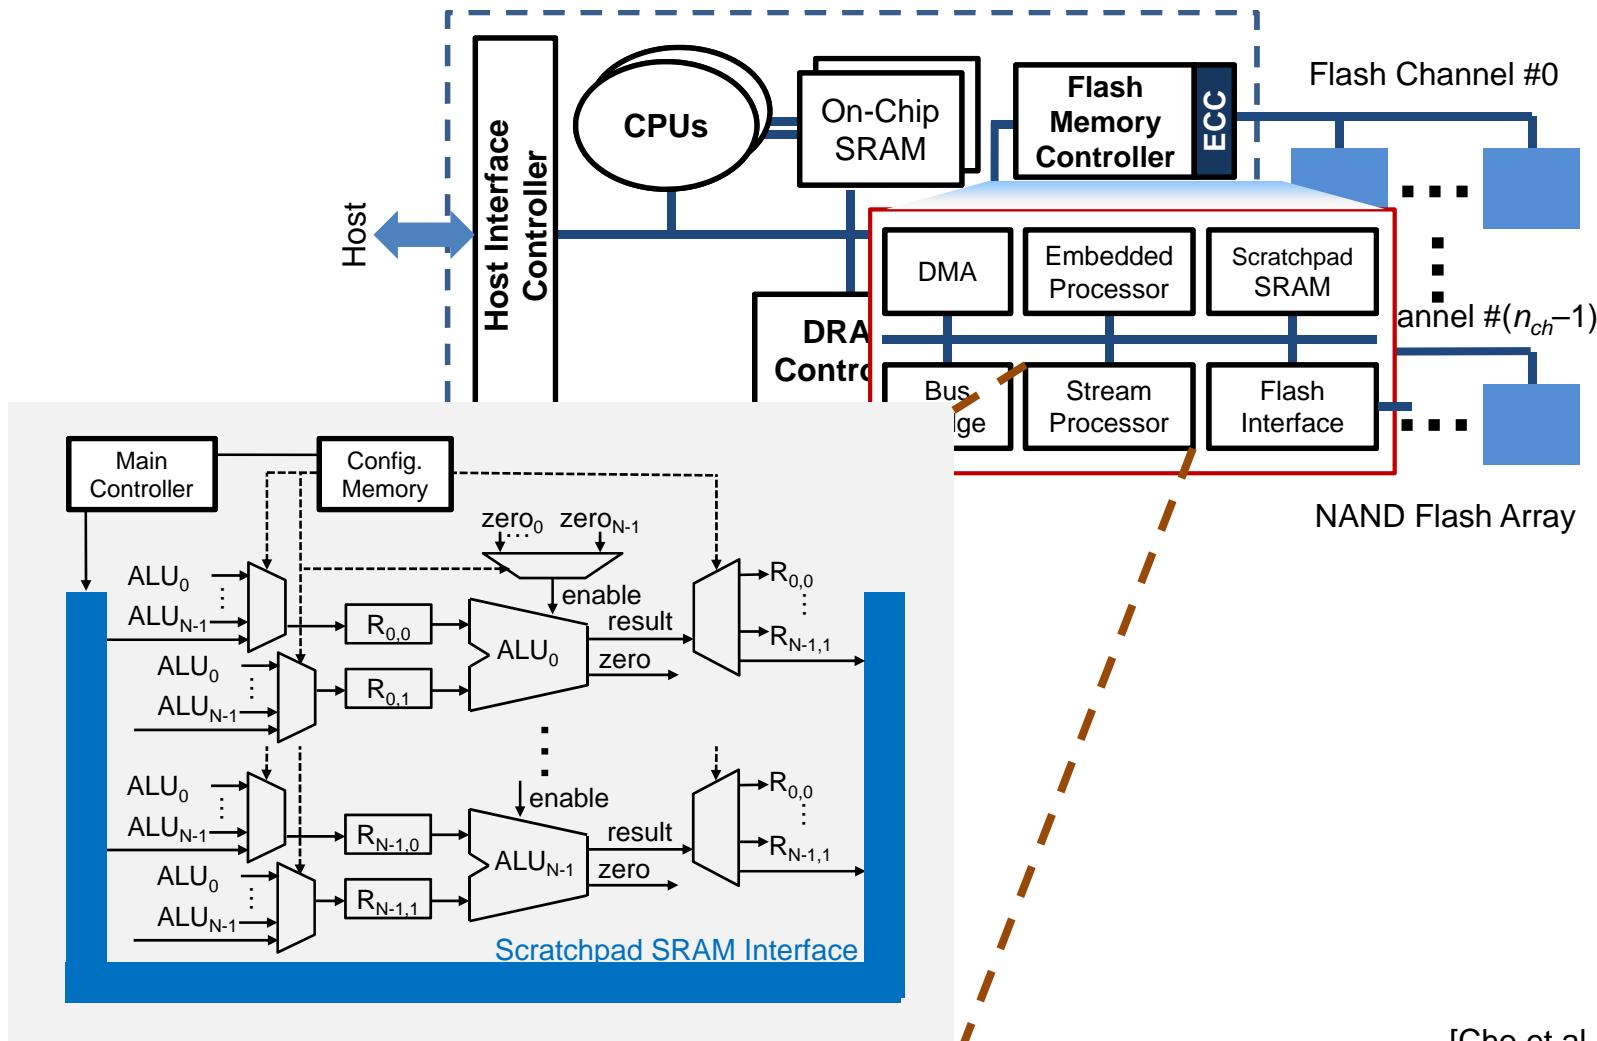
Example fat-tree and HyperX network topologies



[Ranganathan, IEEE Computer 2011]

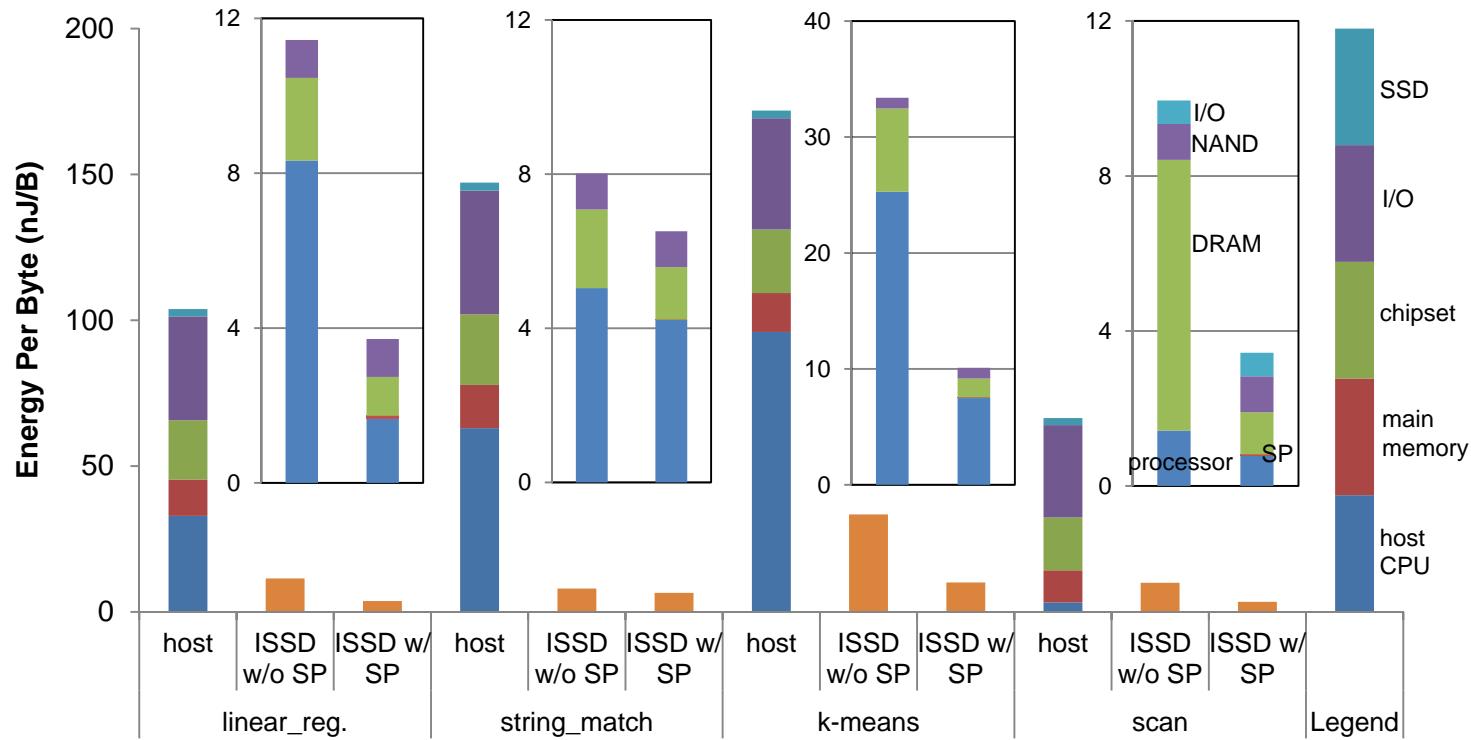


“Intelligent” SSD (iSSD)



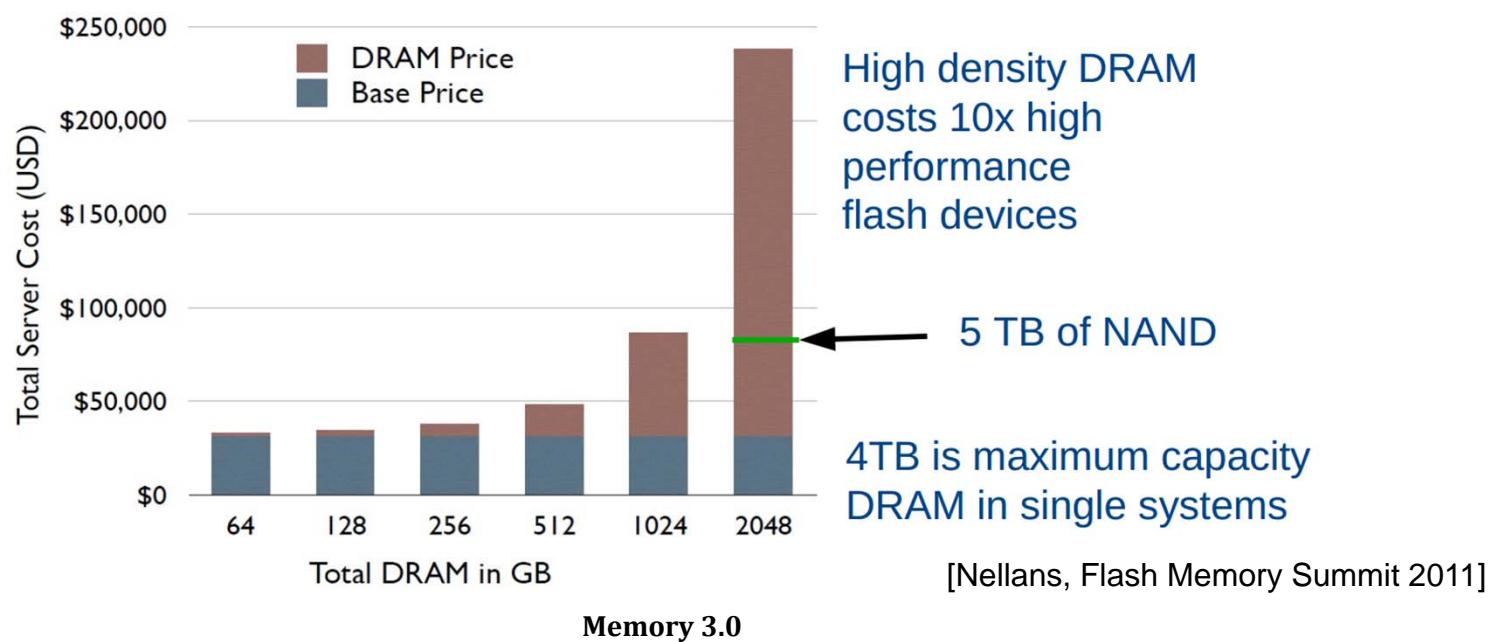
[Cho et al., ICS 2013]

Energy (energy per byte)

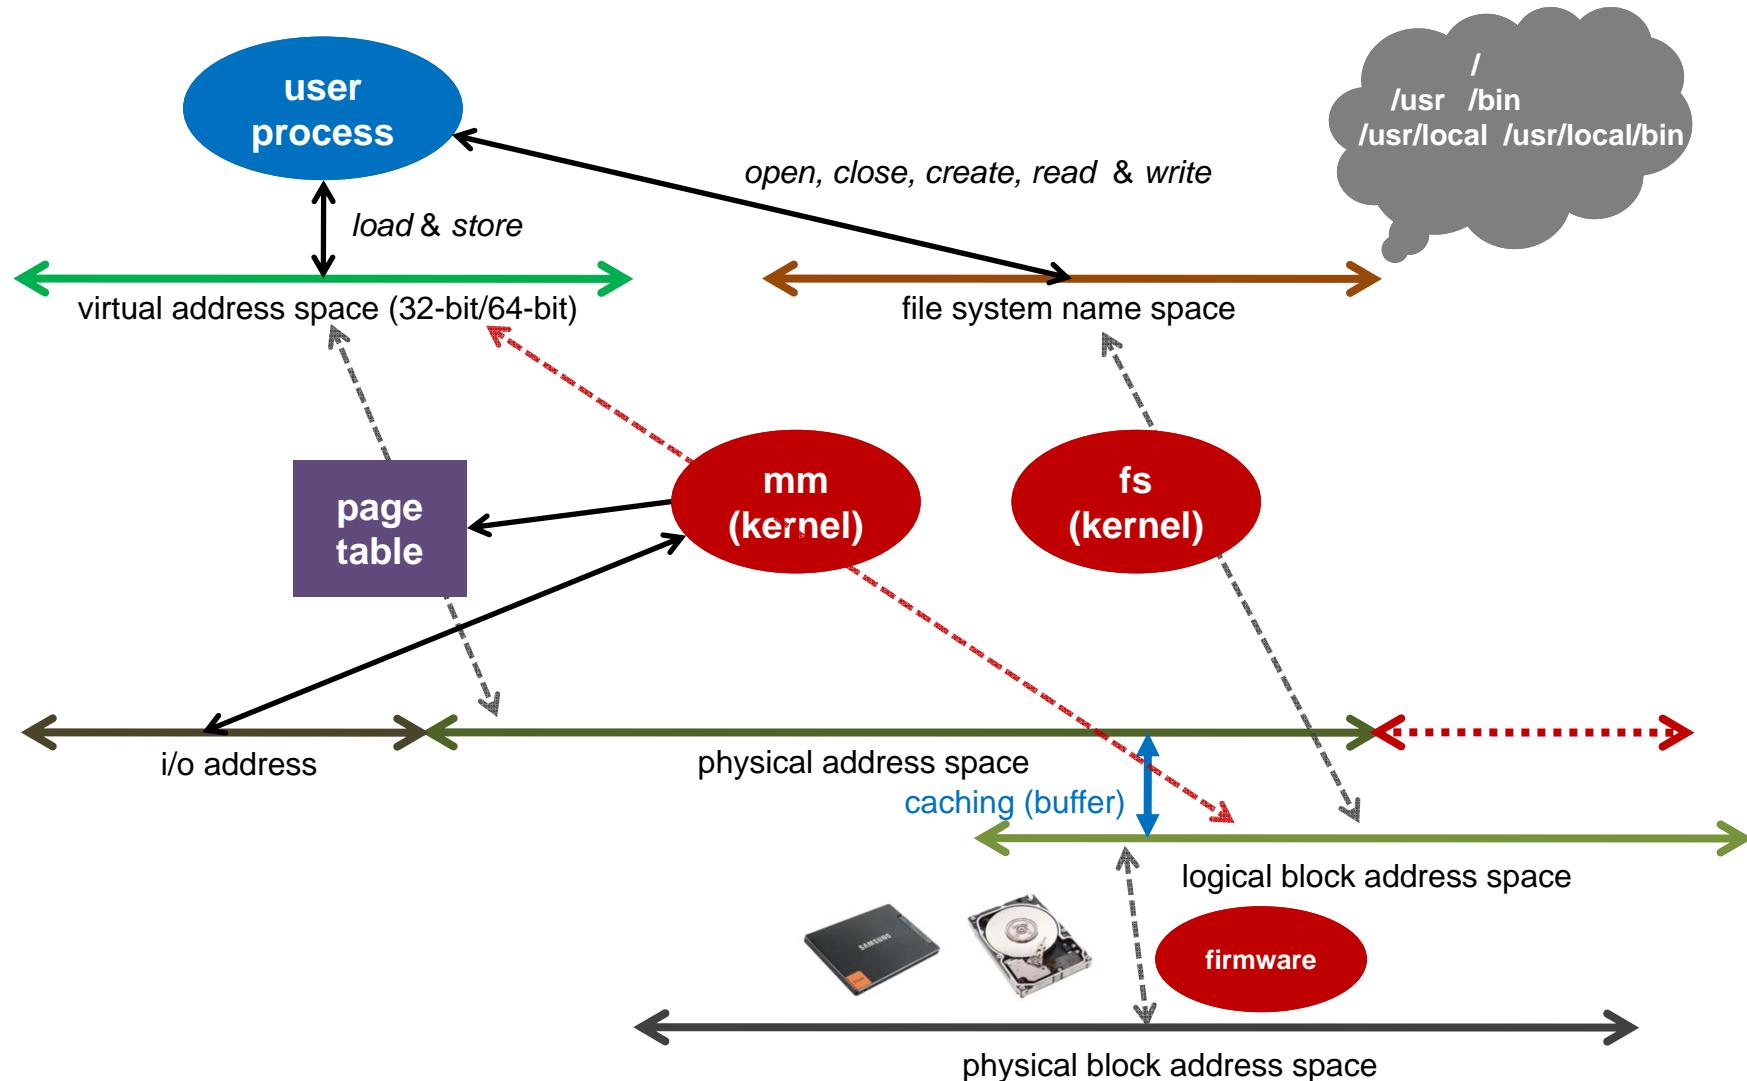


- iSSD energy benefits are large!
 - At least 5× (k-means) and the average is 9+×

4. Cooler and larger, please!

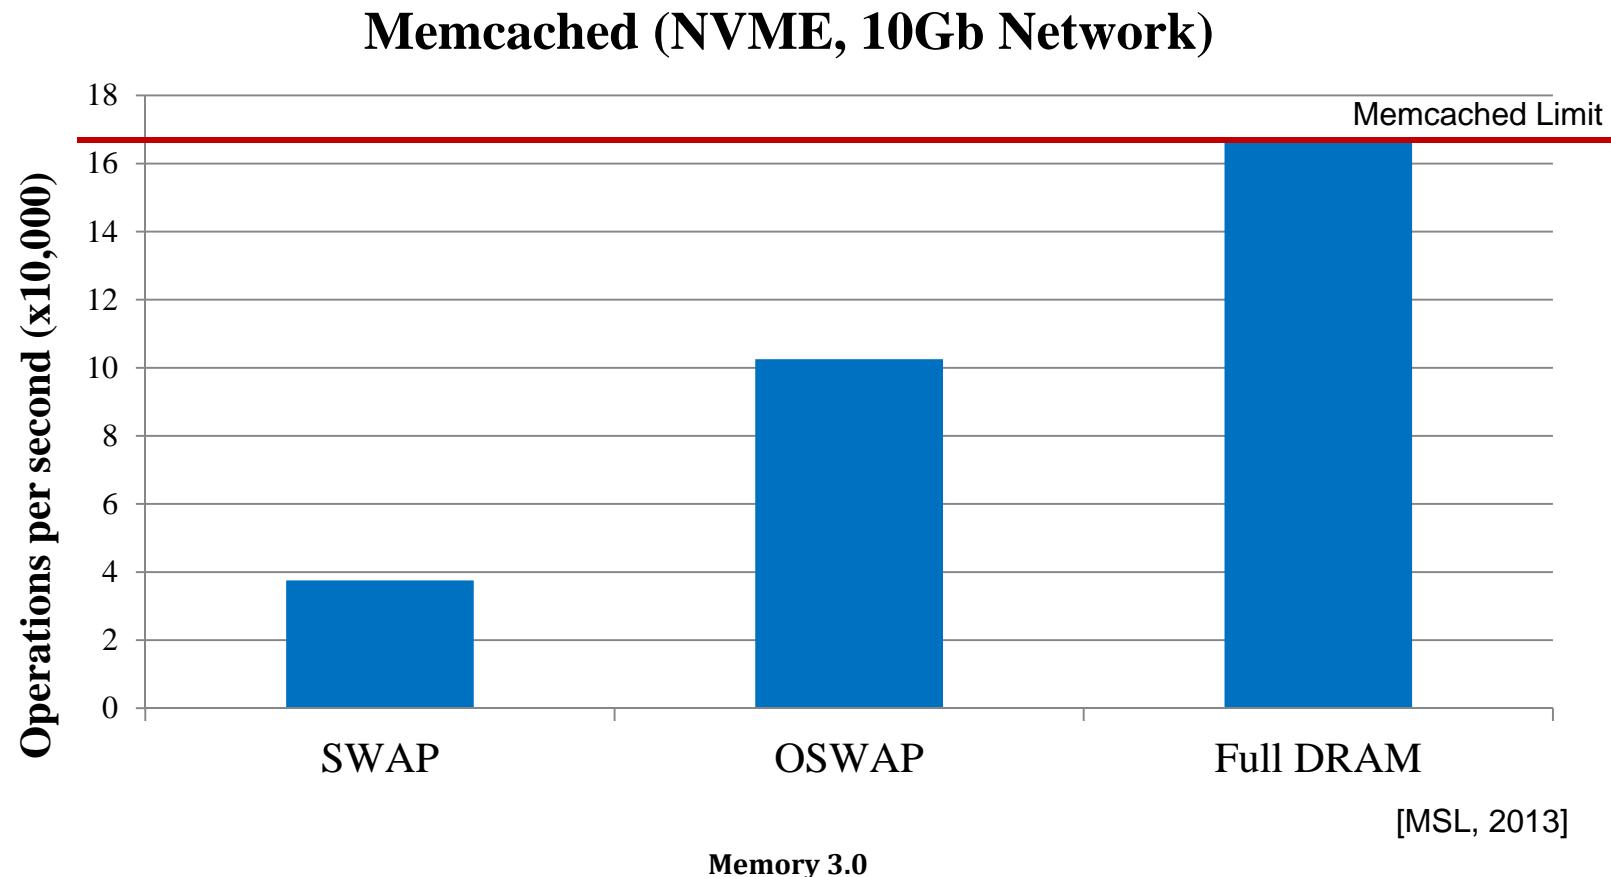


Memory space hierarchy



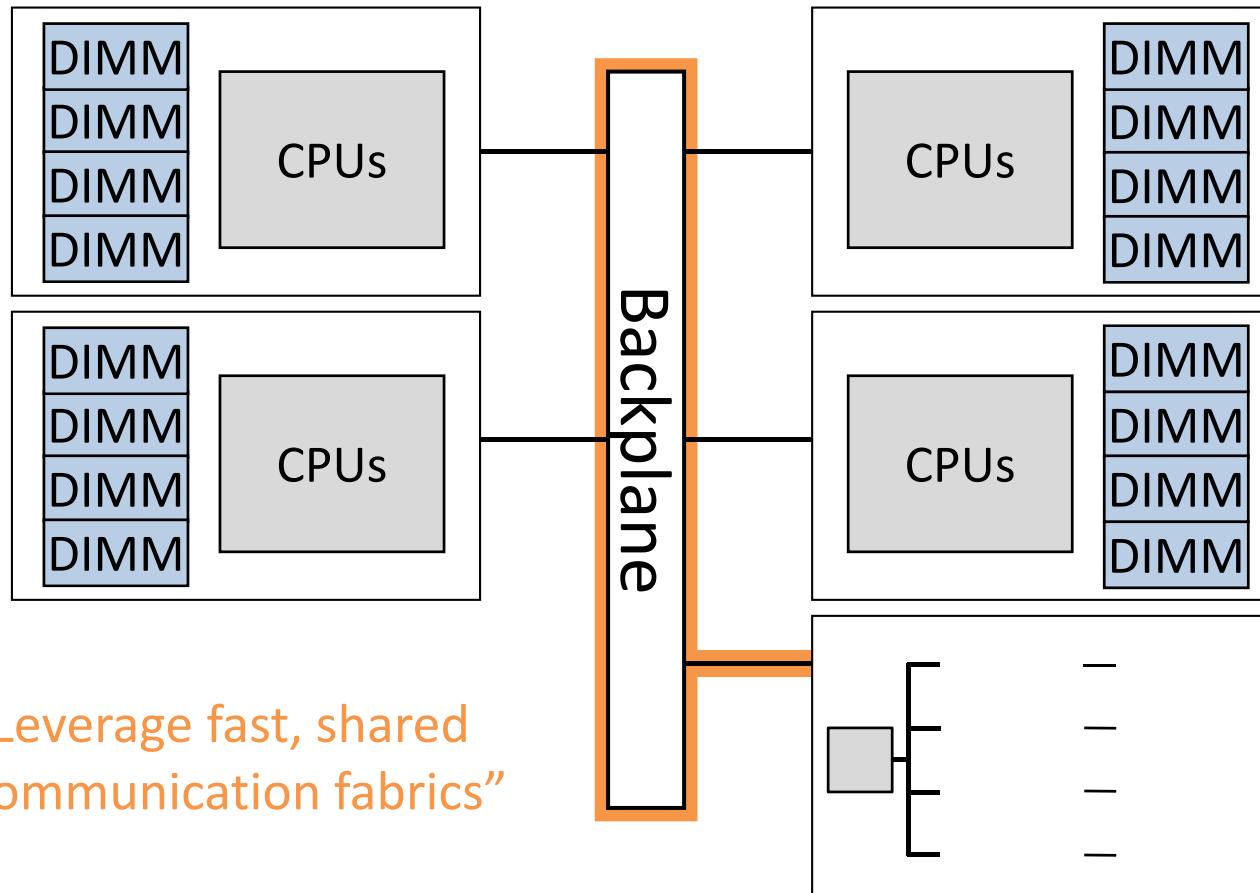
O-SWAP

- **Idea:** Provide a transparent DRAM-flash data exchange path
- **Target metric:** $(\text{QoS throughput} \times \text{data capacity})/\$$



“Memory blade” [Lim et al., ISCA 2009]

Blade system with disaggregated memory



BIG DATA EXPLOSION

2013 will see data become more important than ever before. Take a look at how much data is being generated every day from around the globe and how it could effect the way your business works in the years to come.

FROM THE DAWN OF CIVILISATION THROUGH TO 2003



GOOGLE CALCULATES HUMANS PRODUCED 5 EXABYTES OF DATA

WE NOW PRODUCE 5 EXABYTES OF DATA EVERY 2 DAYS

WHAT IS AN EXABYTE? AN EXABYTE IS 1 BILLION, BILLIONS



IF PRINTED IT WOULD FILL 20 BILLION
4 DRAW FILING CABINETS

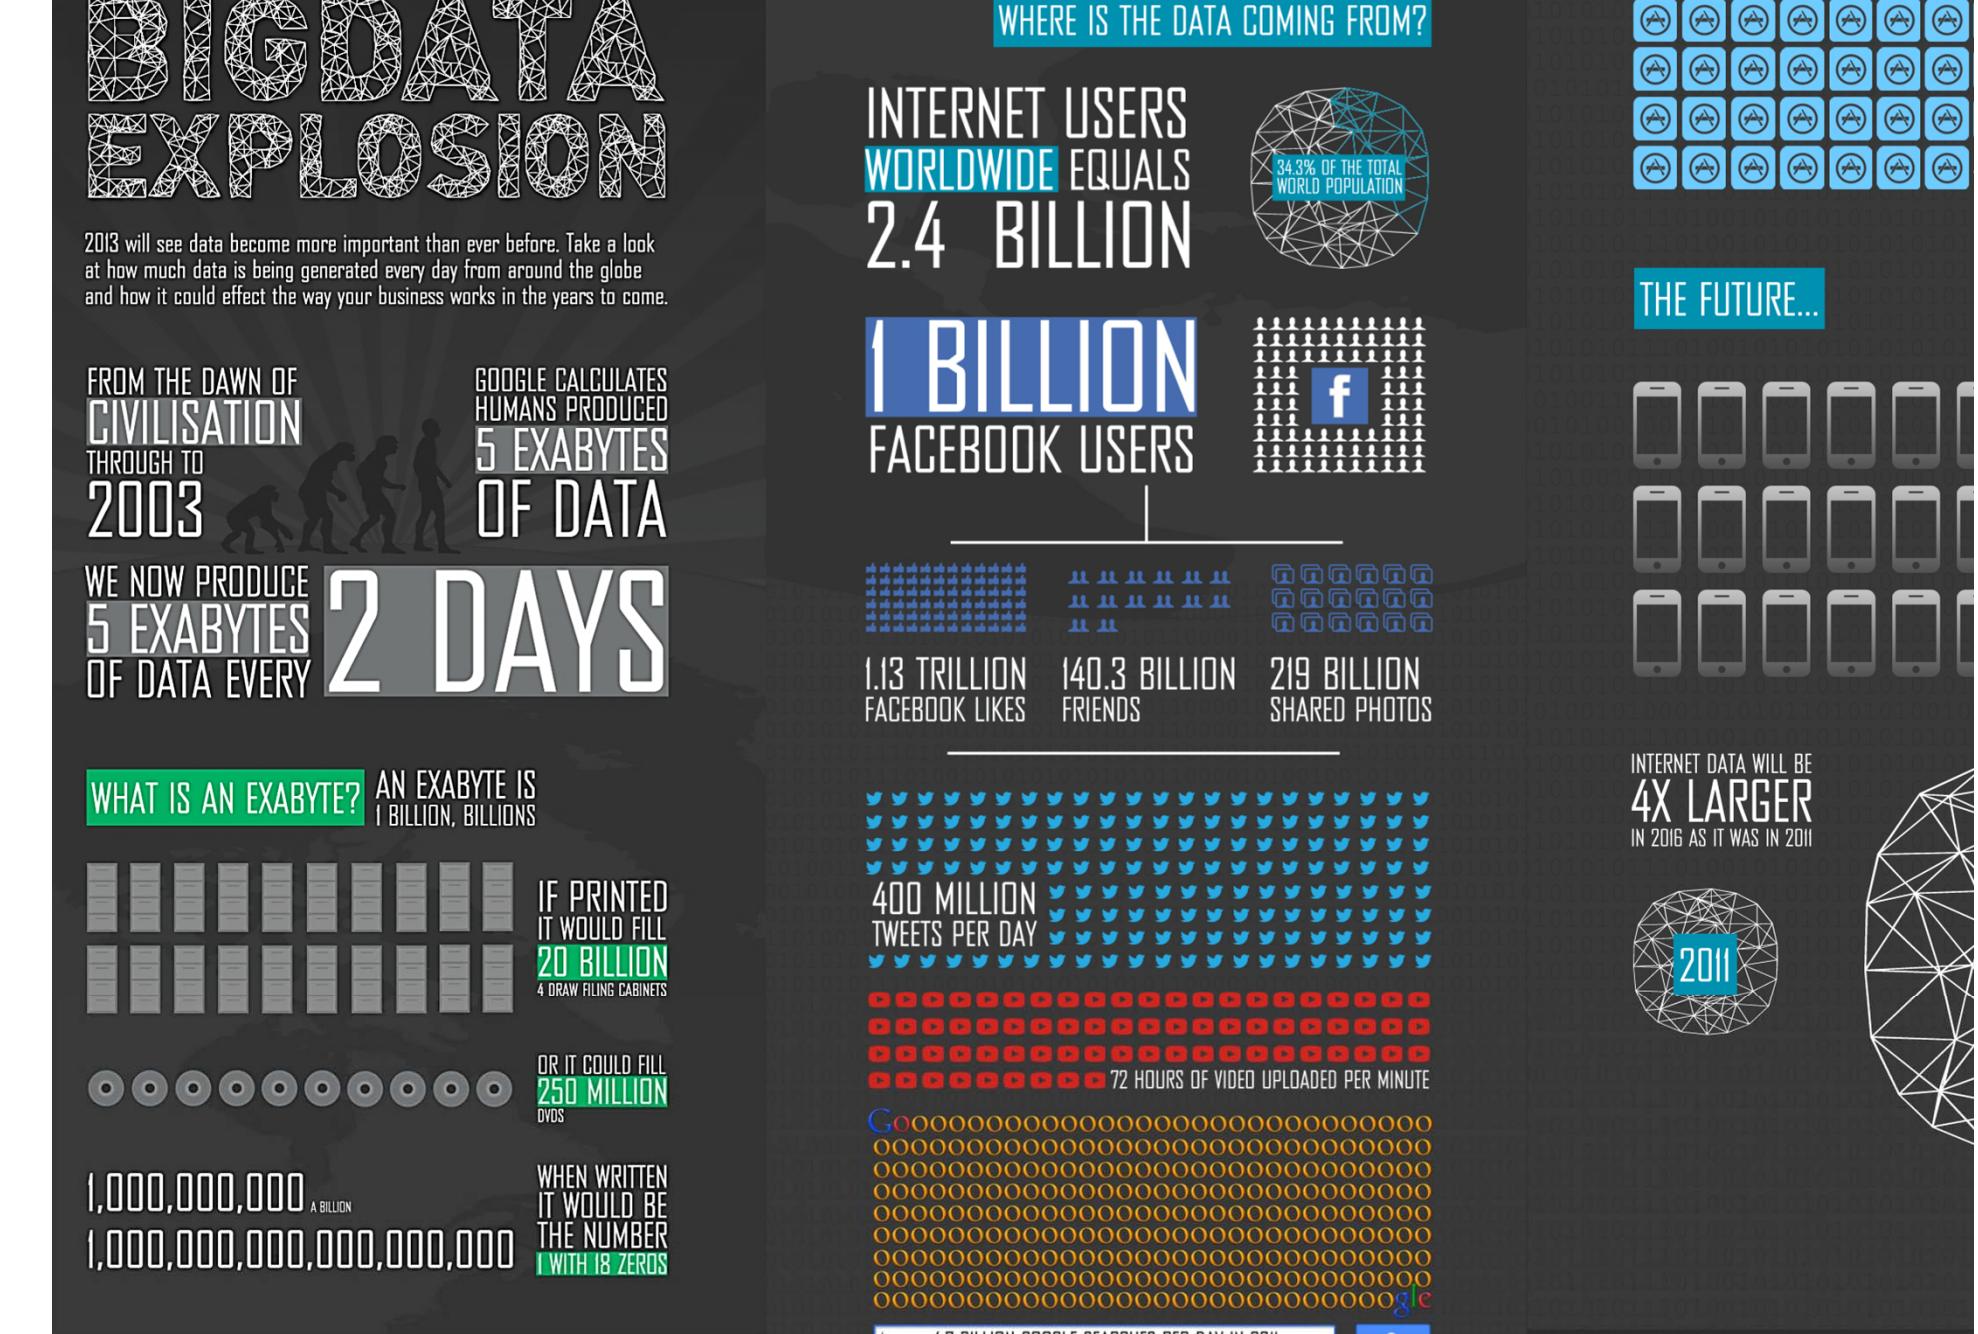


OR IT COULD FILL 250 MILLION DVDS

1,000,000,000 A BILLION

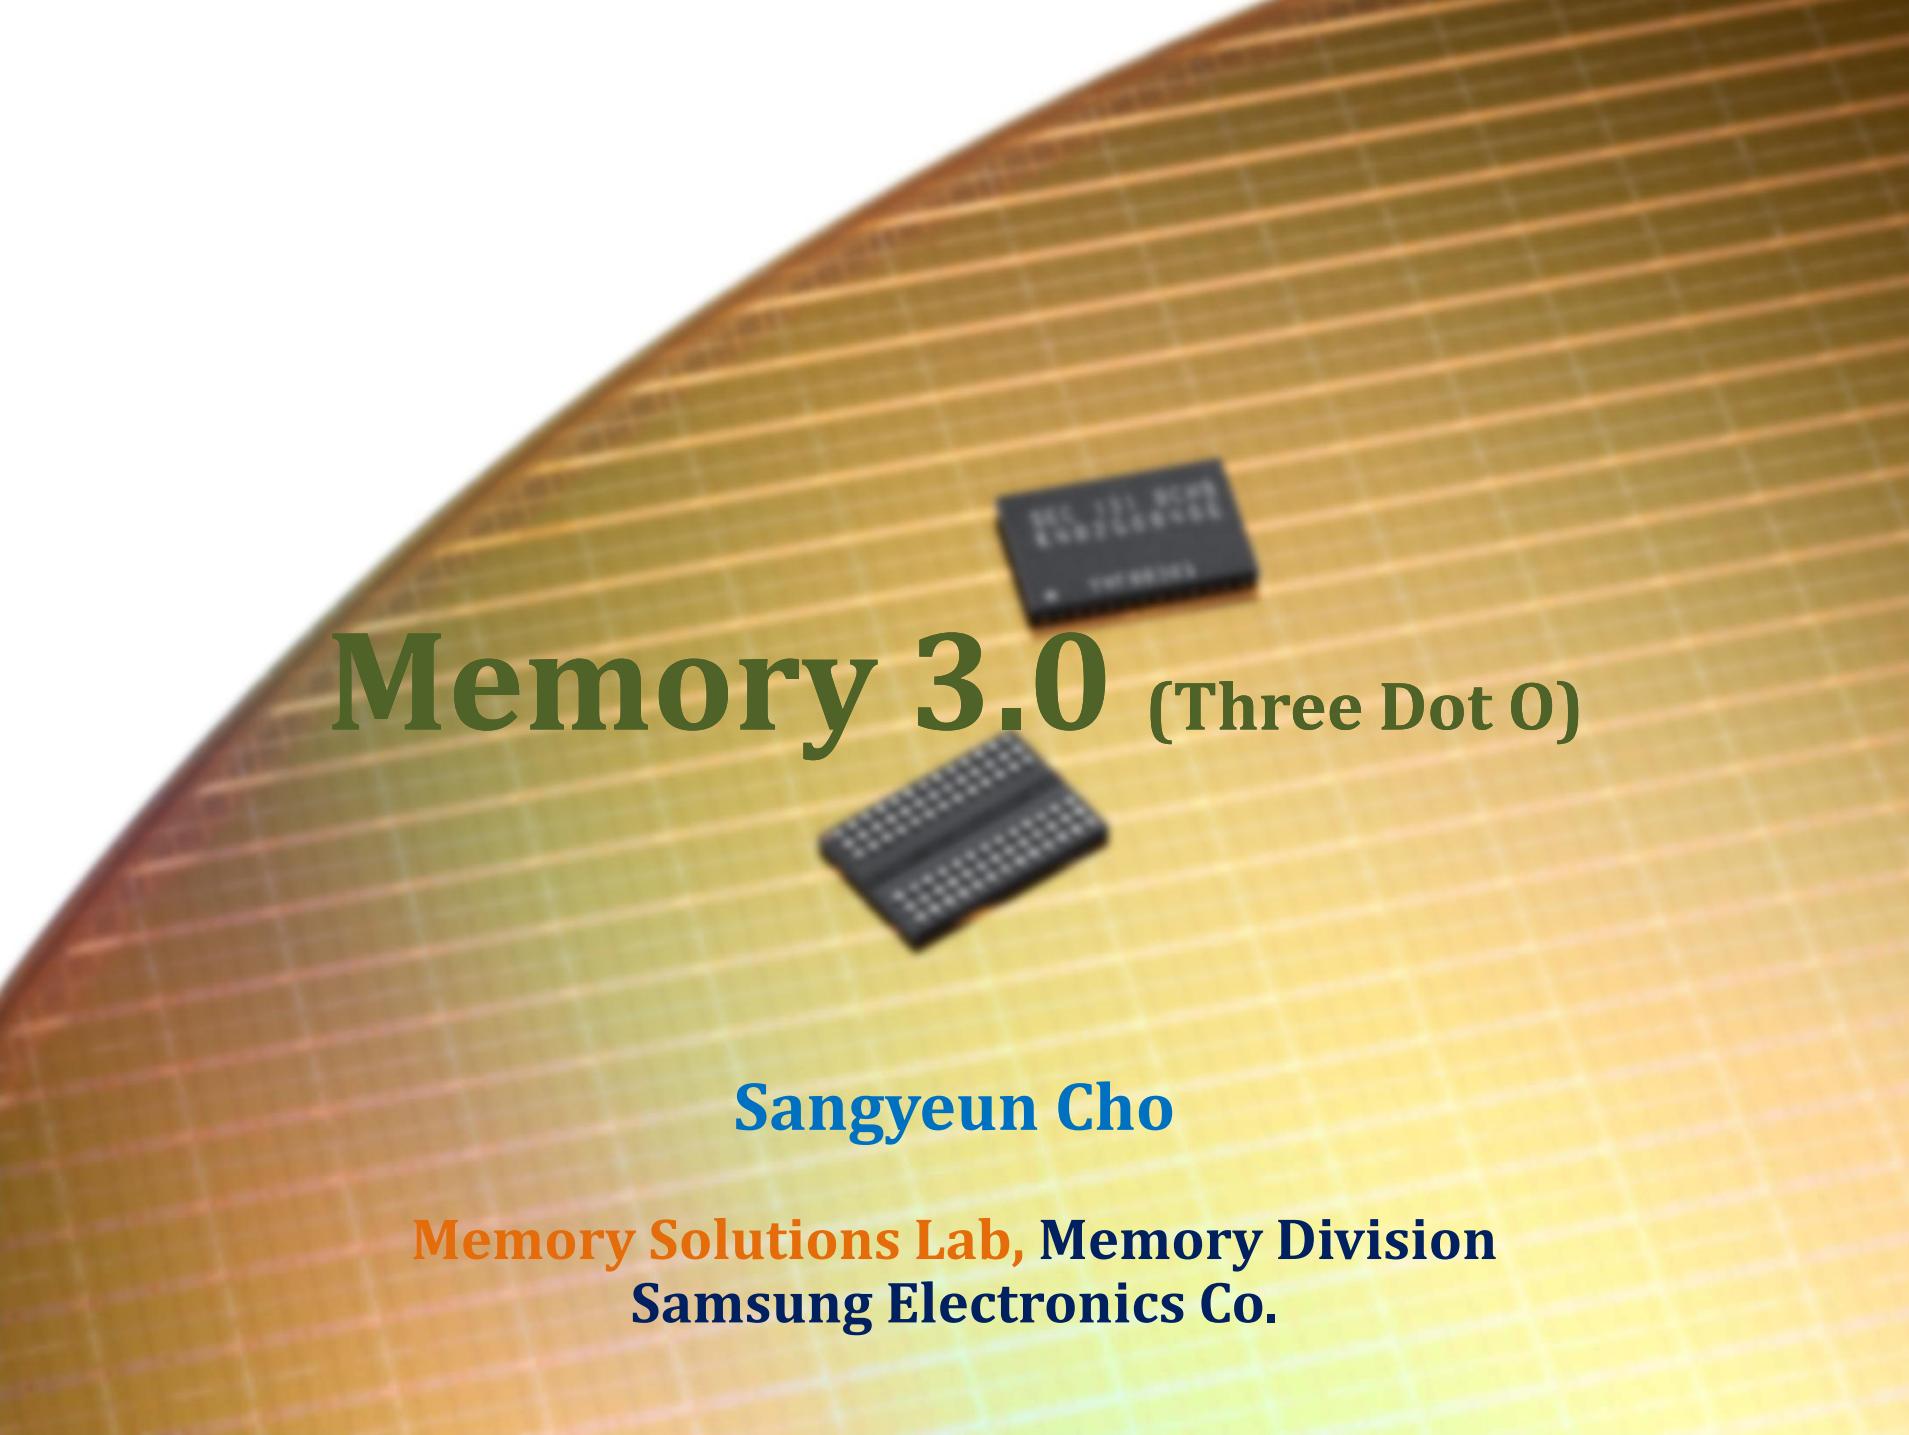
1,000,000,000,000,000,000,000

WHEN WRITTEN IT WOULD BE THE NUMBER 1 WITH 18 ZEROS



Summary of memory 3.0

- We are transitioning from memory 2.0 to a new era, when
 - Economic planar scaling of DRAM and flash becomes hard; creative scaling (e.g., 3D) expected;
 - New memory technologies are more interesting; and
 - New primary and secondary storage subsystems that increase the system capabilities and values will be of increasing importance (e.g., co-location vs. disaggregation)
- To succeed...
 - We need more **creativity** in defining and delivering new system-level memory solutions
 - We need far more **collaboration** in the systems areas



Memory 3.0 (Three Dot O)

Sangyeun Cho

Memory Solutions Lab, Memory Division
Samsung Electronics Co.