

# Advanced Memories

## 8<sup>th</sup> Fault-Tolerant Spaceborne Computing Employing New Technologies

May 28, 2015

Matt Marinella  
Sandia National Laboratories  
[mmarine@sandia.gov](mailto:mmarine@sandia.gov)

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

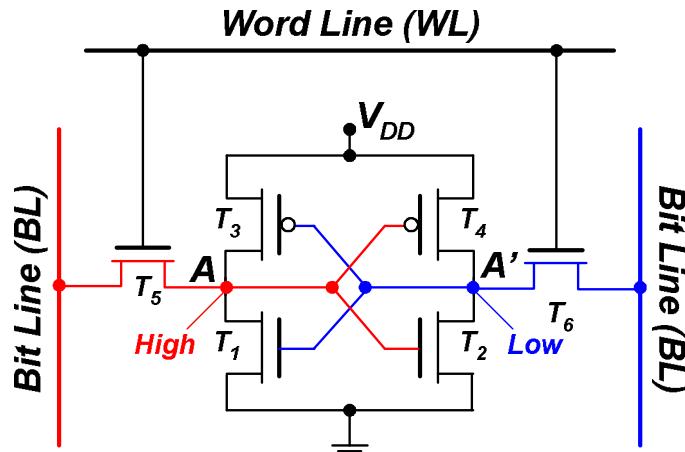
# Outline

- The Status Quo: SRAM, DRAM, and Flash
- Emerging Memory Technologies
- Session Overview

# Static Random Access Memory

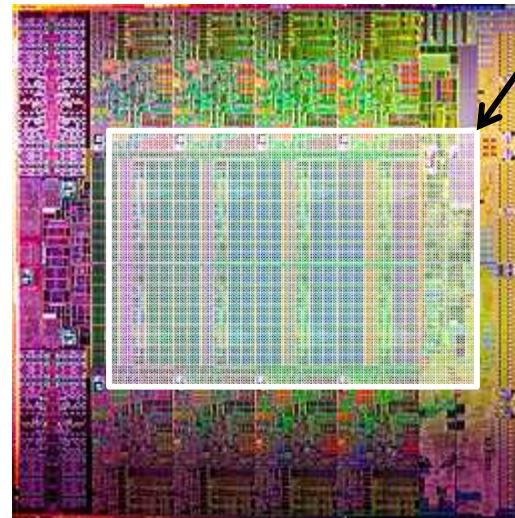
- Highest speed – <1 ns write time
- Larger cell size – typically requires 6 transistors
- Lower density than DRAM
- P depends strongly on  $f$
- Max SRAM on chip: Intel Xeon E7 – **45 MB SRAM cache**

## SRAM Cell Schematic



Courtesy Dieter Schroder, ASU

Intel Xeon E5

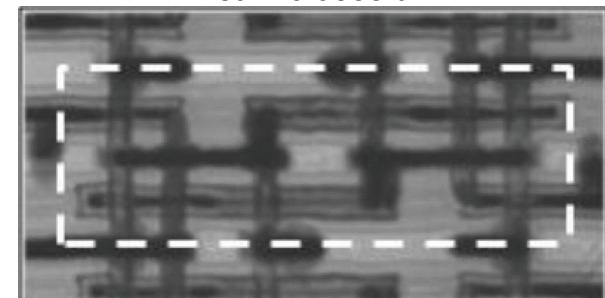


Intel.com

SRAM

Intel 14nm SRAM Cell

Area =  $0.0588 \mu\text{m}^2$

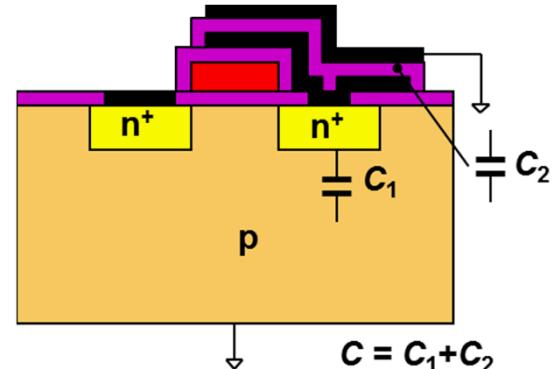


S Natarajan, IEDM 2014

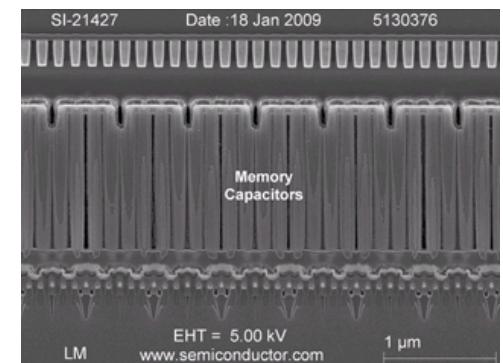
# Dynamic Random Access Memory

- State stored in capacitor charge
- Lower cost, higher density than SRAM
- Volatile and changes memory state if not refreshed periodically (64 ms)
- ~20 nm cells in production as of 2015
  - Is there a path to continued scaling?
- DRAM Challenges:
  - DRAMs struggling to maintain reasonable equivalent oxide thickness
  - Dielectric for cells below 20 nm still TBD
  - DDRx interfaces have high power requirements (although DDR4 is an improvement)

**Stacked DRAM Cell**

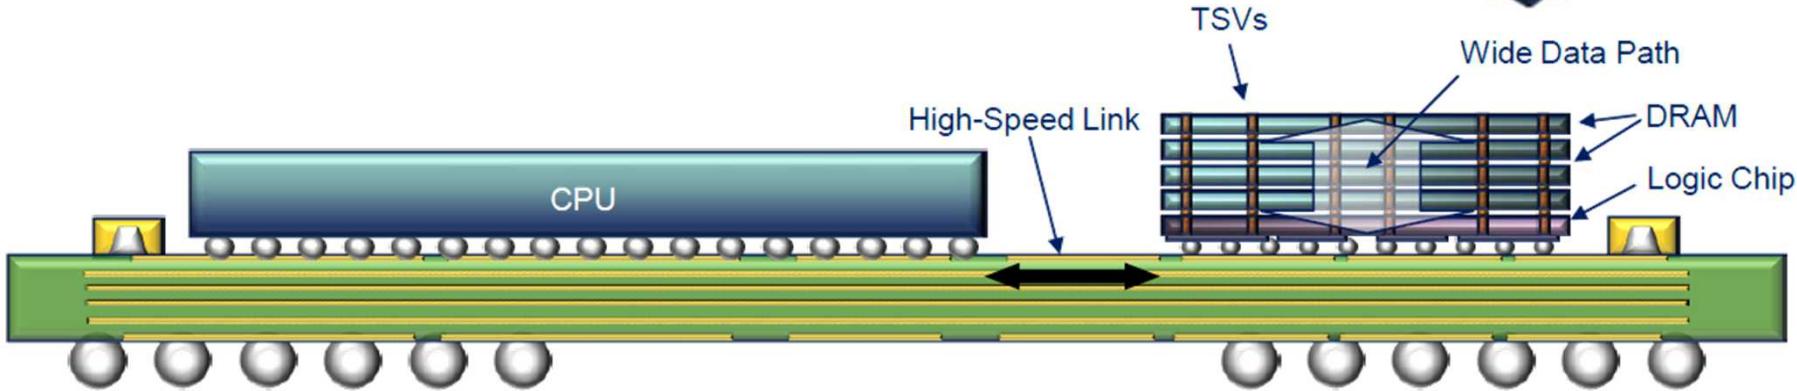
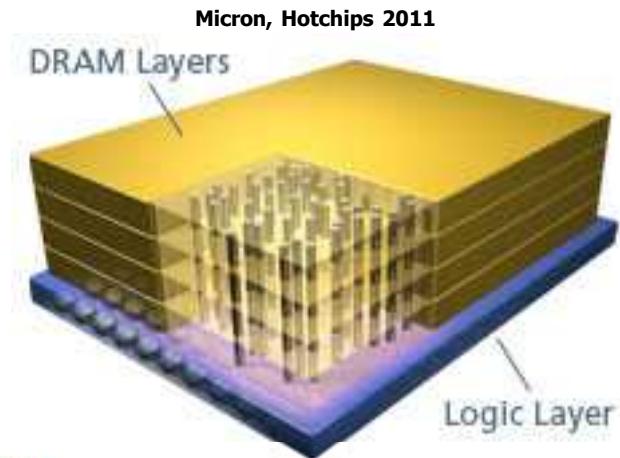


**Micron Stacked DRAM**



# 3D DRAM

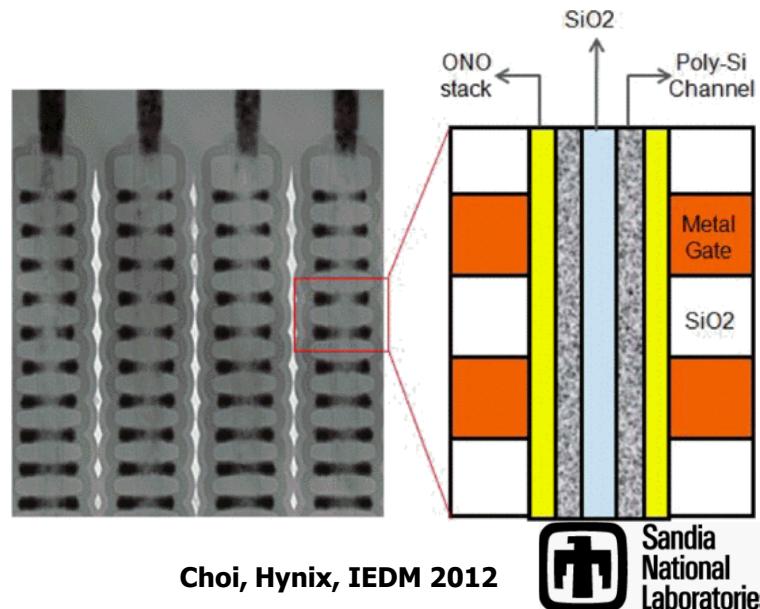
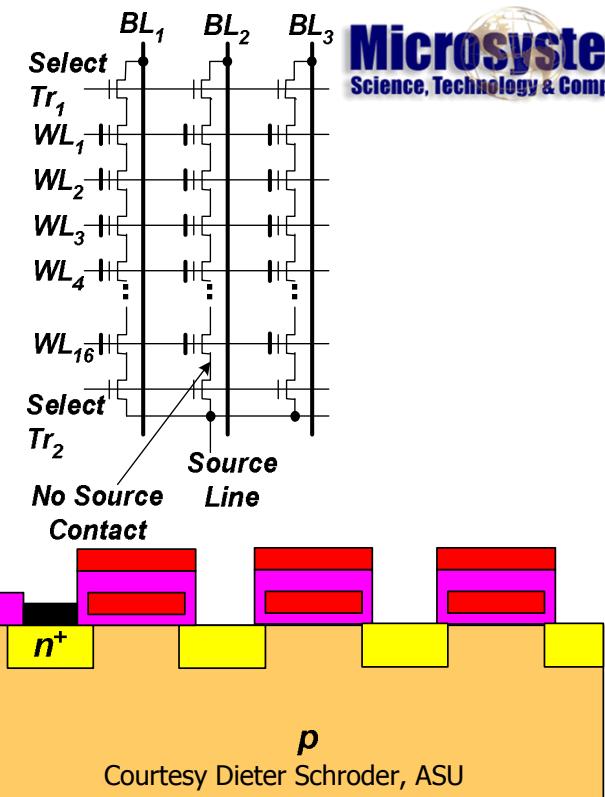
- Micron/Intel Hybrid Memory Cube
- DRAM die stacked on logic
- Connected via through-silicon-via
- Major energy savings



Technology	VDD	IDD	BW GB/s	Power (W)	mW/GB/s	pJ/bit	real pJ/bit
SDRAM PC133 1GB Module	3.3	1.50	1.06	4.96	4664.97	583.12	762
DDR-333 1GB Module	2.5	2.19	2.66	5.48	2057.06	257.13	245
DDRII-667 2GB Module	1.8	2.88	5.34	5.18	971.51	121.44	139
DDR3-1333 2GB Module	1.5	3.68	10.66	5.52	517.63	64.70	52
DDR4-2667 4GB Module	1.2	5.50	21.34	6.60	309.34	38.67	39
HMC, 4 DRAM w/ Logic	1.2	9.23	128.00	11.08	86.53	10.82	13.7

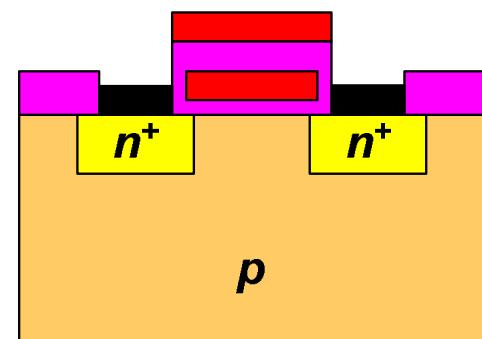
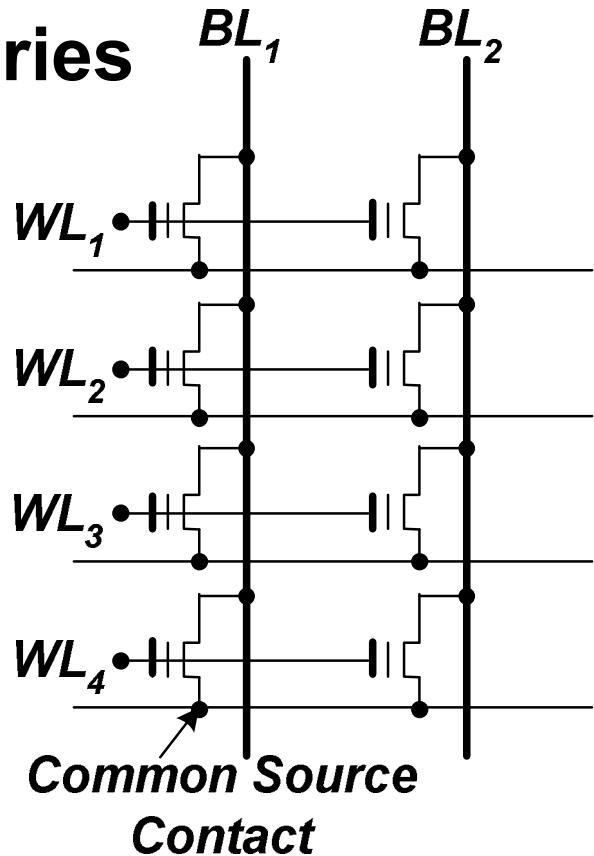
# NAND Flash Memories

- Serial access; slower than NOR
- Low bit cost
- High density:  $F \approx 18 \text{ nm}$  in 2015
- Small cell size (5-6  $F^2$ ), since no source contact required
- Monolithically stacked introduced in 2014 (Ex. Samsung VNAND)
- Block erase required
- Write/Erase: Fowler-Nordheim
- Challenges:
  - Non-scalable tunneling dielectric need  $> 6 \text{ nm}$  for retention
  - Floating gate interference: capacitance coupling between floating gates
  - Reduced coupling ratio with scaling



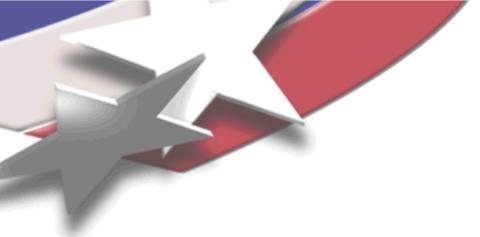
# NOR Flash Memories

- Fast random access, similar to RAM
- Lower voltage (7-10V)
- Write: Hot electron injection, high  $V_D$
- Erase: Fowler-Nordheim
- Erased as blocks
- Area:  $9-11F^2$  (need source contact)
- Embedded code (cellular phones, etc.)
- Challenges:
  - More severe drain disturbance with continuous scaling
  - Severely limited scaling below 32nm



# Outline

- The Status Quo: SRAM, DRAM, and Flash
- Emerging Memory Technologies
- Session Overview



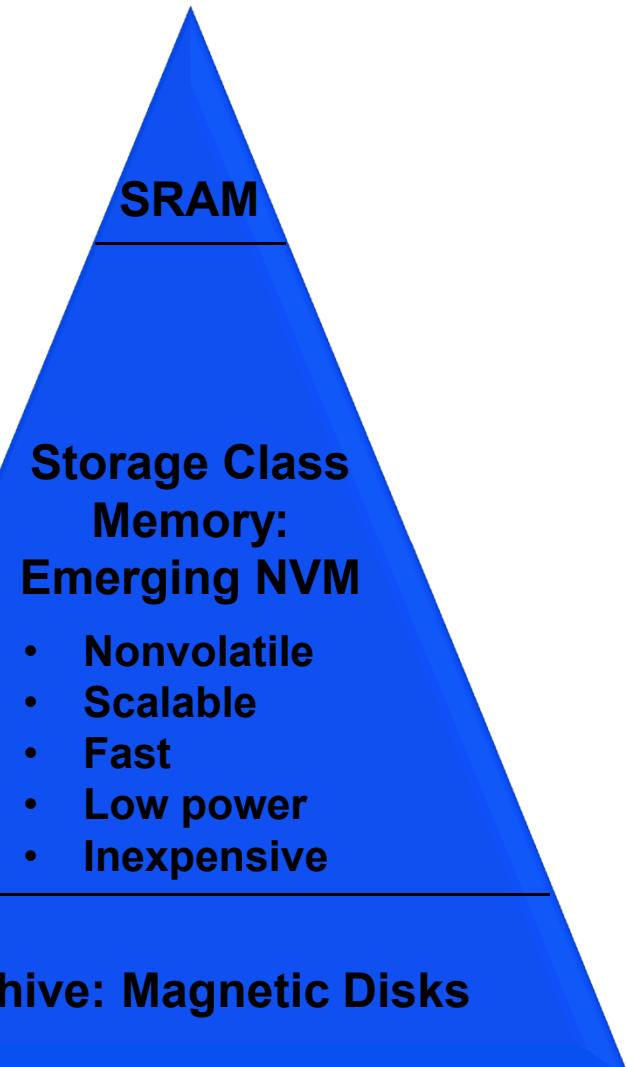
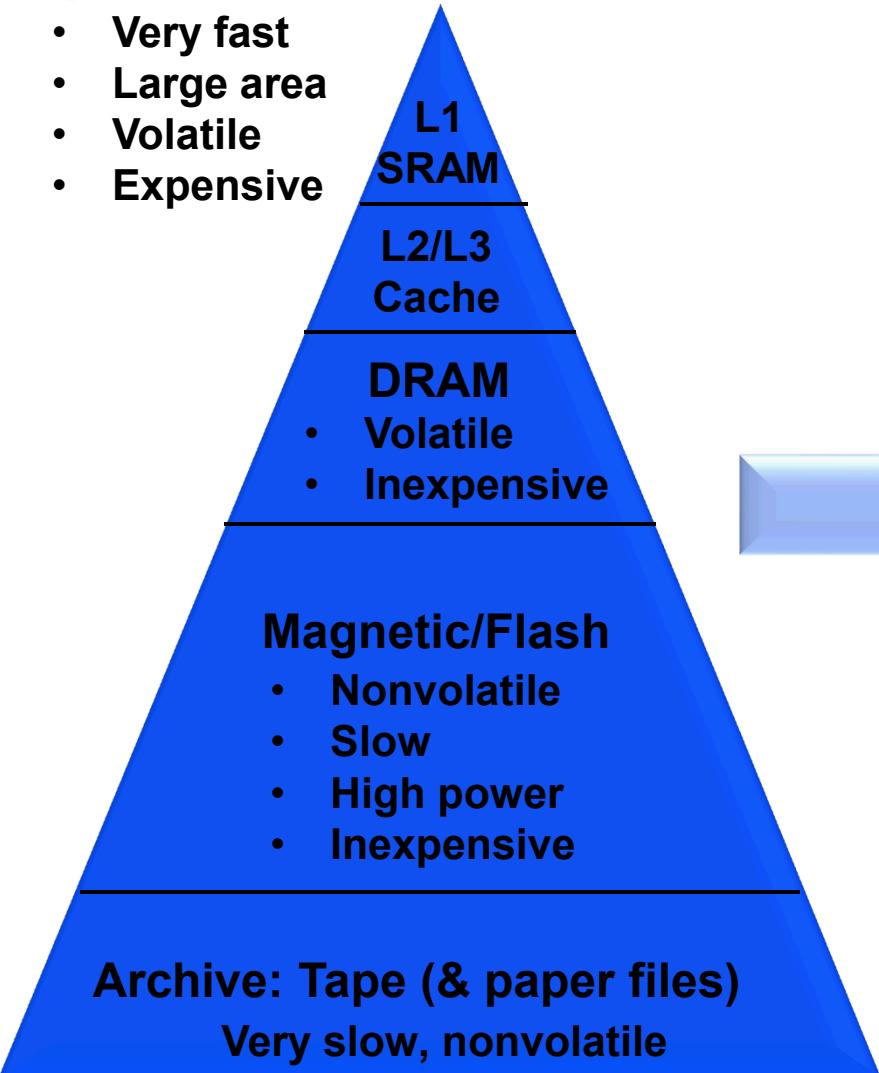
# Emerging Memory

- We are in a significant era for memory
- NAND Scaling:
  - Amazing progress in recent years: Samsung has a 32 layer process enabling 256 megabit per die
  - 3D will quench density issues temporarily
  - Reliability suffers with scaling; 12 nm is theoretical FG limit
- DRAM Scaling:
  - Struggling to maintain reasonable eq. oxide thickness
  - Dielectric for cells <20 nm still TBD
- Limitations in sight for both of these giants!
- Storage Class Memory
  - Magnetic to DRAM latency gap
- End of transistor scaling: no obvious new technology
- End of flash/DRAM scaling: several new technologies on the horizon!

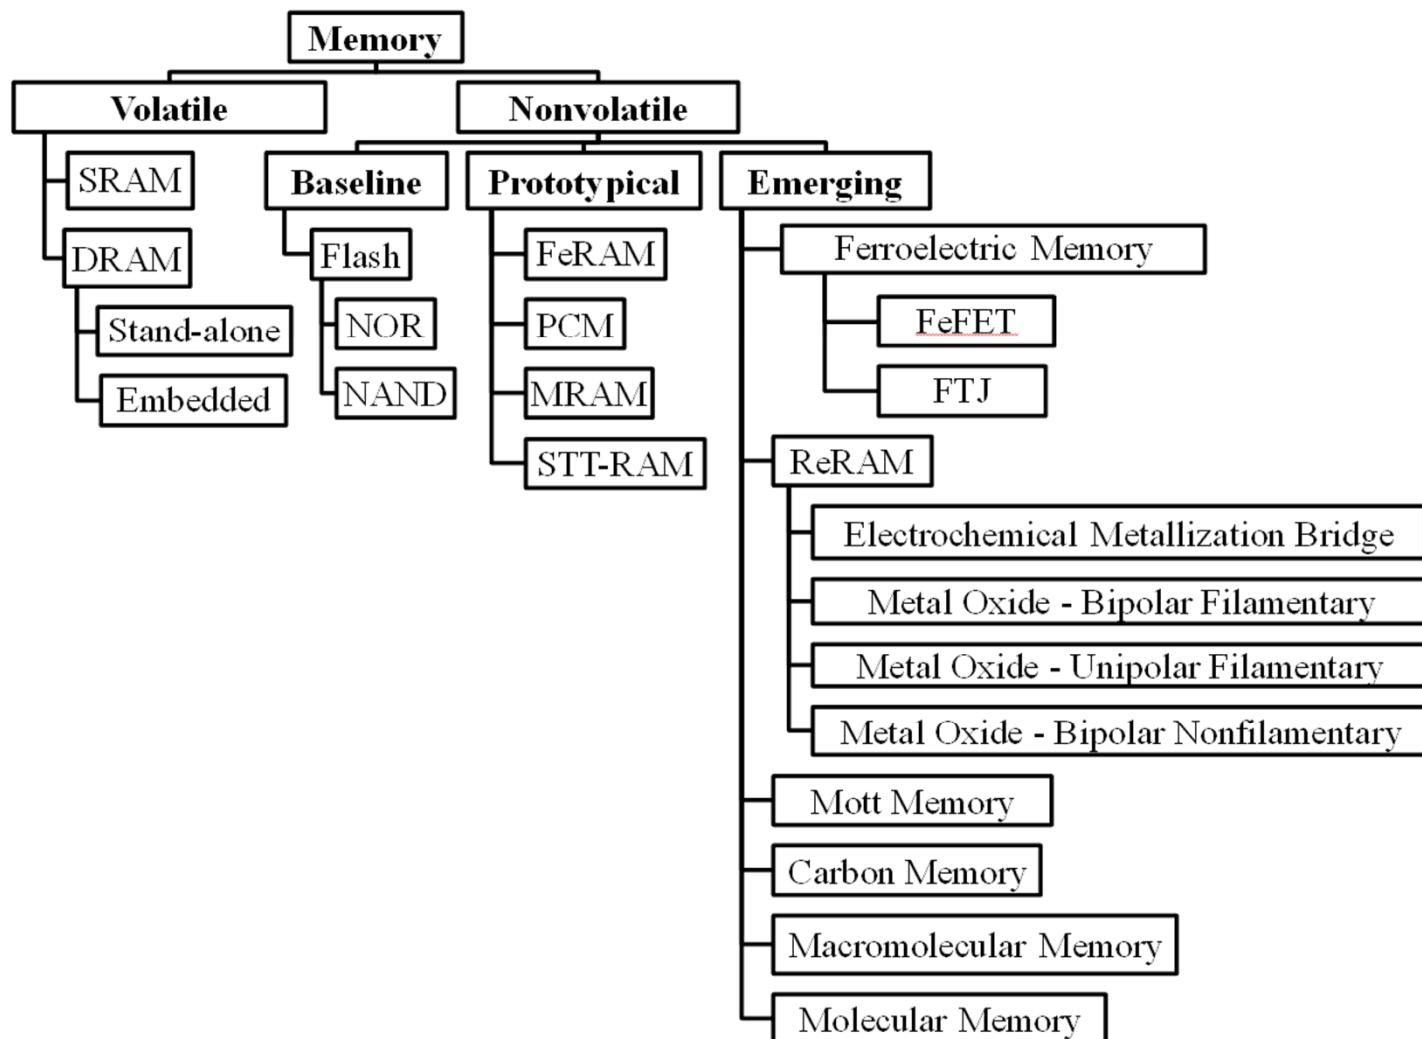


# Storage Class Memory: A Game Changer

- Very fast
- Large area
- Volatile
- Expensive

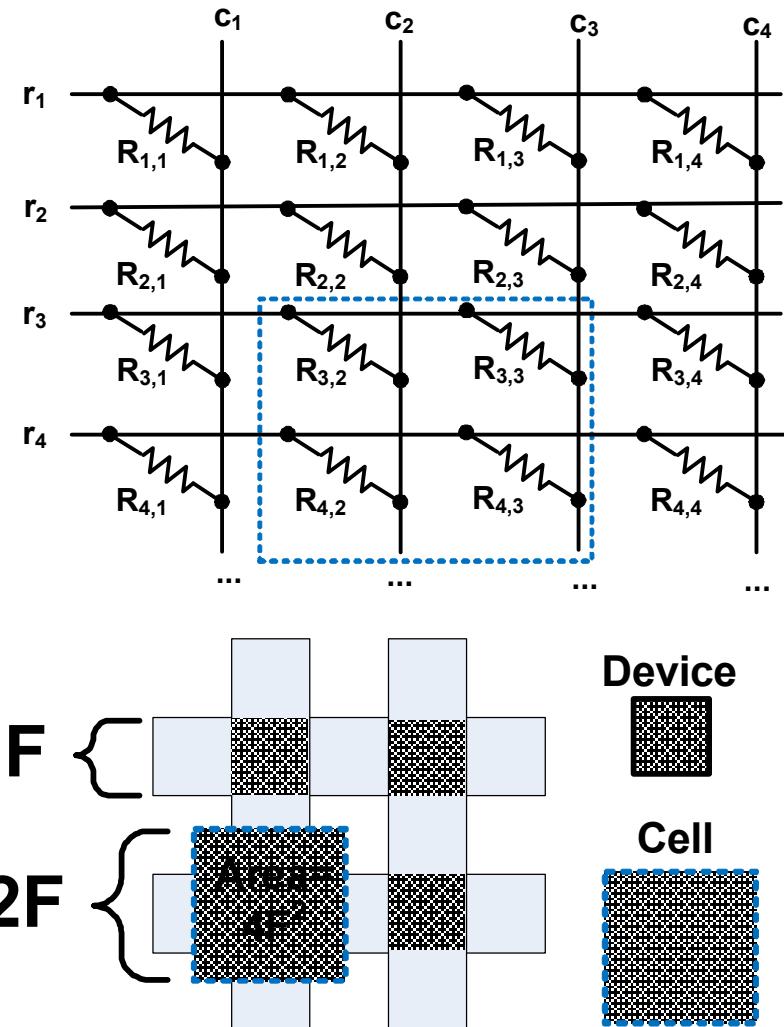
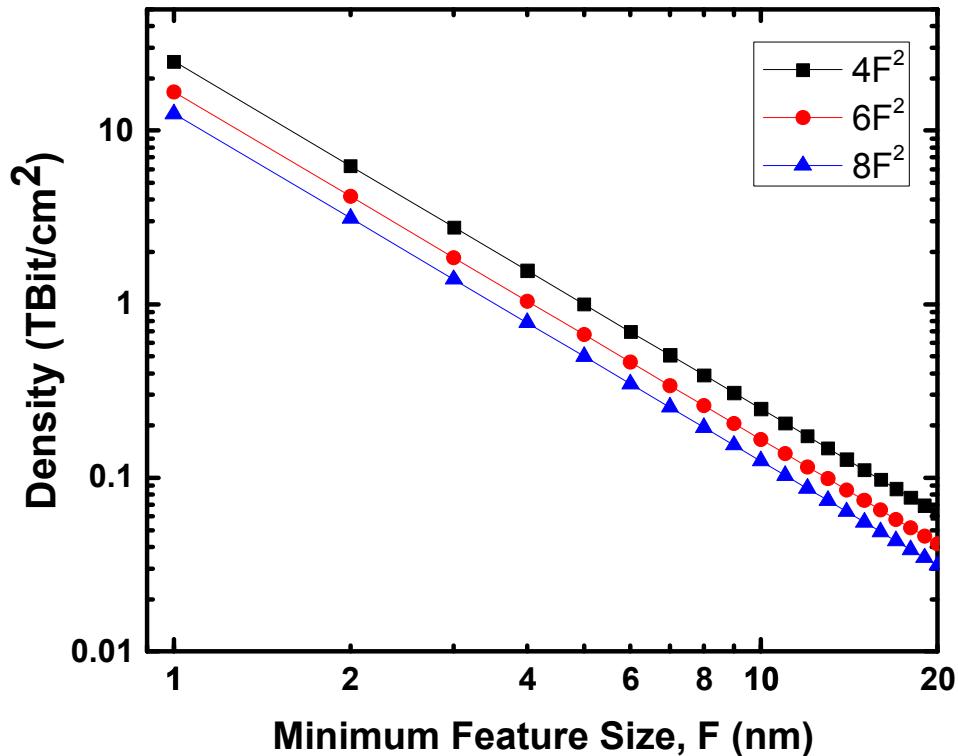


# Emerging Memory Taxonomy



# Resistive Crossbar Memories

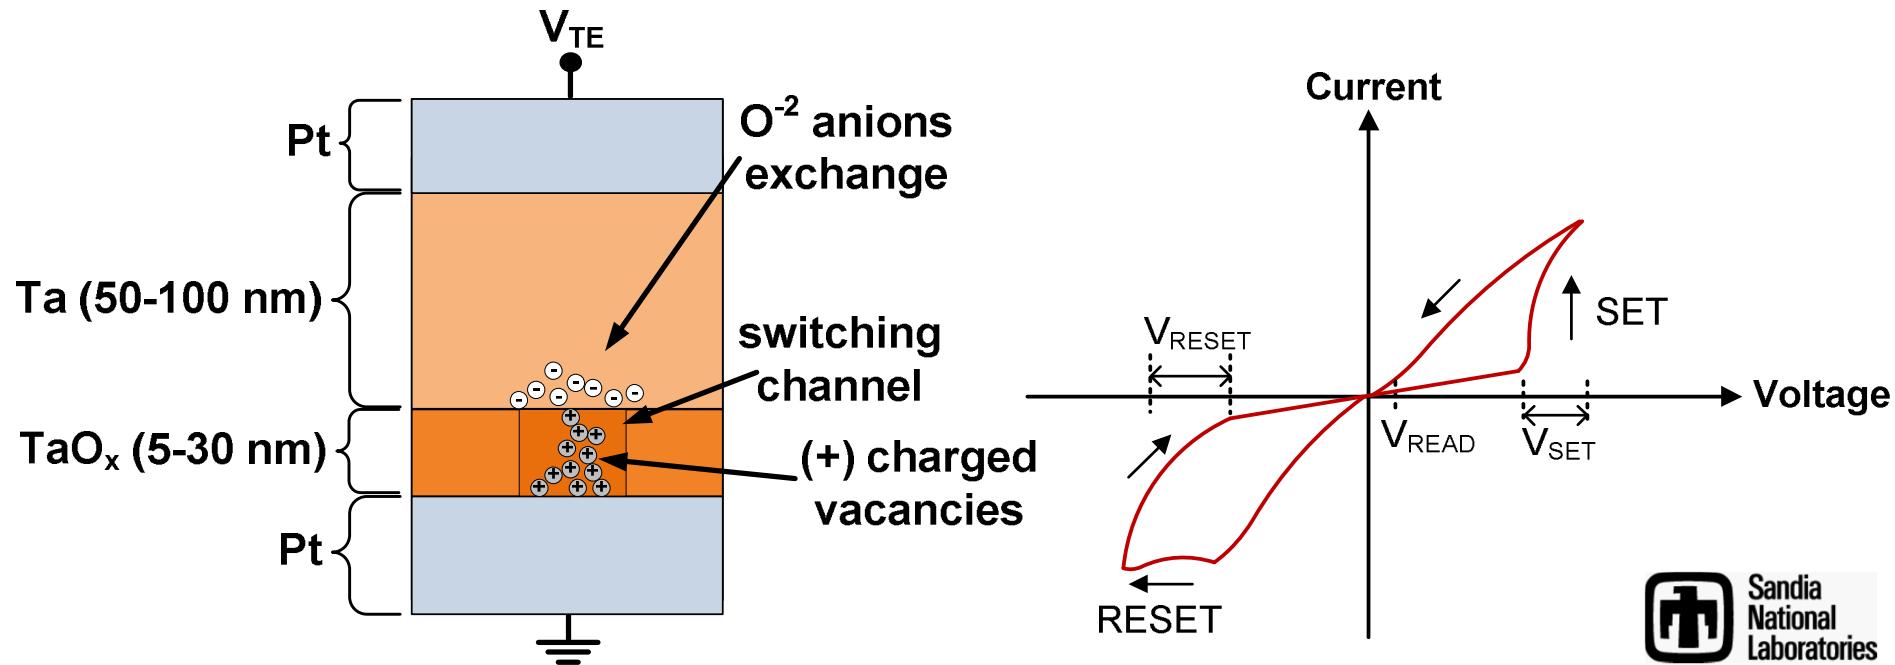
- $F$ =Feature size
- Max areal density possible  $\rightarrow 4F^2$



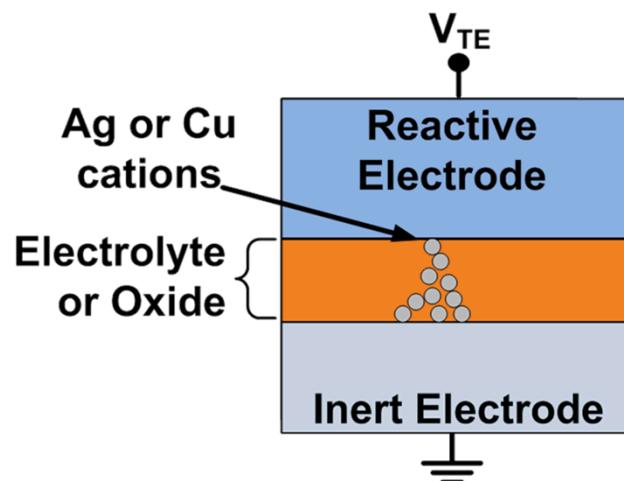
Marinella and Zhrinov, in Emerging  
Nanoelectronic Technologies, Wiley, 2014.

# Bipolar Metal Oxide ReRAM

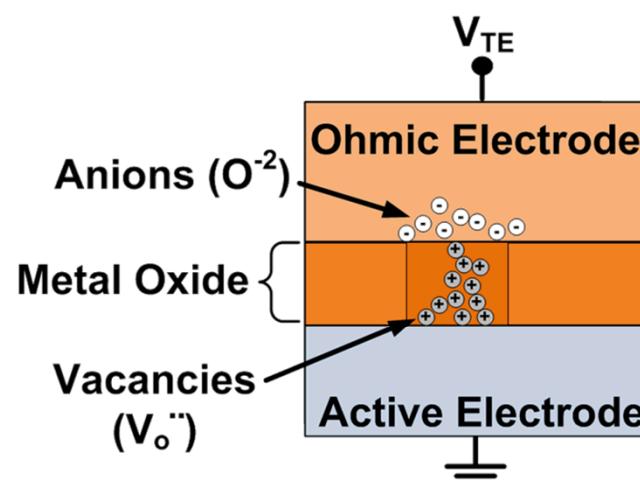
- “Hysteresis loop” is simple method to visualize operation
  - (real operation through positive and negative pulses)
- Hypothesized oxide resistance switching mechanism
  - Positive voltage/electric field: low R – O<sup>2-</sup> anions leave oxide
  - Negative voltage/electric field: high R – O<sup>2-</sup> anions return
- Common switching materials: TaO<sub>x</sub>, HfO<sub>x</sub>, TiO<sub>2</sub>, ZnO
- Despite progress, details of switching mechanism still debated



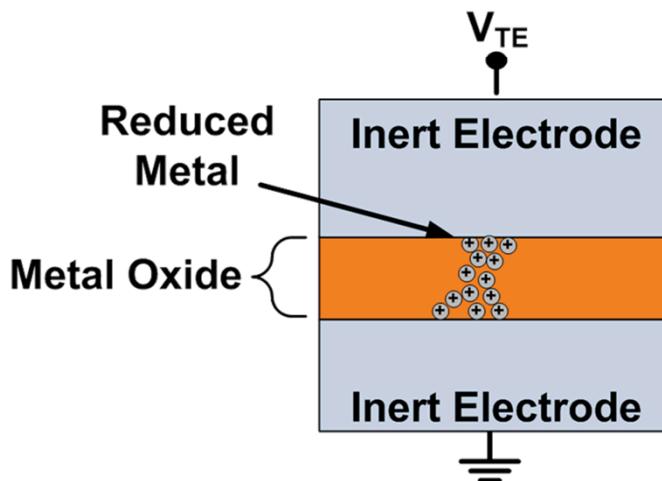
### Electrochemical Metallization Bridge



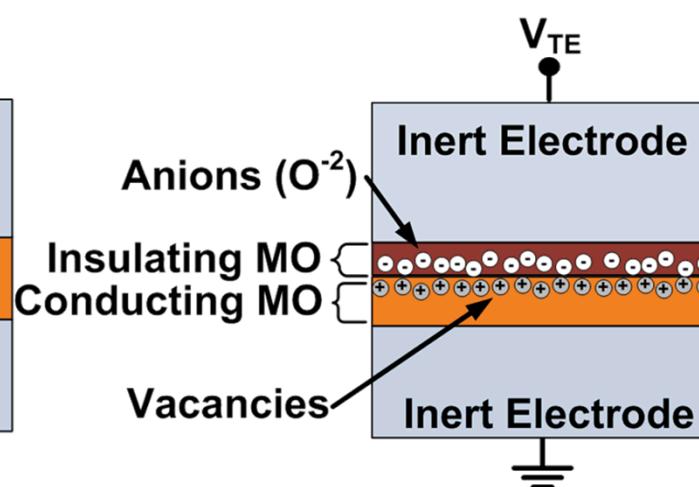
### Metal Oxide: Bipolar Filamentary



### Metal Oxide: Unipolar Filamentary

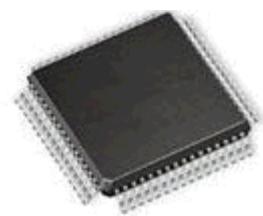
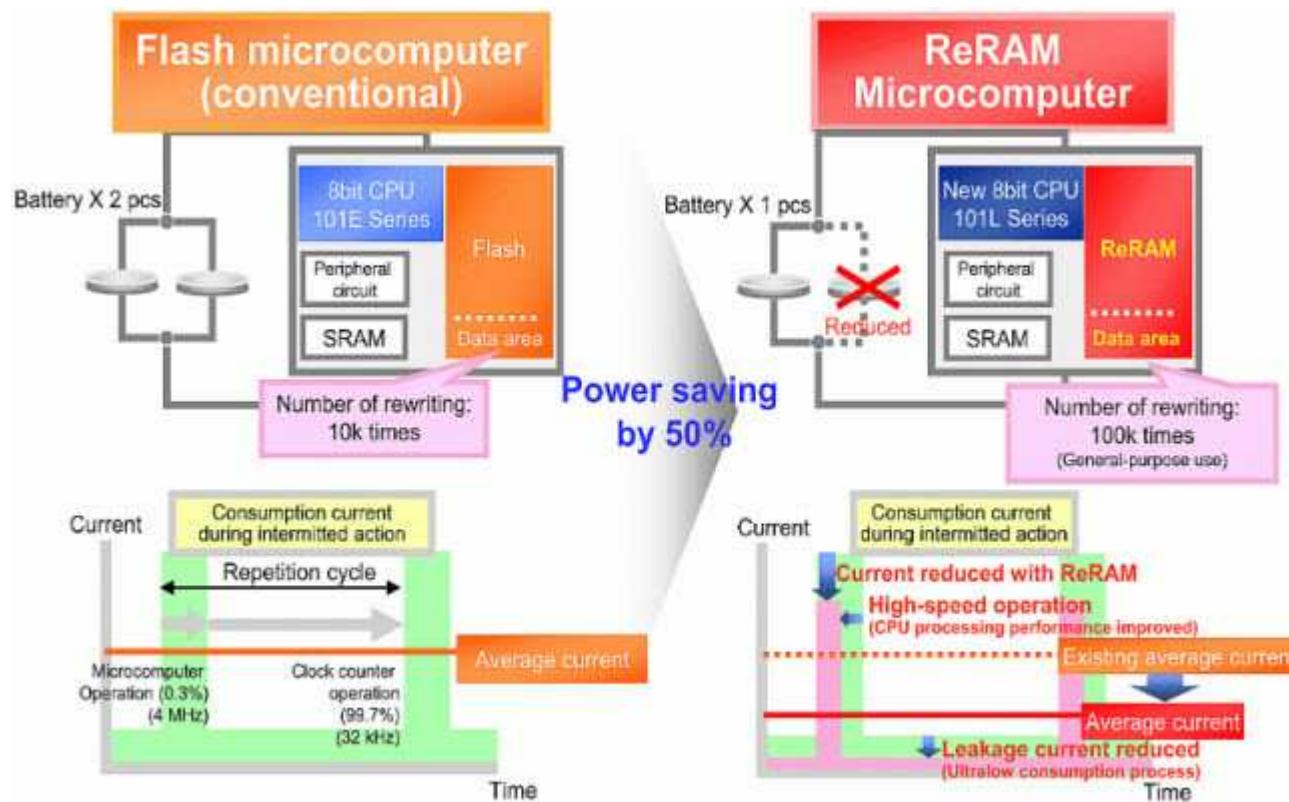


### Metal Oxide: Bipolar Non-Filamentary



# Panasonic MN101L ReRAM MCU

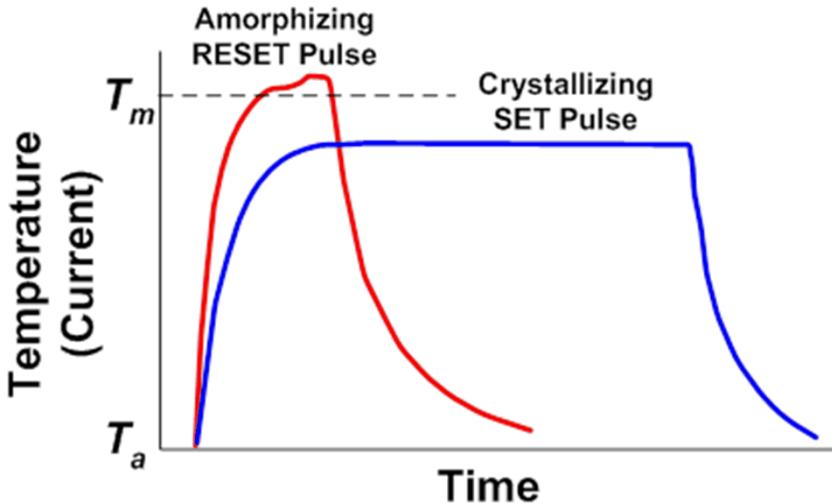
- First bipolar metal oxide commercial product
- Power and time saving over flash MCU



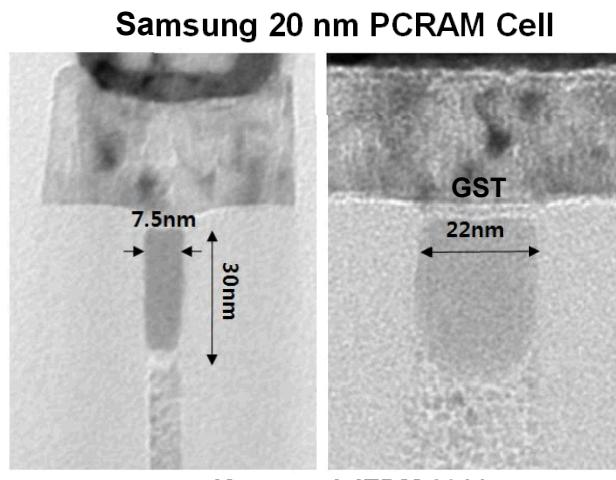
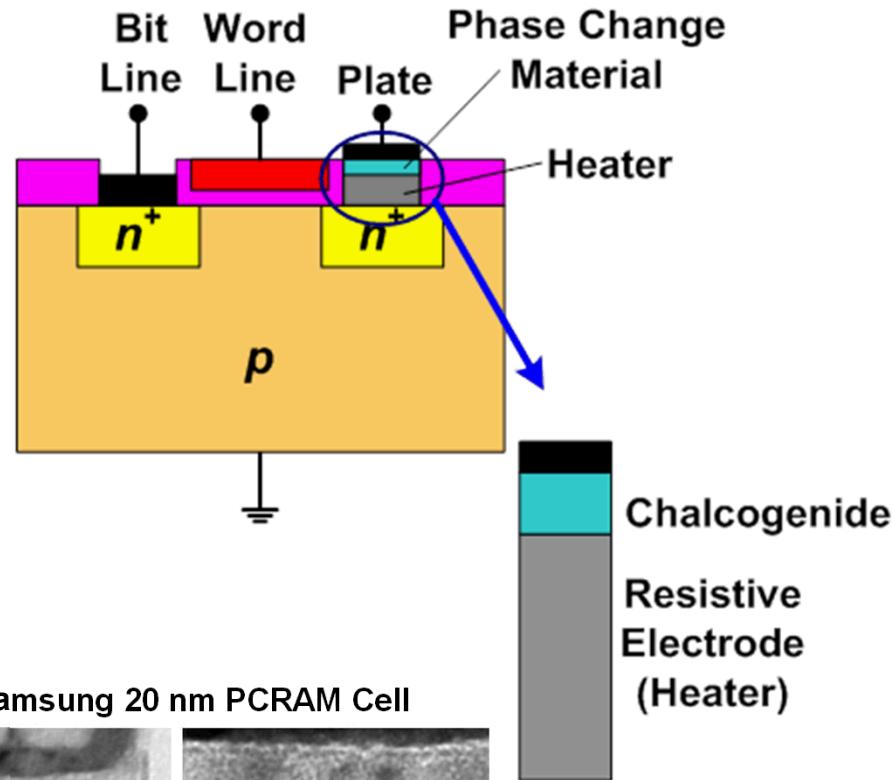
\* Please note that these values are subject to change without prior notice.

# Phase Change RAM

- Type of Resistive RAM
- GST most common material
- In commercial production
  - Samsung, Micron
- Set – crystallize, long pulse
- Reset – amorphize, short high current pulse



PCRAM Cell Schematic and Plot  
Courtesy D.K. Schroder, ASU



Kang et al, IEDM 2011

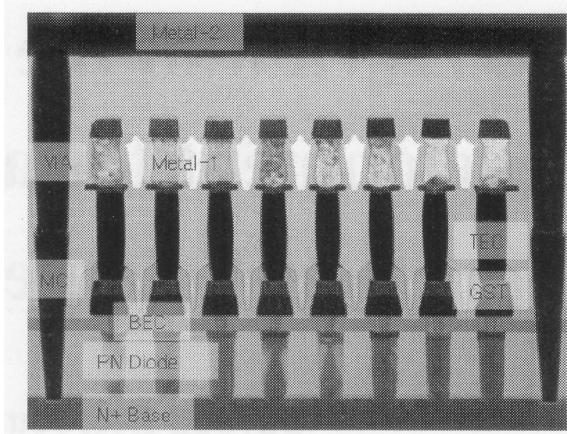
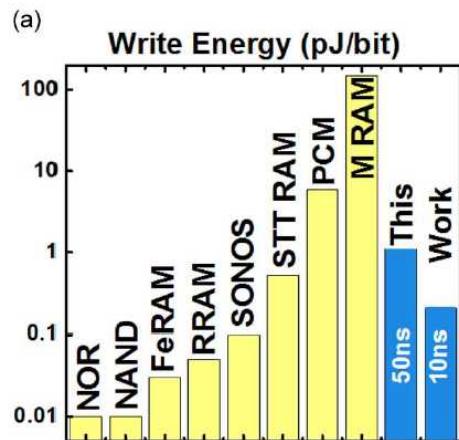
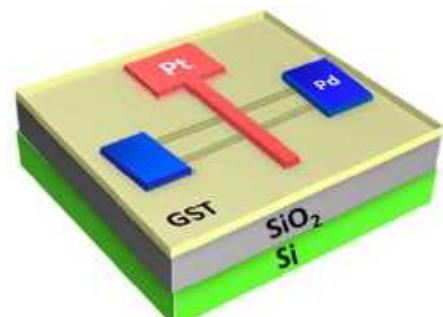
# Phase Change RAM

- Challenges

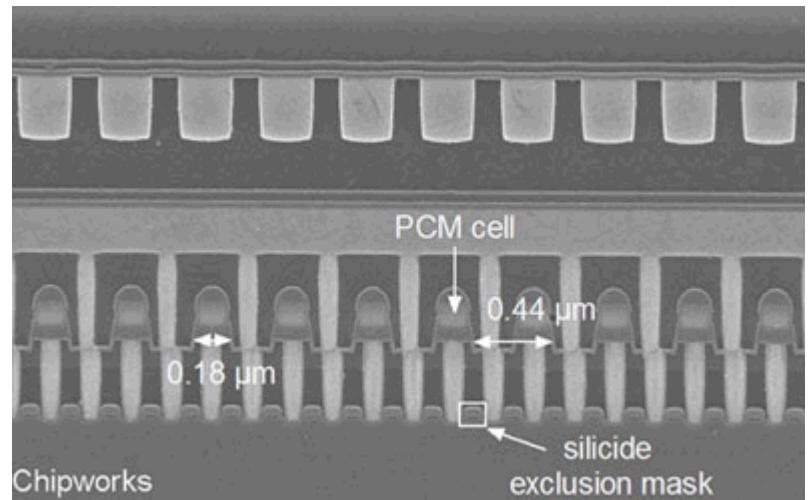
- High reset current (~ 500  $\mu$ A)
- Retention loss with scaling

- Possible solution

- Small contact area
- Heat confinement



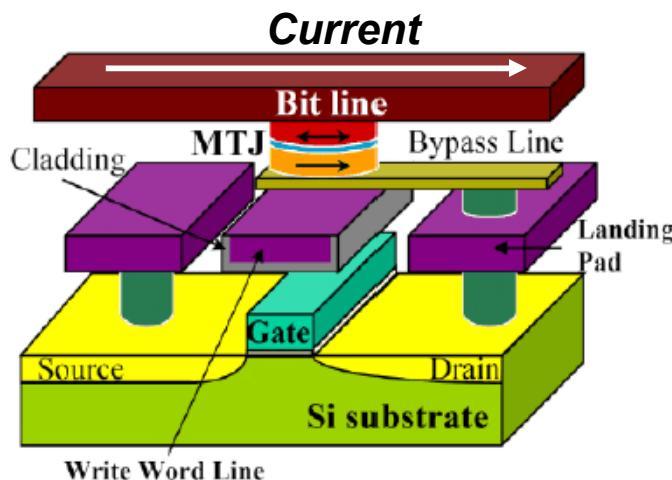
Samsung 512 Mb Array



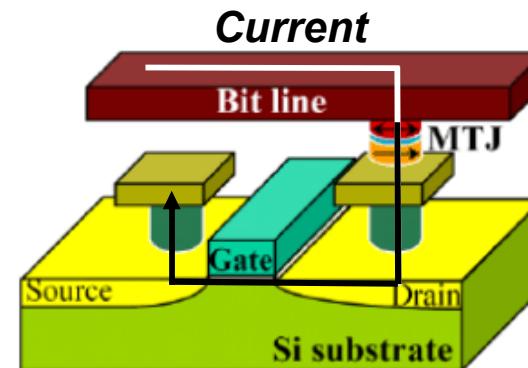
Numonyx PCM cell consists of a layer of  $\text{Ge}_2\text{Sb}_2\text{Te}_5$ , embedded in a dielectric structure and in contact with two electrodes

# Magnetic RAM

- Magnetic tunnel junction (MTJ)
- Field switched MRAM: complex cell architecture, high write current ( $\sim$  mA)
- Spin Transfer Torque: current through MTJ, much lower switching current ( $\sim$   $\mu$ A)
  - Given new life to the MRAM industry



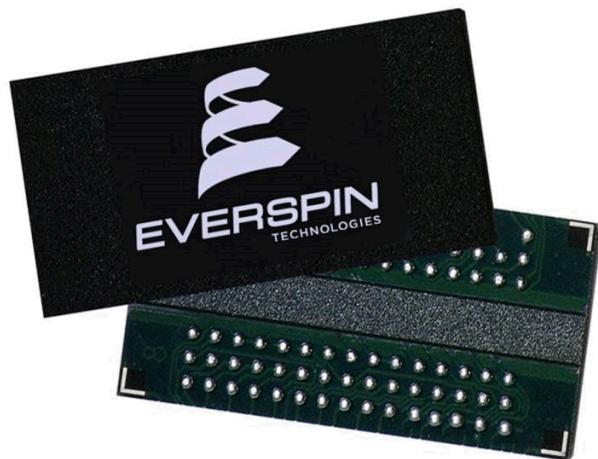
Field Switched MRAM



SST MRAM

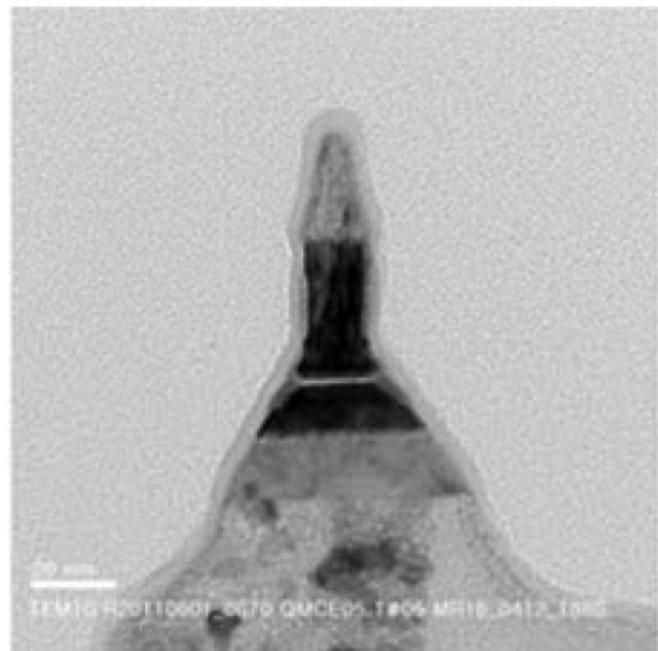
# MRAM: Current State of the Art

**Everspin DDR3  
Compatible STT-MRAM**



[Everspin.com](http://Everspin.com)

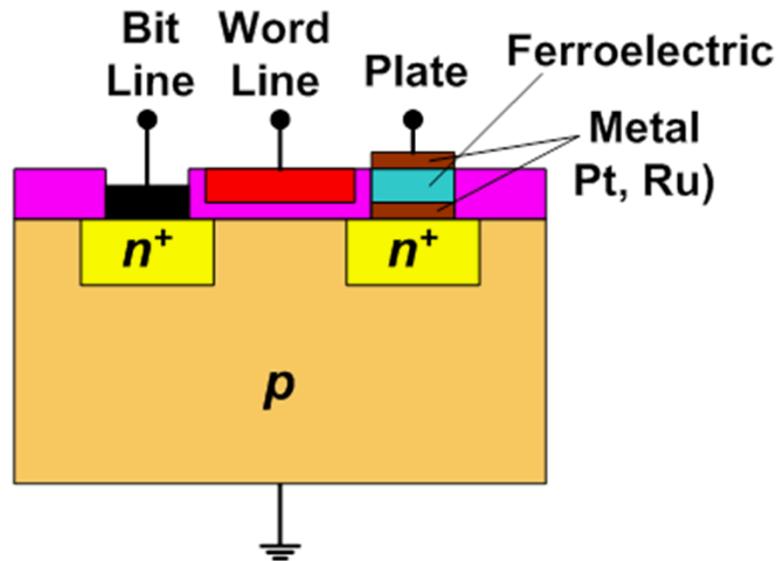
**Samsung 17 nm MTJ**



Kim et al, IEDM 2011

# Ferroelectric RAM

- Similar to DRAM cell
- Uses ferroelectric film capacitor
- State is stored as the polarization of the FE film
- Nearly unlimited endurance
- Moderate retention
- Process very finicky
- Commercial devices:
  - TI
  - RAMTRON (now Cypress)
  - Fujitsu

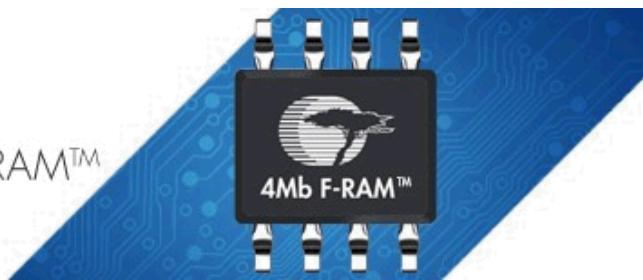


Courtesy D.K. Schroder

**CYPRESS INTRODUCES**

The industry's first 4-Mbit serial F-RAM™

[Cypress.com](http://Cypress.com)

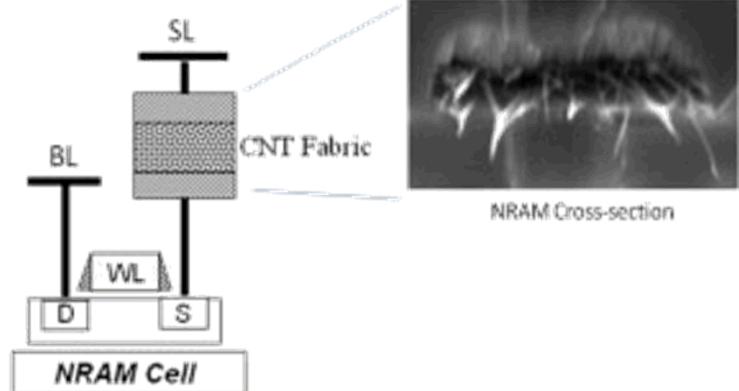
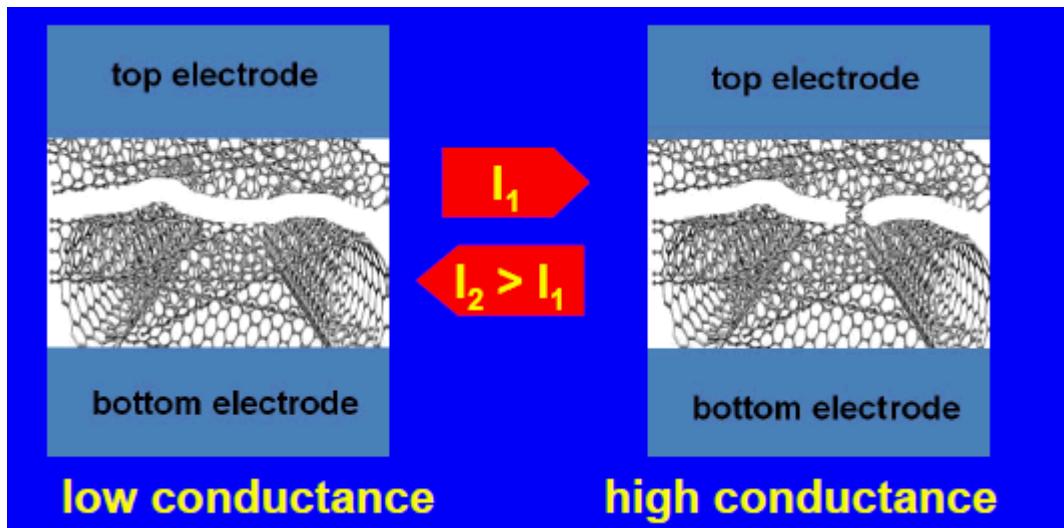


# Carbon Memory

Three material systems

1. Nanotube (single nanotube and layers)
2. Graphene
3. Amorphous carbon based resistive memory

Many possible mechanisms!



[nantero.com](http://nantero.com)

Kreupl, ERD Memory Workshop, 2014

# Emerging Memory Comparison



**Biggest challenge for ReRAM:  
Catch-up**

	DRAM	Flash (NOR-NAND)	ReRAM/Memristor	STT-MRAM	PC-RAM
	Production (30 nm)	Production (16 nm)	Development	Production (65 nm)	Production (45 nm)
Min device size (nm)	20	18	<10	16	<10
Density ( $F^2$ )	6	4+	4	8-20	$4F^2$
Read Time (ns)	< 10	$10^5$	2	10	20
Write Time (ns)	< 10	$10^6$	2	13	50
Write Energy (pJ/bit)	0.005	100	<1	4	6
Endurance (W/E Cycles)	$>10^{16}$	$10^4$	$10^{12}$	$10^{12}$	$>10^9$
Retention	64 ms	> 10 y	> 10 y	weeks	> 10 y
BE Layers	FE	FE	4	10-12	4
Process complexity	High/FE	High/FE	Low/BE	High/BE	Low/BE

**Biggest challenge for STT-MRAM: Balancing  
Retention/Scaling/Temperature/Write current**

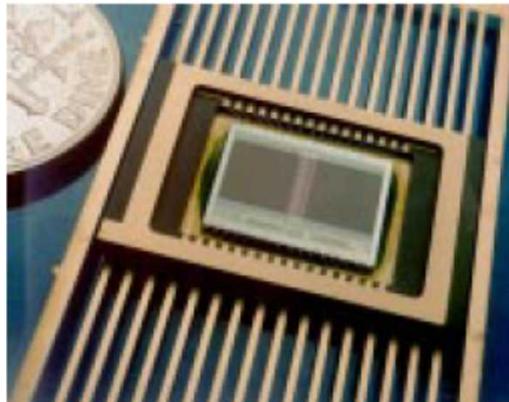
**Biggest challenge for PCM:  
High erase current**

**\*\*\*DISCLAIMER: Due to 10s of thousands of references on these technologies –  
many of these numbers are not universally agreed on!**

# What are the implications for space computing?

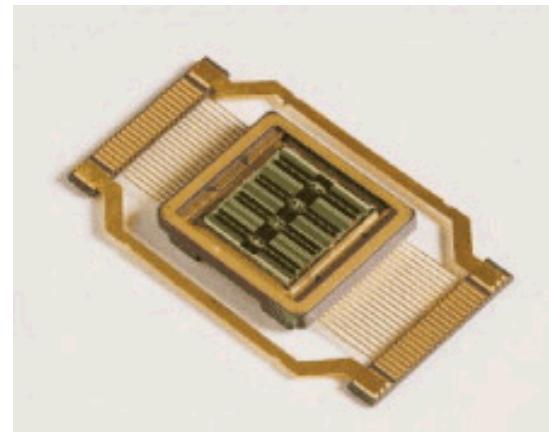
- Where have many emerging memories ended up?
  - As a rad-hard product targeting aerospace applications!
- Commercially available rad-hard nonvolatile memories
- NG EEPROM: 1Mbit, 100ms write,  $10^4$  cycles, 1.25µm RHC莫斯
- BAE C-RAM: 4Mbit (planned 20 Mbit), 70ns write
- Honeywell MRAM: 16Mbit die, 140ns write,  $10^{12}$  cycles
- *Rad-hard memory requires a rad-hard CMOS base process*

**NG Rad-hard EEPROM**



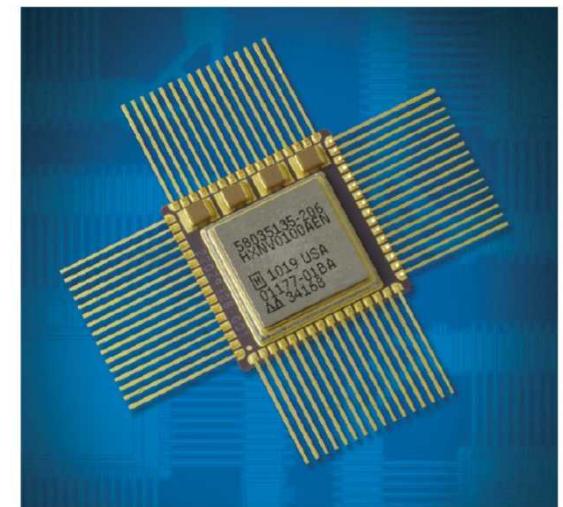
Rad Hard 256K EEPROM  
[northropgrumman.com](http://northropgrumman.com)

**BAE C-RAM**



[baesystems.com](http://baesystems.com)

**Honeywell M-RAM**



[honeywell.com](http://honeywell.com)

# Questions?

# Outline

- The Status Quo: SRAM, DRAM, and Flash
- Emerging Memory Technologies
- Session Overview

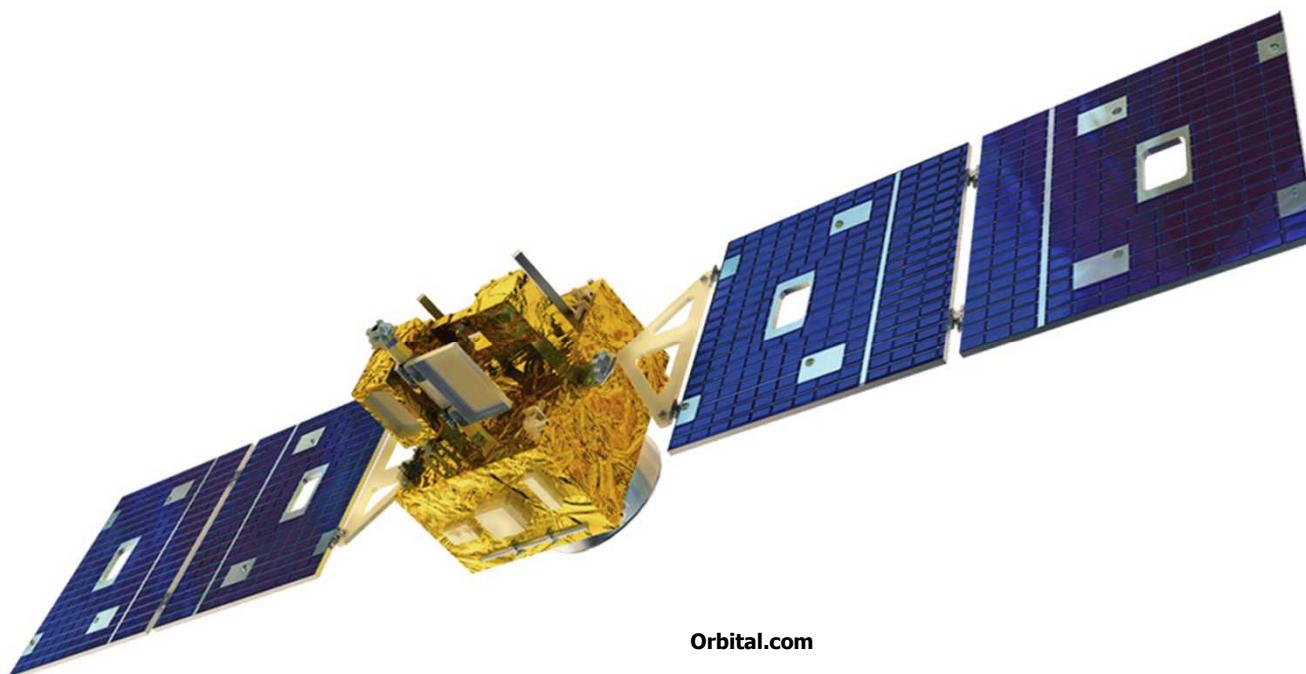
# Session Overview

- **7:00 Introduction and Emerging Memory Technologies**
  - Matt Marinella
- **7:20 Resistive Memory for Space Applications**
  - David Hughart
- **7:30 Hybrid Memory Cube**
  - Dave Resnick
- **8:00 Processor in Memory and Storage**
  - Erik DeBenedictis
- **8:30 Discussion**
- **8:45 Wrap-up, conclusions and next steps...**

# Discussion

## Key Constraints for Memory Systems in Space

1. Environmental constraints
2. Interfacing and SWaP (size weight and power) constraints
3. Usage patterns



# Environmental Constraints

- **Memory must be able to withstand**
  - Radiation effects
  - Thermal cycling
  - Launch survival
  - Mission Duration
    - CubeSat life is 18months
    - As high as 100 years for deep space missions
- **Also requires:**
  - Adequate cooling
  - Hermetic sealing, esp if a device has problems operating in a vacuum

# SWaP and Interface Considerations

- **Size, Weight, and Power**
  - Cubesat: 10x10x10 cm, 3 lbs
- **Must consider redundancy requirements**
- **Emerging memories have significantly lower read and write power than magnetic, flash, and DRAM**
- **Few space compatible high speed interfaces**



Nasa.gov

# Usage Patterns

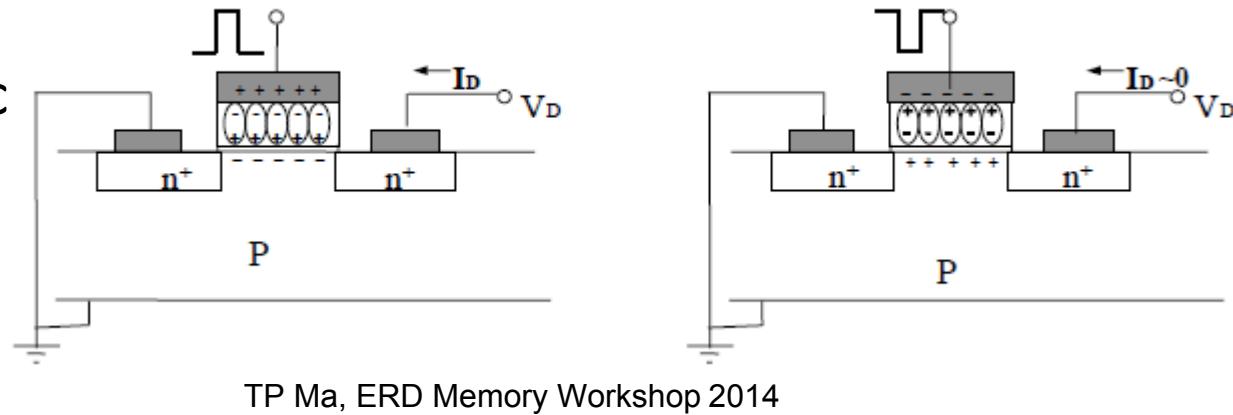
- Terrestrial usage patterns of storage devices:  $R \gg W$ 
  - More reads than writes
  - Standard disk devices are tuned for  $R \gg W$
  - File system layout assumes  $R \gg W$
- Space usage patterns:  $W \geq R$ 
  - Number of write can be equal to or greater than number of reads
  - Volatile and non-volatile storage is important
  - Patterns can be large block streaming or small entities
  - Random access can occur for either of those patterns. Sometimes there are more writes than reads as data is discarded after initial assessment.
- These requirements impact both device technologies and also on the file system layout
- Hardware/Software codesign

# Final Points, Summary, Next Steps

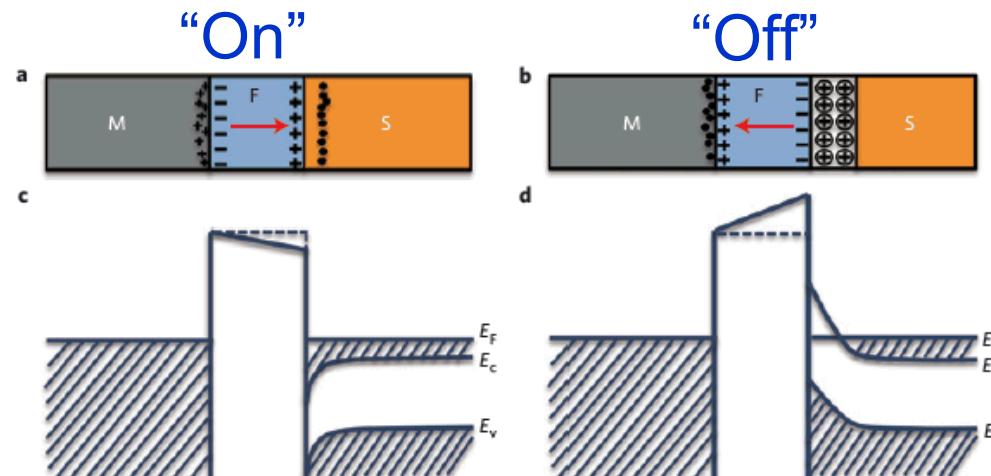
# Backup Slides

# Emerging Ferroelectric Memories

Ferroelectric  
FET



Ferroelectric  
Tunnel  
Junction

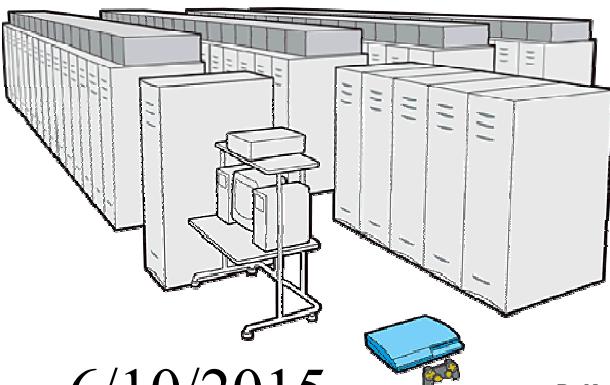


# ITRS Requirements for SCM

Parameter	Benchmark [A]			Target	
	HDD [B]	NAND flash [C]	DRAM	Memory-type SCM	Storage-type SCM
<i>Read/Write latency</i>	3-5 ms	~100µs (block erase ~1 ms)	<100 ns	<100 ns	1-10µs
<i>Endurance (cycles)</i>	unlimited	$10^4$ - $10^5$	unlimited	$>10^9$	$>10^6$
<i>Retention</i>	>10 years	~10 years	64 ms	>5 days	~10 years
<i>ON power (W/GB)</i>	~0.04	~0.01-0.04	0.4	<0.4	<0.04
<i>Standby power</i>	~20% ON power	<10% ON power	~25% ON power	<1% ON power	<1% ON power
<i>Areal density</i>	$\sim 10^{11}$ bit/cm <sup>2</sup>	$\sim 10^{10}$ bit/cm <sup>2</sup>	$\sim 10^9$ bit/cm <sup>2</sup>	$>10^{10}$ bit/cm <sup>2</sup>	$>10^{10}$ bit/cm <sup>2</sup>
<i>Cost (\$/GB)</i>	0.1	2	10	<10	<3-4

# Supercomputers

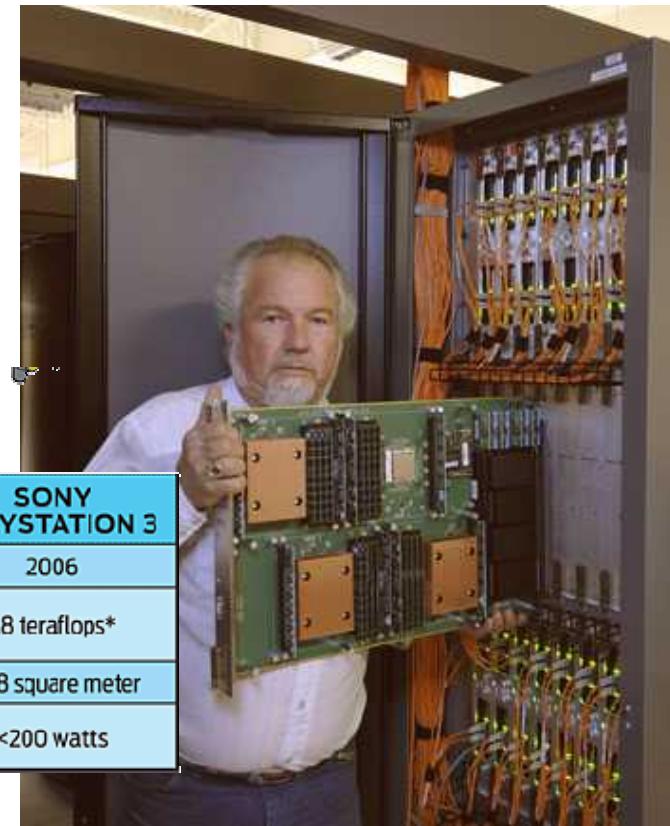
- FLOPS: floating point operations per second
- Exaflop:  $10^{18}$  operations per second
- US would like to have an Exascale Computer by 2018(ish)
- **Exascale computers will have a lot of hardware**
- 10-100 petabytes main memory
  - 10-100 million DRAM chips
- 100's of exabytes storage
  - Millions of hard drives



TALE OF THE TAPE:  
SUPERCOMPUTER  
VS. GAME CONSOLE

	SANDIA LAB'S ASCI RED	SONY PLAYSTATION 3
DATE OF ORIGIN	1997	2006
PEAK PERFORMANCE	1.8 teraflops	1.8 teraflops*
PHYSICAL SIZE	150 square meters	0.08 square meter
POWER CONSUMPTION	800 000 watts	<200 watts

\* For GPU: CPU adds another 0.3 teraflops



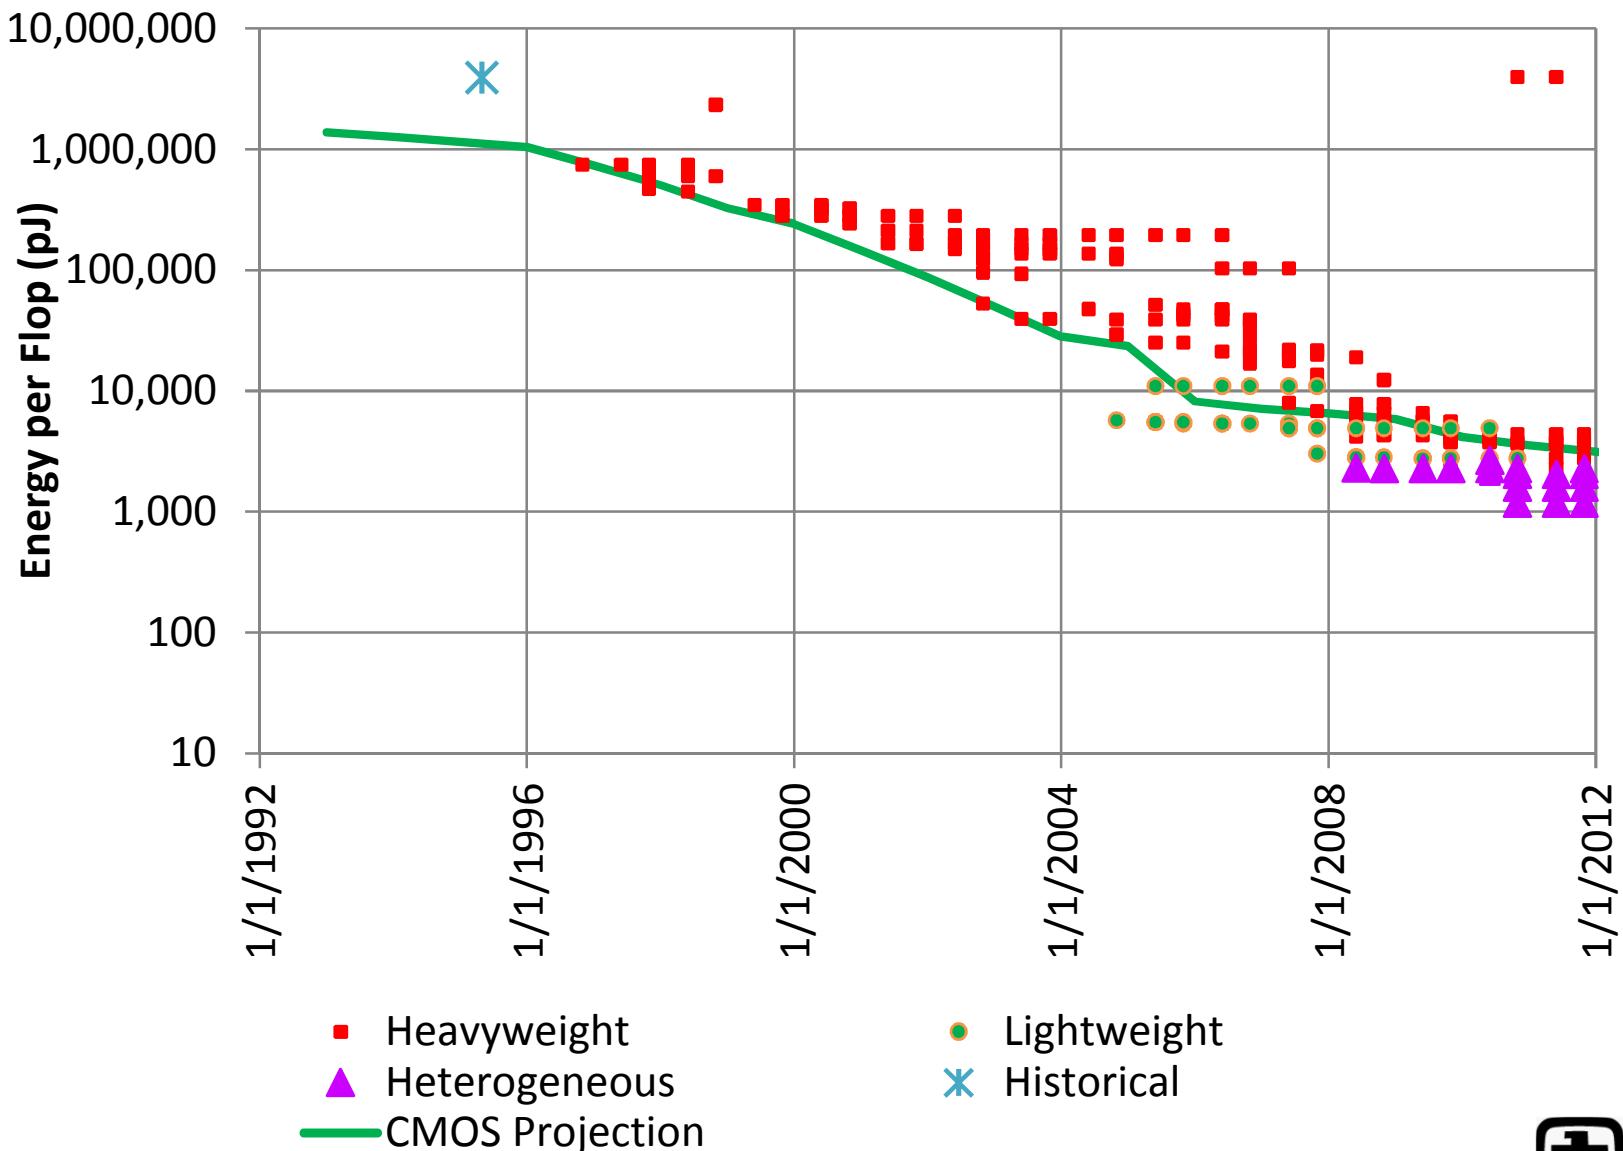
# Power

- K computer
  - Power: 13 MW
- Tianhe
  - Power: 4 MW
- Roadrunner
  - Power: 7 MW → enough to power 5000 homes
- Palo Verde Nuclear Generating Station
  - Power: 3 GW
- Typical Coal Fired Power Plant
  - Power: 500 MW
- $1 \text{ MW} = \$1,000,000/\text{year power bill}$
- $X \text{ pJ per operation} = X \text{ MW per } 10^{18} \text{ operations/sec (Exaflop)}$

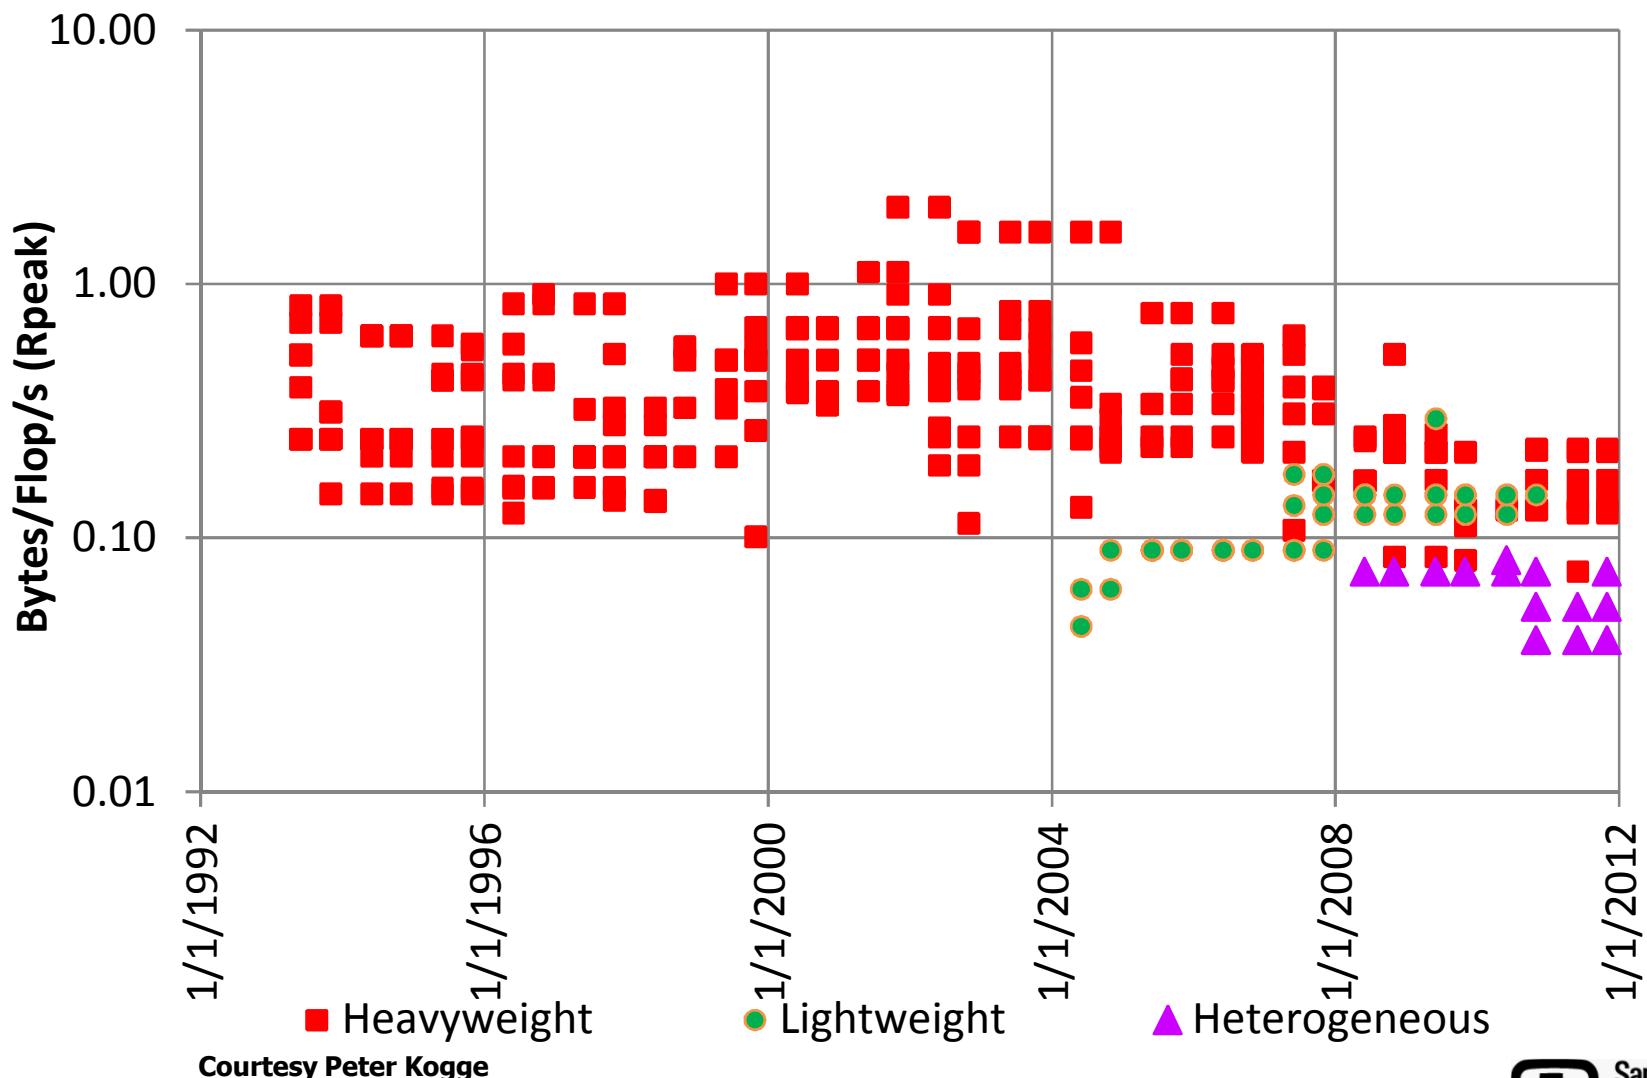


**Will Exascale need dedicated Nuclear Power Plant?**

# Energy per Flop



# DRAM Bytes per Flop



# Common Requirements

- Space and supercomputing stand to benefit from commercial progress in emerging NVMs:
  - Low power
  - Fast read/write
  - High endurance
  - High density
  - Long retention
  - Non-volatility
- Resiliency and fault-tolerance
- HPC and space benefit from radiation hard
  - SEU *is* a problem for supercomputers

# Array Architecture

- How do we architect ReRAM as a main memory array?
- What new issues will we face when converting from DRAM array → ReRAM
- This process has been started for PCM
  - Example – PCM architecture and write scheme below
- Do we need wear leveling?
- Work needed for ReRAM (can learn from PCM techniques)

