IBM Research

Storage Class Memory

Towards a disruptively low-cost solid-state non-volatile memory

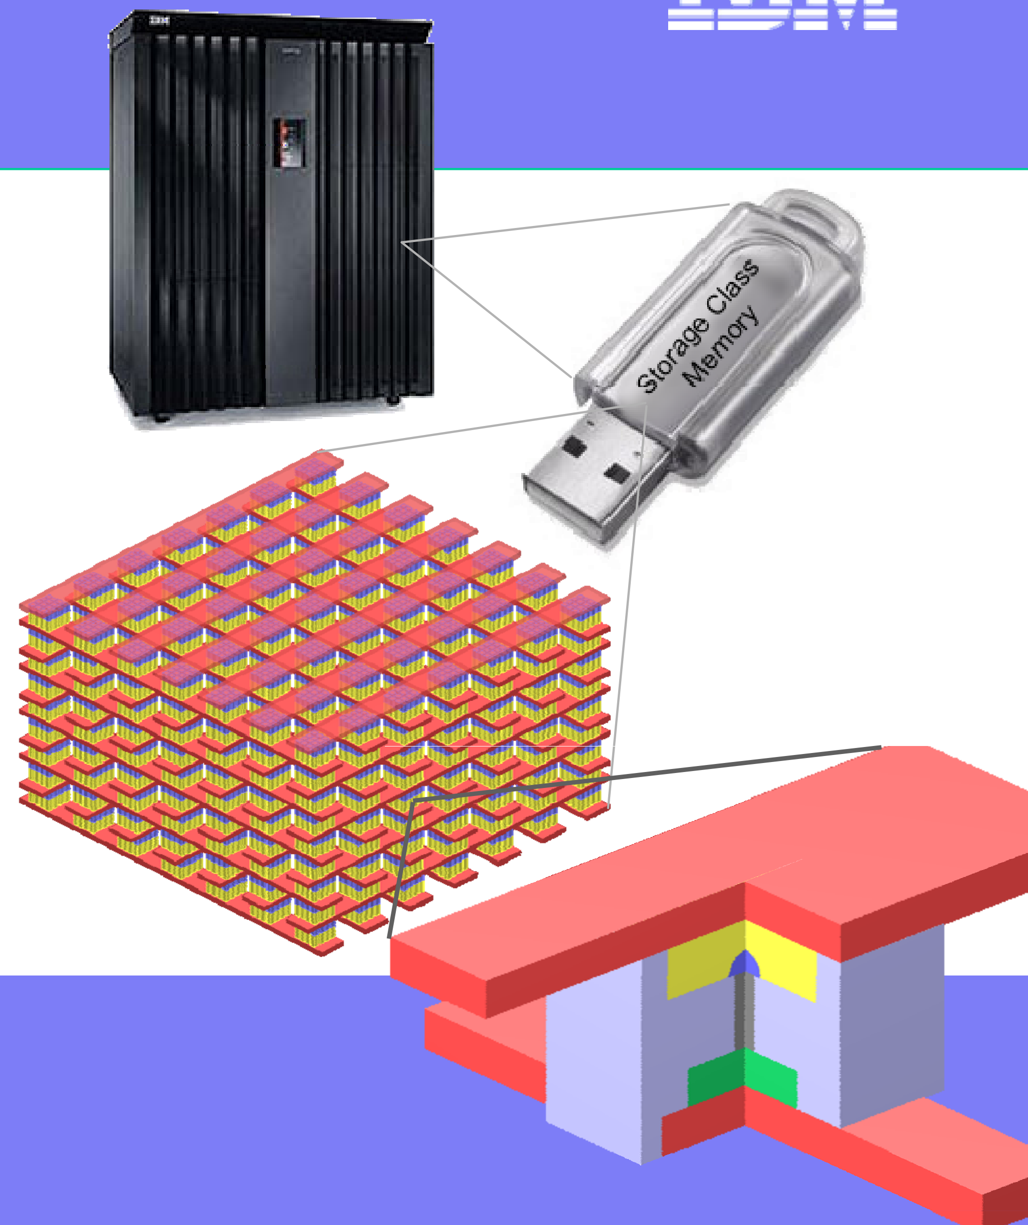

Geoffrey W. Burr

IBM Almaden Research Center

June 15, 2011

Power & space in the server room

The cache/memory/storage hierarchy is rapidly becoming the **bottleneck for large systems**.

We know how to create MIPS & MFLOPS cheaply and in abundance,
but **feeding them with data** has become
the performance-limiting *and* most-expensive part of a system (in **both $ and Watts**).

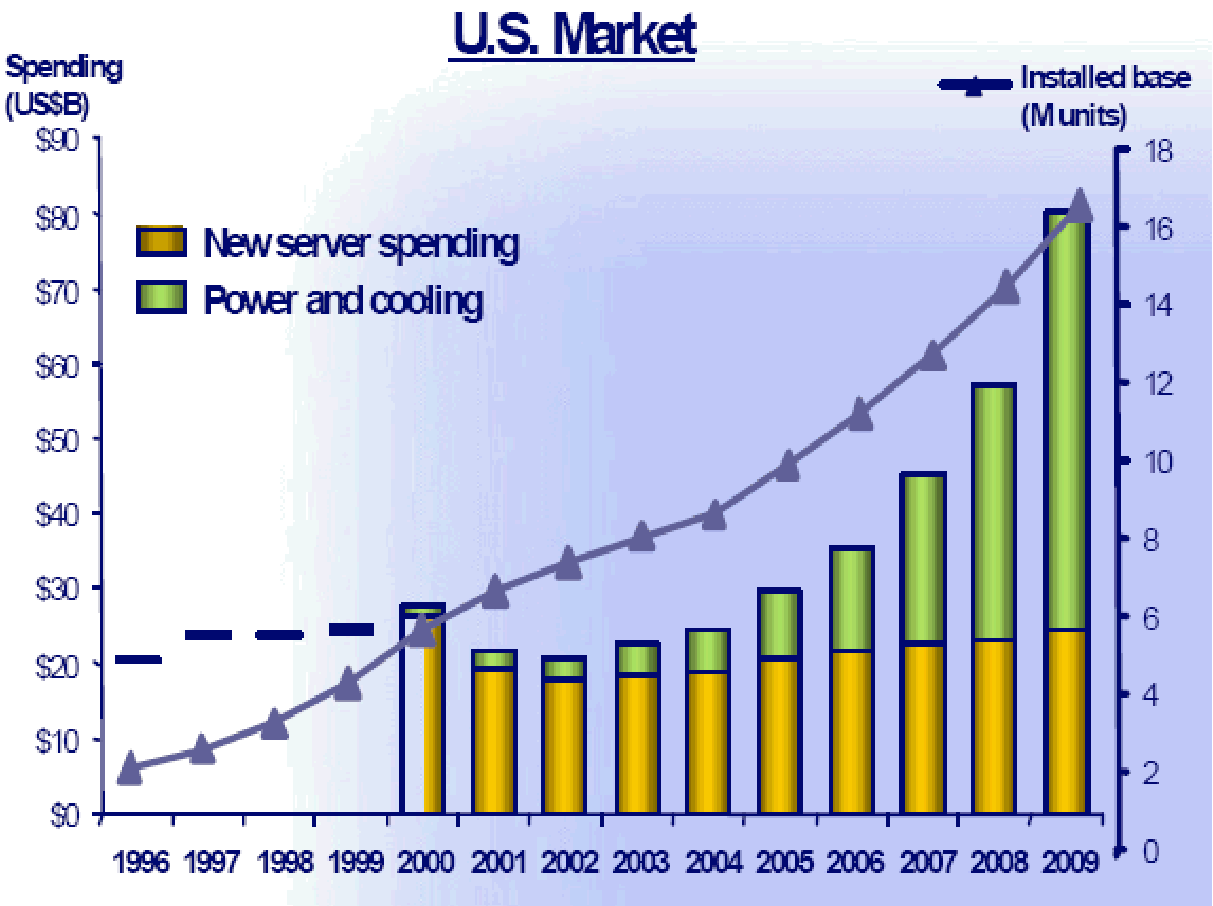

Source IDC: 2006, Document # 201722, "The Impact Of Power and Cooling On Data Center Infrastructure", John Humphreys, Jed Scaramella

Extrapolation to 2020

(at 70% CGR → need **2 GIOP/sec**)

- **5 million HDD**
  - **16,500** sq. ft. !!
  - **22 Mega**watts

R. Freitas and W. Wilcke, *Storage Class Memory: the next storage system technology* –to appear in "Storage Technologies & Systems" special issue of the IBM Journal of R&D.

…yet critical applications are also undergoing a paradigm shift

	Compute-centric paradigm	**Data-centric paradigm**
Main Focus:	**Solve differential equations**	**Analyze petabytes of data**
Bottleneck:	***CPU / Memory***	***Storage & I/O***
Typical Examples:	Computational Fluid Dynamics Finite Element Analysis Multi-body Simulations	Search and Mining Analyses of social/terrorist networks Sensor network processing Digital media creation/transmission Environmental & economic modeling

Extrapolation to 2020

(at 90% CGR → need **1.7 PB/sec**)

- **5.6 million HDD**
  - **19,000** sq. ft. !!
  - **25 Mega**watts

(at 90% CGR → need **8.4G SIO/sec**)

- **21 million HDD**
  - **70,000** sq. ft. !!
  - **93 Mega**watts

[Freitas:2008]

Problem (& opportunity): The access-time gap between memory & storage

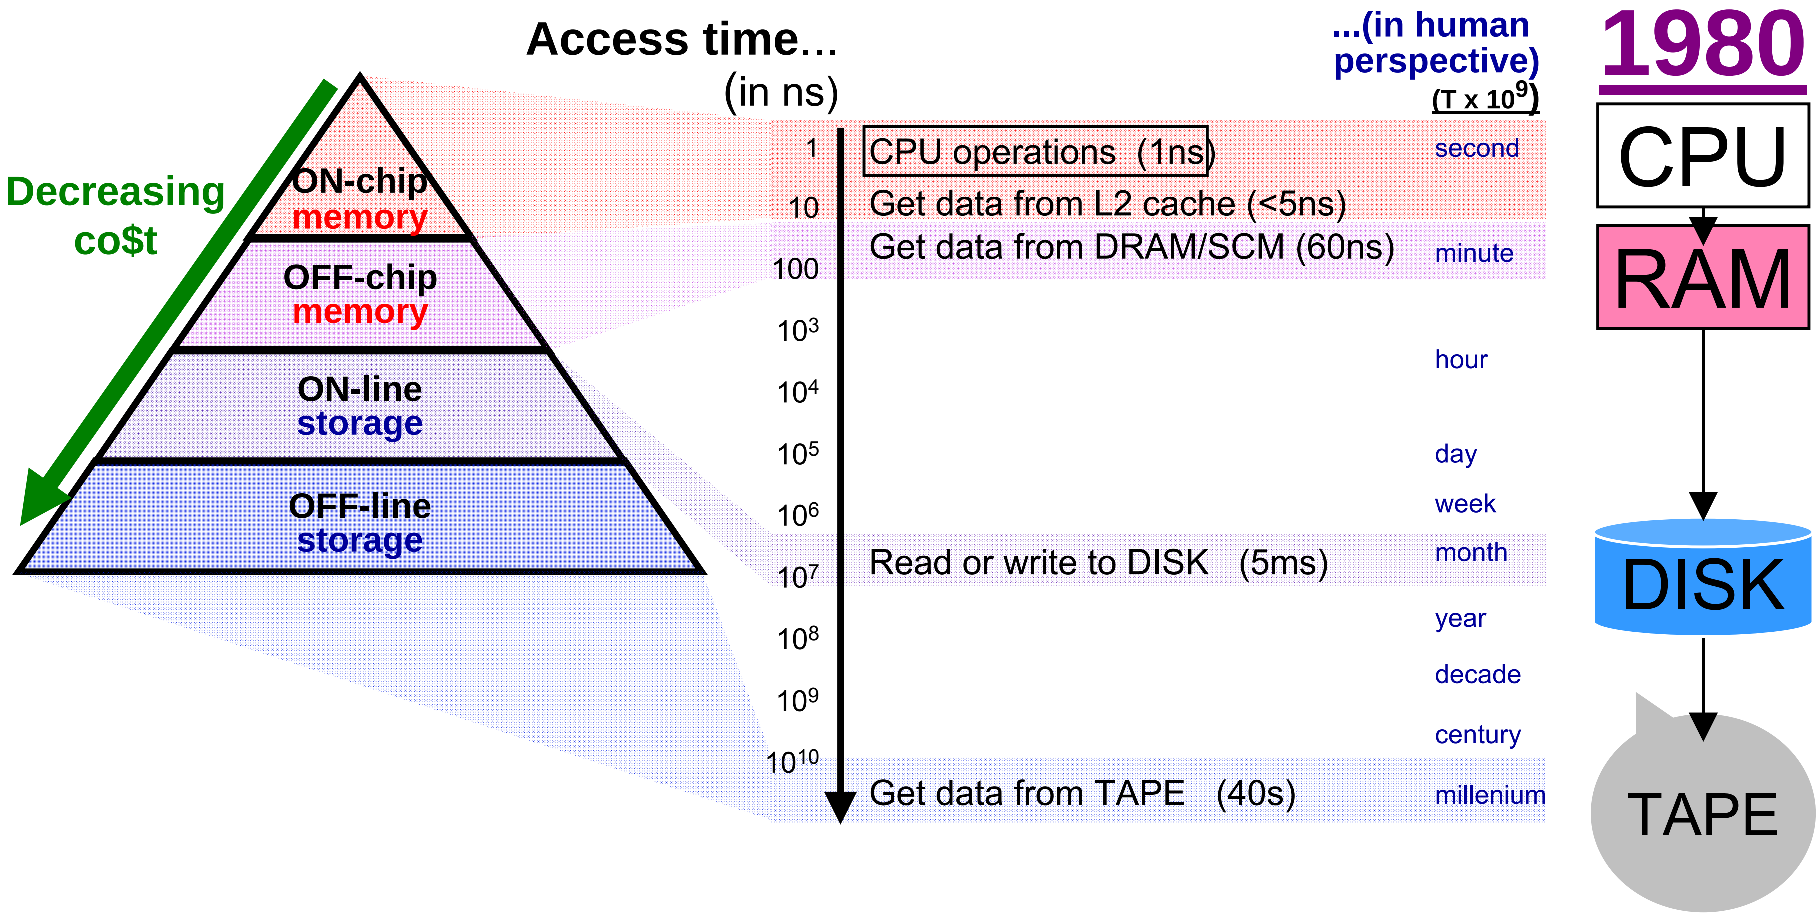

- Modern computer systems have long had to be designed around **hiding the access gap** between **memory** and **storage** → caching, threads, predictive branching, etc.

- "Human perspective" – if a CPU instruction is analogous to a 1-second decision by a human, retrieval of data from off-line tape represents an analogous delay of 1250 years

Problem (& opportunity): The access-time gap between memory & storage

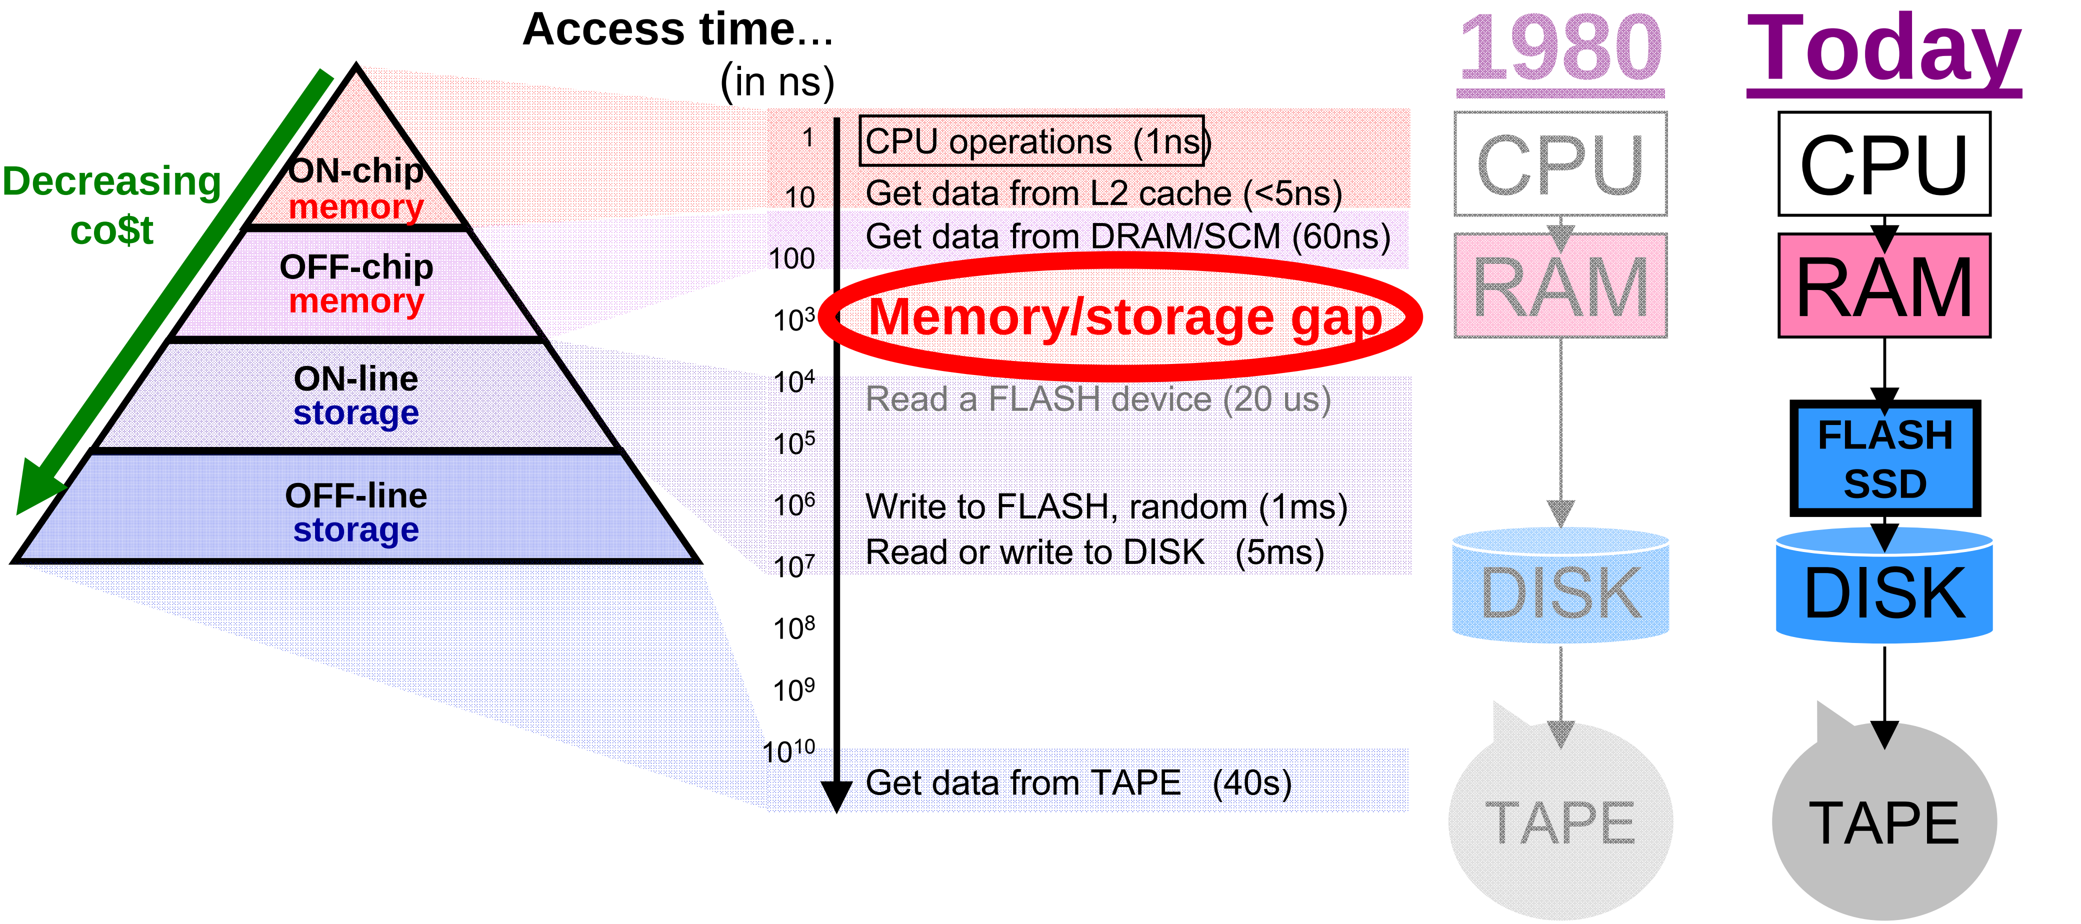

- Today, **Solid-State Disks** based on NAND Flash can offer fast ON-line storage, and storage capacities are increasing as devices scale down to smaller dimensions...

  ...but while prices are dropping, the **performance gap** between memory and storage remains significant, and the already-**poor device endurance** of Flash is getting worse.

Problem (& opportunity): The access-time gap between memory & storage

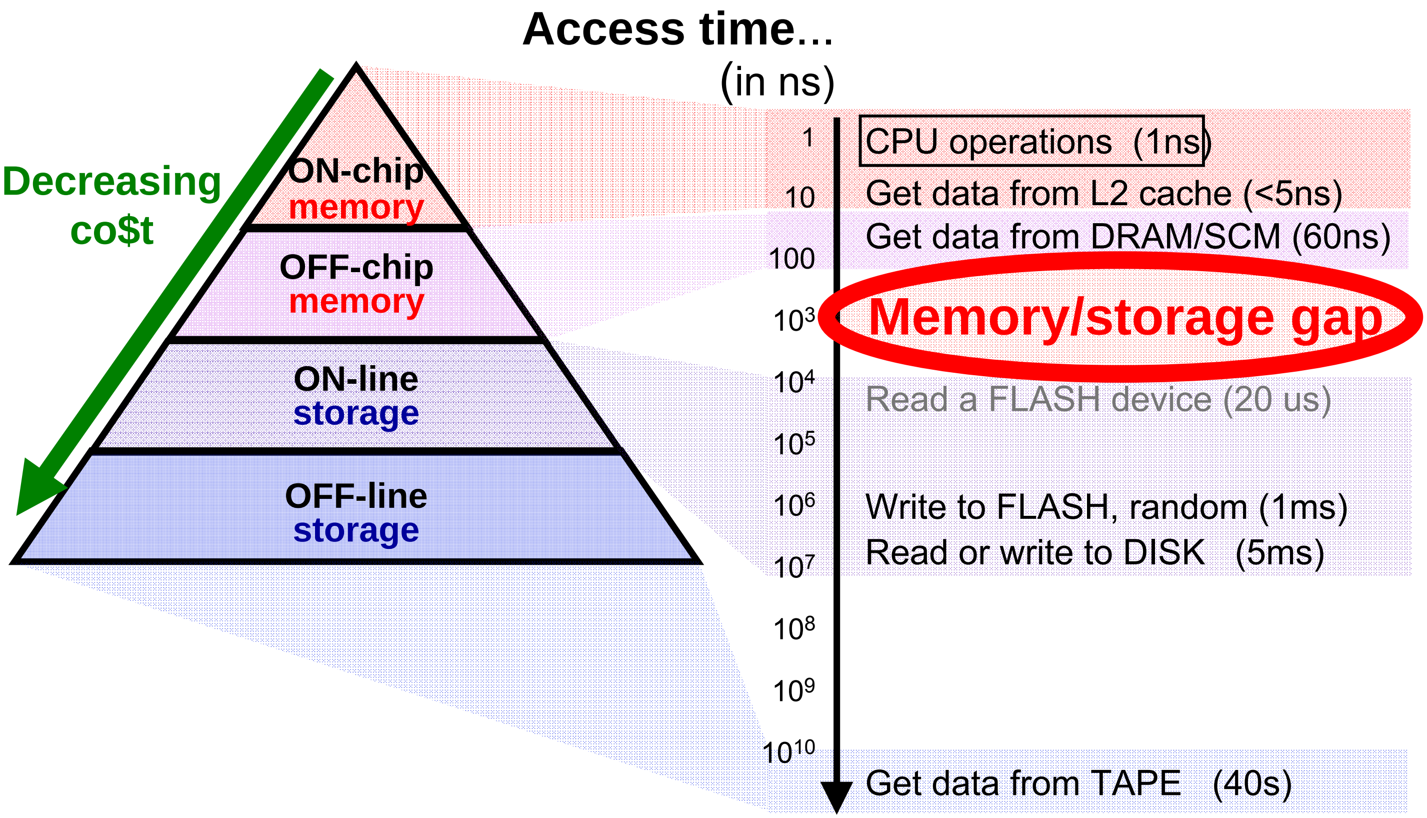

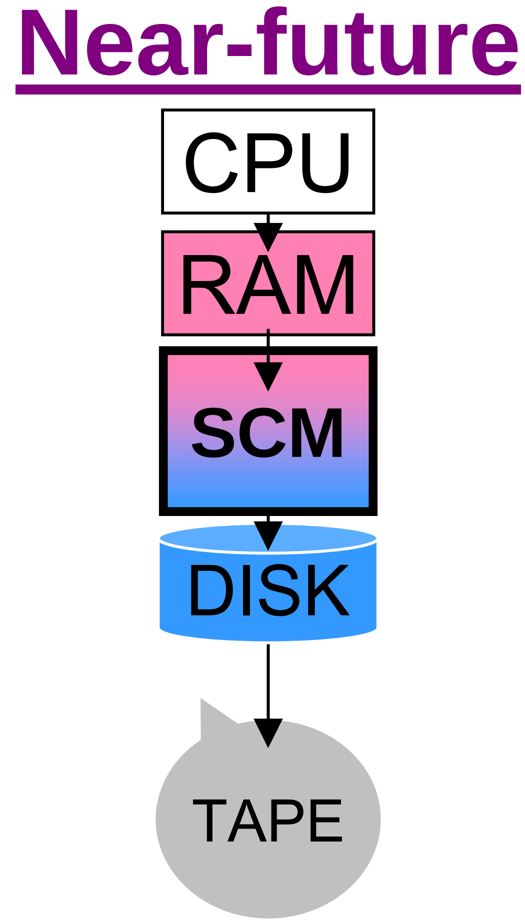

Research into new solid-state non-volatile memory candidates
– originally motivated by finding a "successor" for NAND Flash –
has opened up several interesting ways to change the memory/storage hierarchy...

1) **Embedded** Non-Volatile Memory – low-density, fast ON-chip NVM
2) **Embedded** Storage – low density, slower ON-chip storage
3) M-type **Storage Class Memory** – **high-density**, fast OFF- (or ON*)-chip NVM
4) S-type **Storage Class Memory** – **high-density**, very-near-ON-line storage

* ON-chip using 3-D packaging

Storage-type vs. memory-type Storage Class Memory

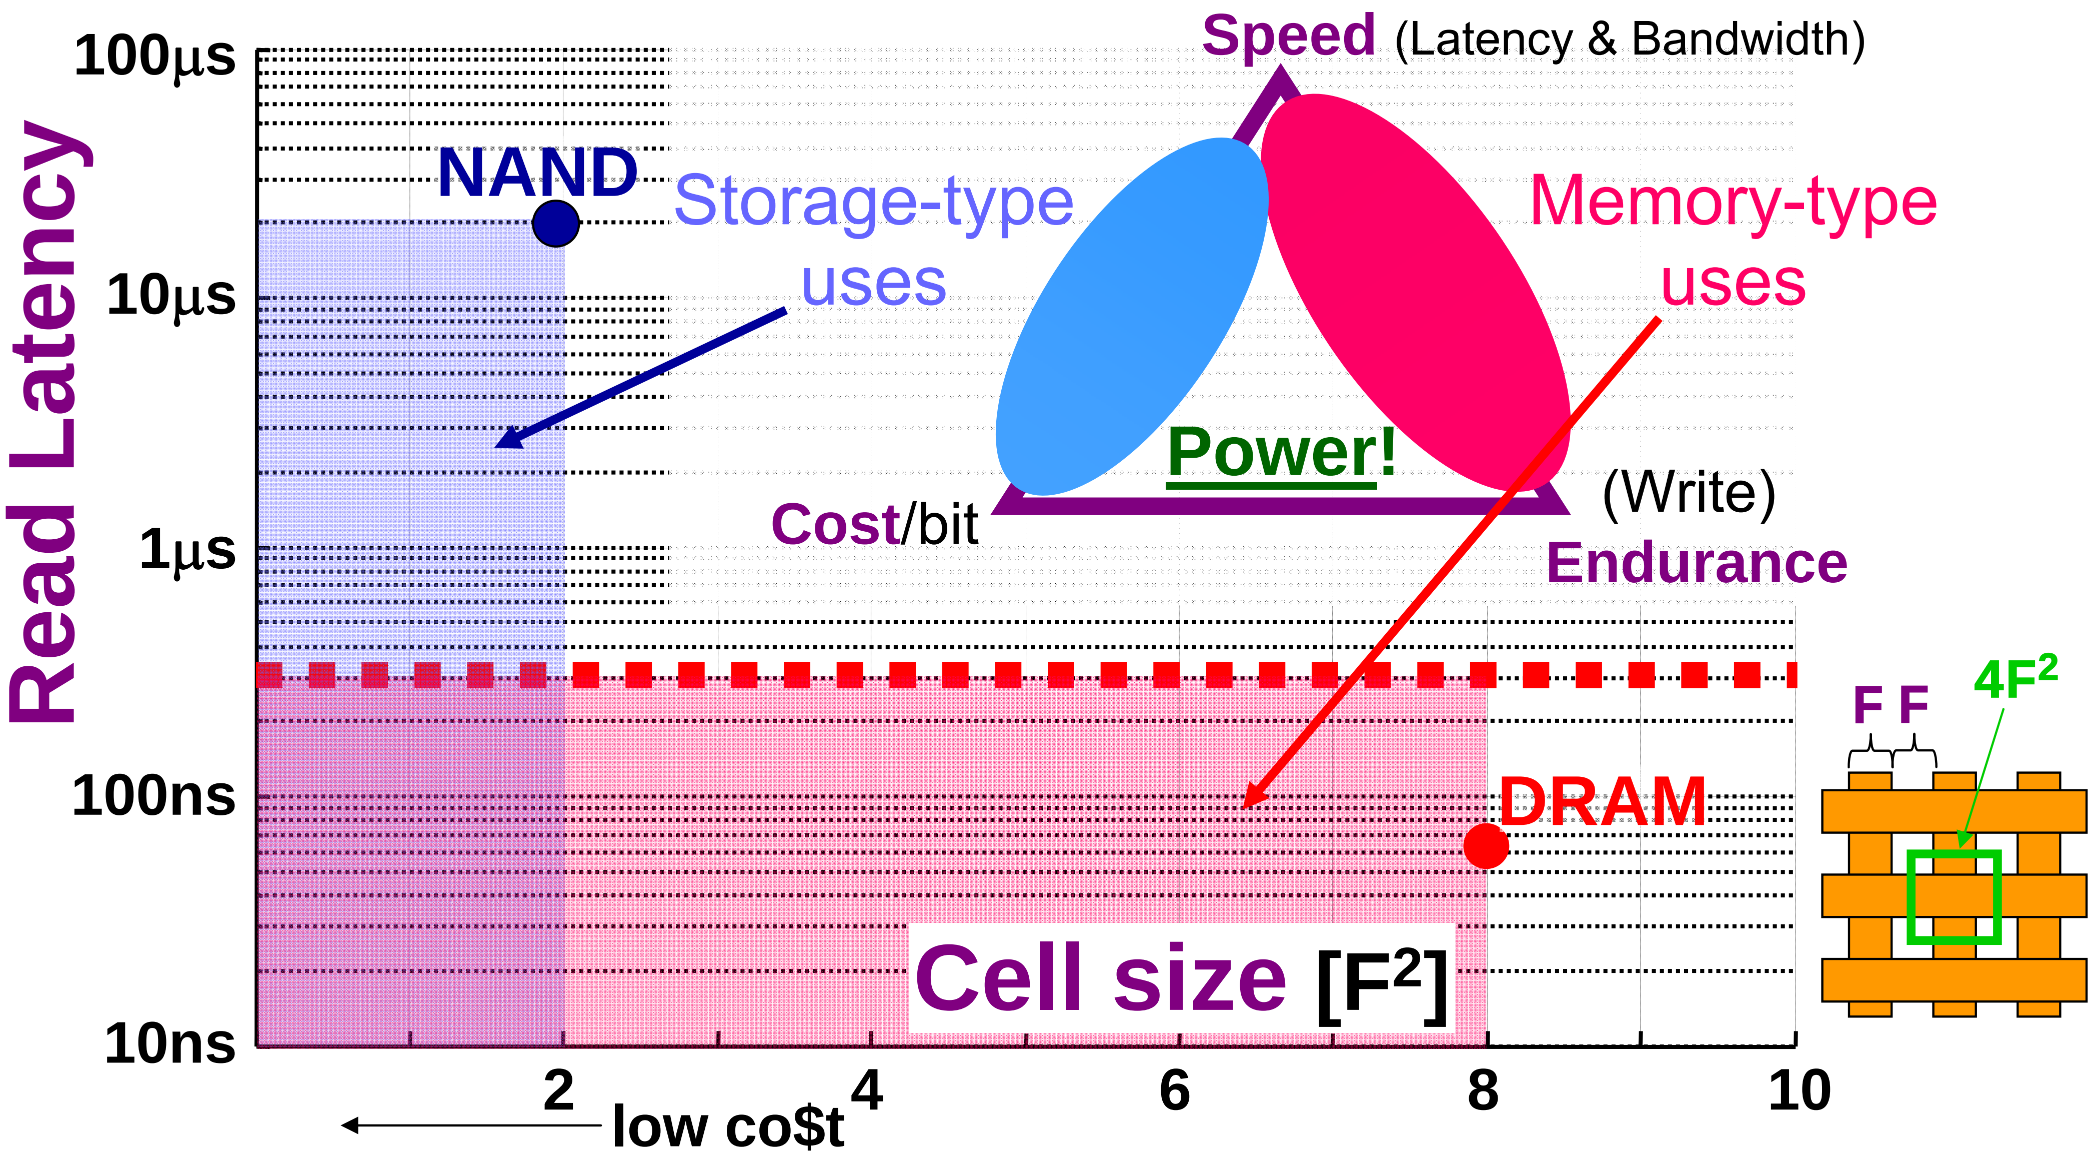

The cost basis of semiconductor processing is well understood – the paths to higher density are 1) shrinking the minimum lithographic pitch **F**, and 2) storing **more bits PER 4F²**

S-type vs. M-type SCM

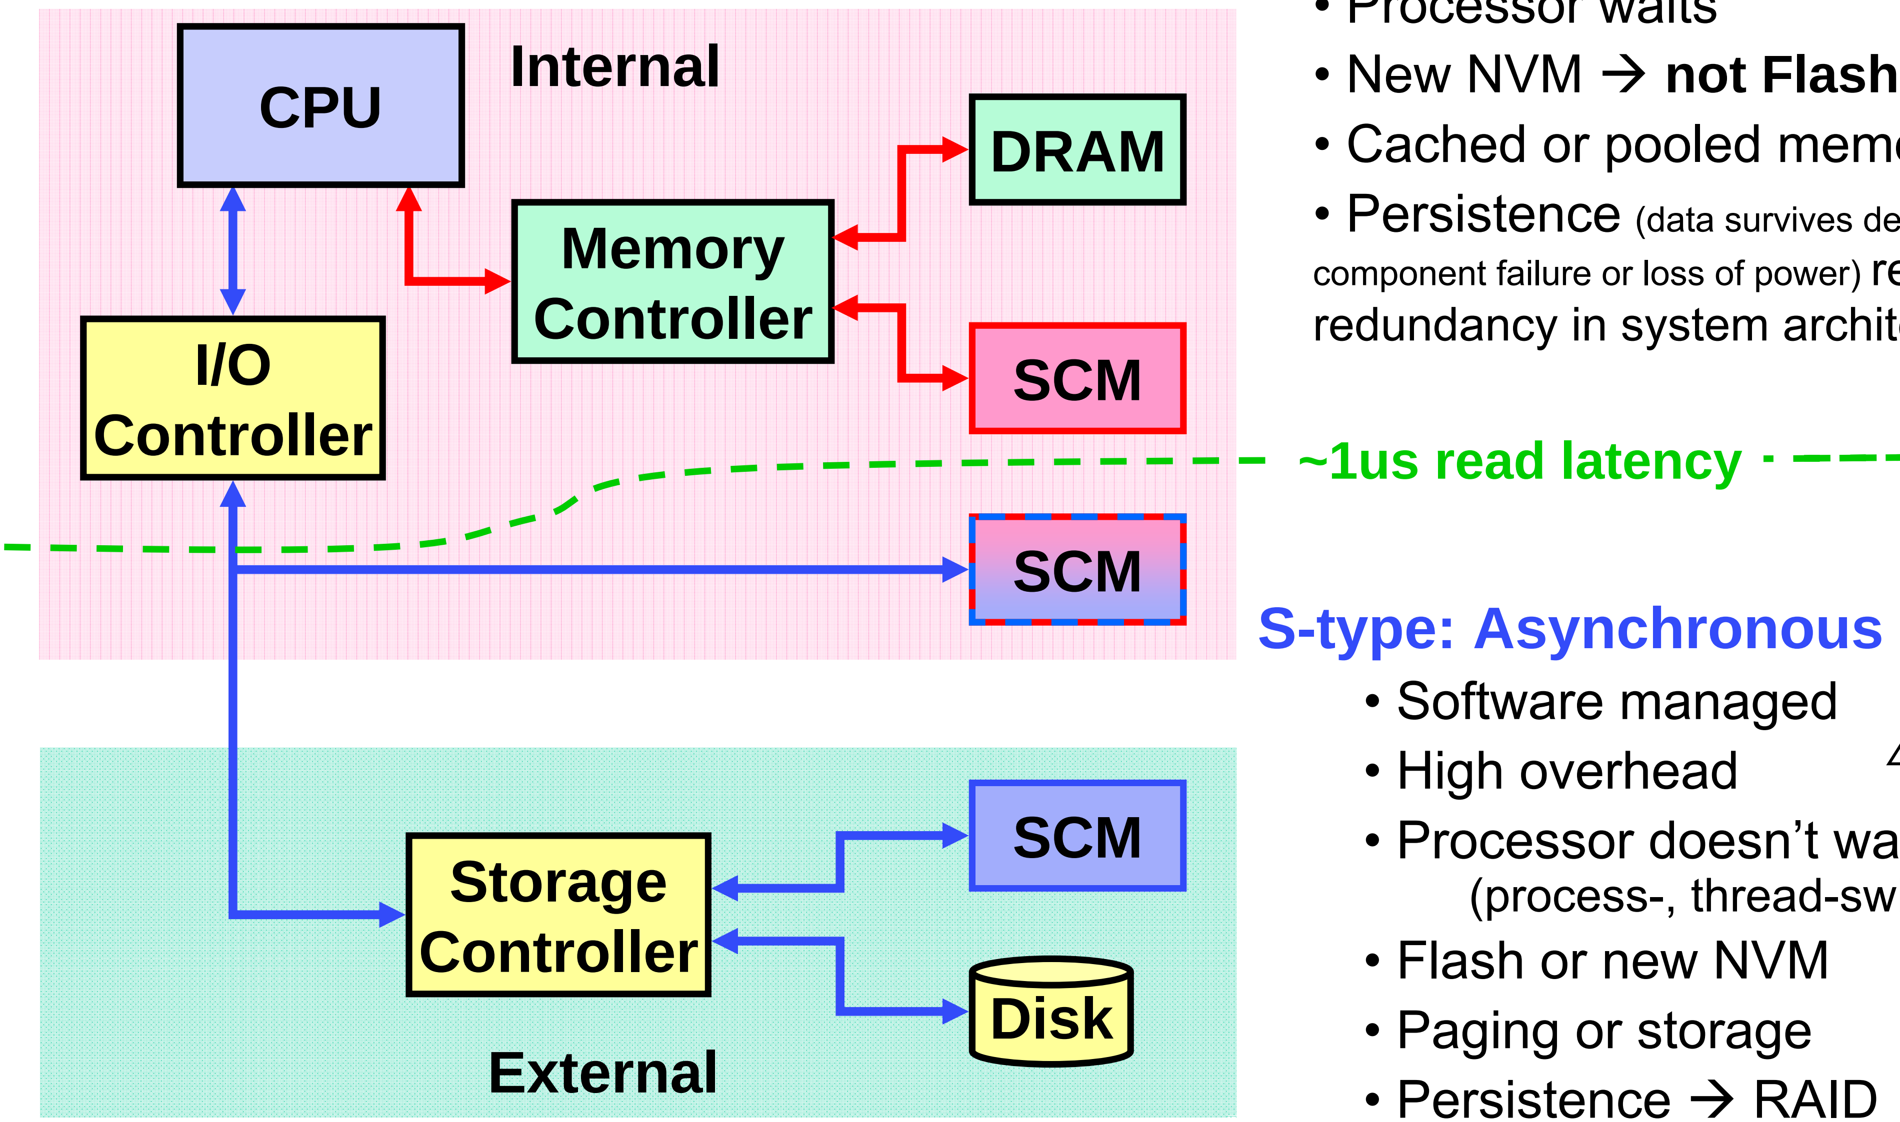

M-type: Synchronous

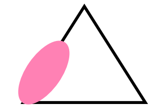

- Hardware managed
- Low overhead
- Processor waits
- New NVM → **not Flash**
- Cached or pooled memory
- Persistence (data survives despite component failure or loss of power) requires redundancy in system architecture

S-type: Asynchronous

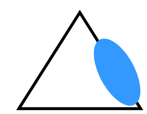

- Software managed
- High overhead
- Processor doesn't wait, (process-, thread-switching)
- Flash or new NVM
- Paging or storage
- Persistence → RAID

NVM candidates for SCM

1) NVM element

- **Improved FLASH**
- **Phase Change RAM**
- **Magnetic Spin Torque Transfer**
  - → **STT-RAM**
  - → **Magnetic Racetrack**
- **Resistive RAM**

2) High-density access device (A.D.)

- **2-D** – silicon transistor or diode
- **3-D** → **higher density per 4F^2**
  - polysilicon diode (but <400^{o}C processing?)
  - MIEC A.D. (Mixed Ionic-Electronic Conduction)
  - OTS A.D. (Ovonic Threshold Switch)
  - Conductive oxide tunnel barrier A.D.

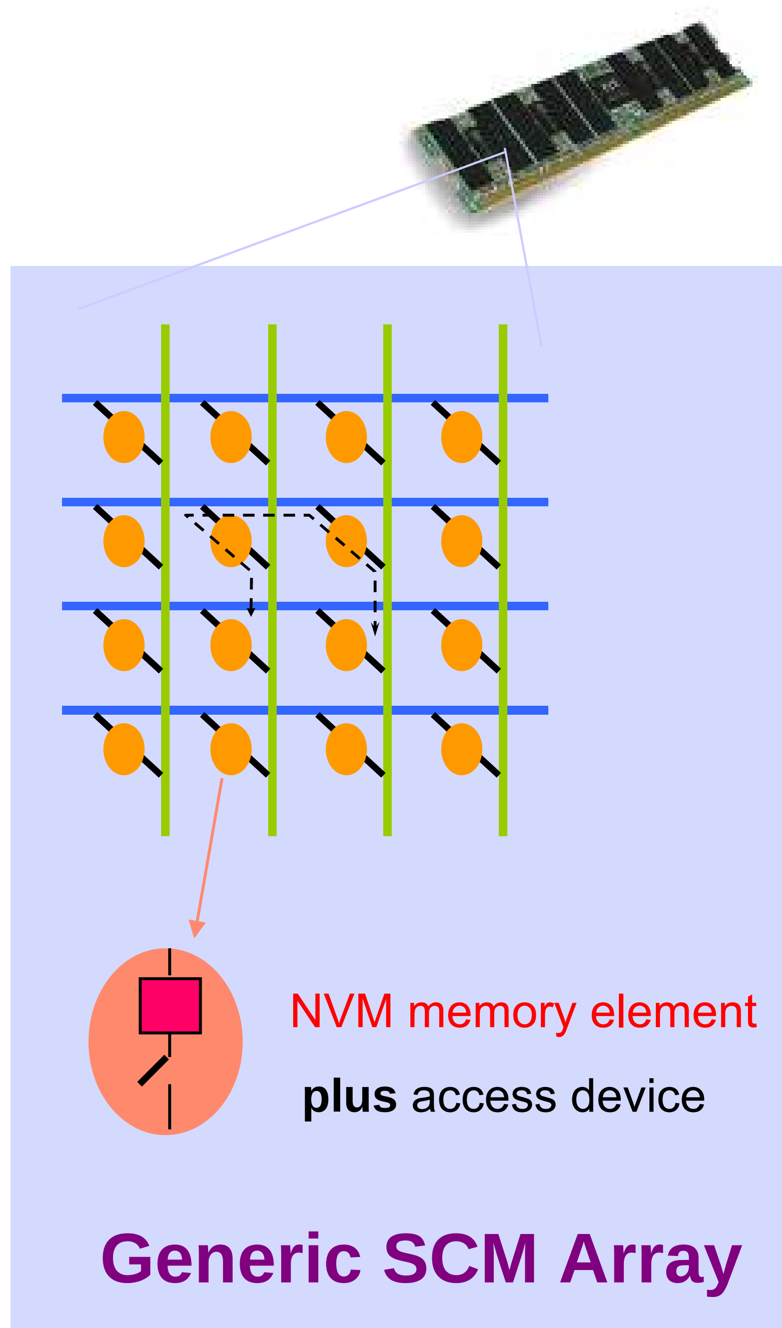

Limitations of Flash

Asymmetric performance

Writes much slower than reads

Program/erase cycle

Block-based, no write-in-place

Data retention and Non-volatility

Retention gets worse as Flash scales down

Endurance

- Single level cell (SLC) → 10^5 writes/cell
- Multi level cell (MLC) → 10^4 writes/cell
- Triple level cell (TLC) → ~300 writes/cell

Future outlook

- Scaling focussed solely on density
- 3-D schemes exist but are complex

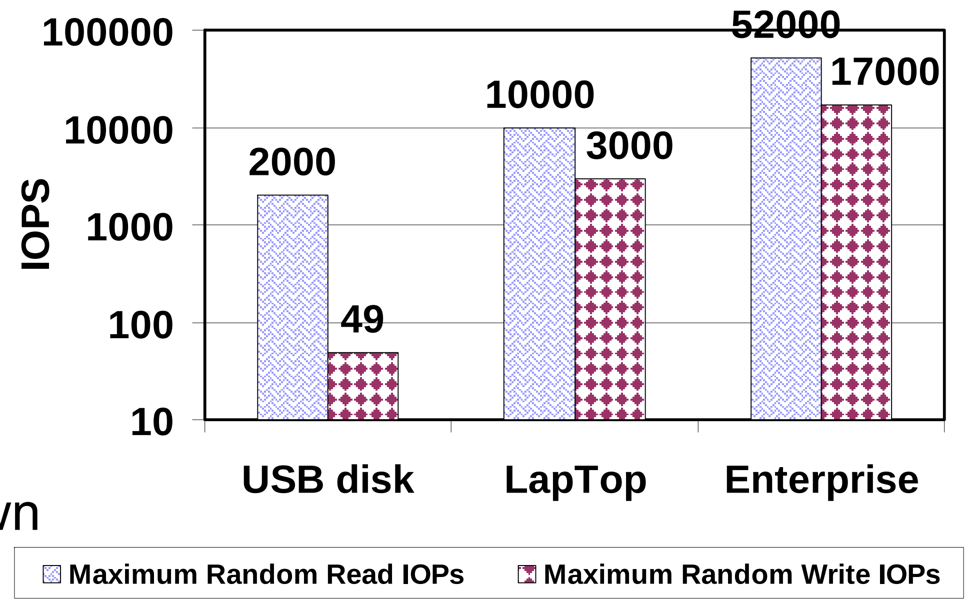

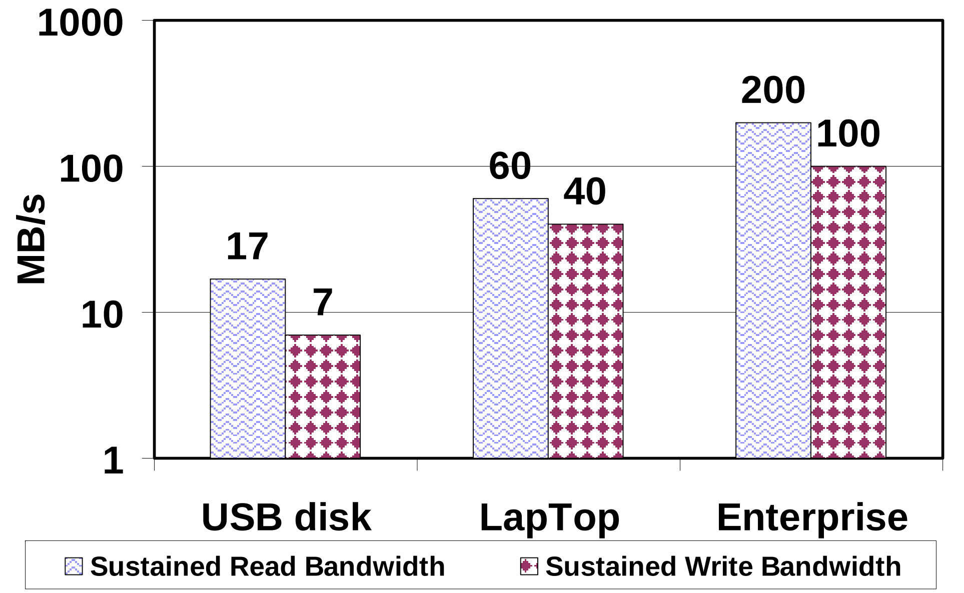

STT (Spin-Torque-Transfer) RAM

- Controlled switching of free magnetic layer in a magnetic tunnel junction using current, leading to two distinct resistance states

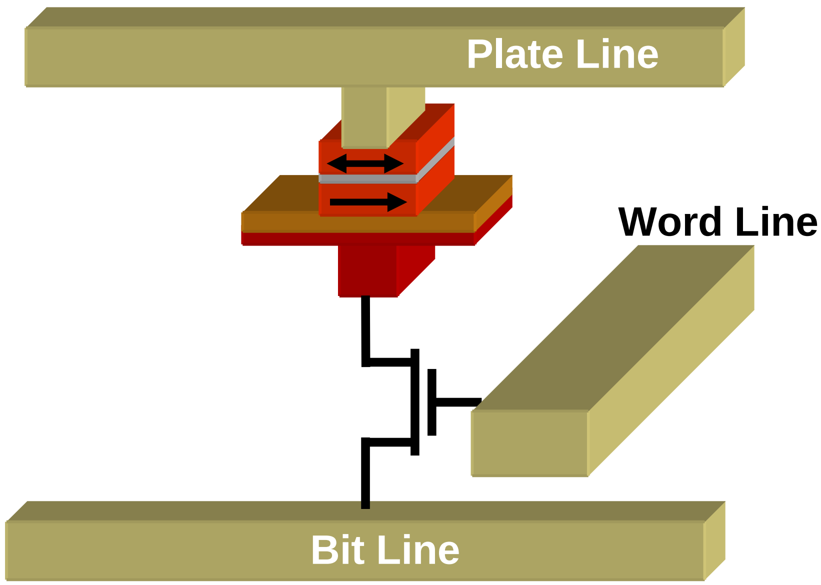

Strengths

- Inherently very fast → **almost as fast as DRAM**
- **Much better endurance** than Flash or PCM
- Radiation-tolerant
- Materials are Back-End-Of-the-Line compatible
- Simple cell structure → reduced processing costs

Weaknesses

- Achieving low switching current/power is not easy
- Resistance contrast is quite low (2-3x) → achieving **tight distributions** is ultra-critical
- High-temperature retention **strongly affected by scaling** below F~50nm
- Tradeoff between fastest switching and switching reliability

Outlook:

Strong outlook for an Embedded Non-Volatile Memory to replace/augment DRAM.

While near-term prospects for high-density SCM with STT-RAM may seem dim, **Racetrack Memory** offers hope for using STT concepts to create vertical "shift-register" of domain walls → potential densities of 10-100 bits/F^2

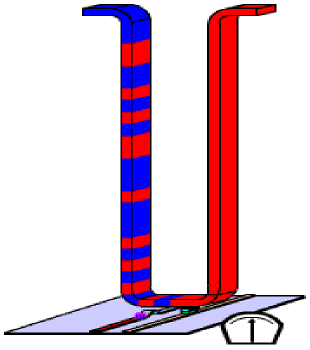

Resistive RAM

Voltage-controlled formation & dissipation of an oxygen-vacancy (or metallic) filament through an otherwise insulating layer

Strengths

- Good retention at elevated-temperatures
- Simple cell structure → reduced processing costs
- Both fast and ultra-low-current switching have been demonstrated
- Some RRAM materials are Back-End-Of-the-Line compatible
- Relatively new field → high hopes for improved material concepts
- Less "gating" Intellectual Property to license
- Some RRAM concepts offer co-integrated NVM & Access Device
- Numerous ongoing development efforts

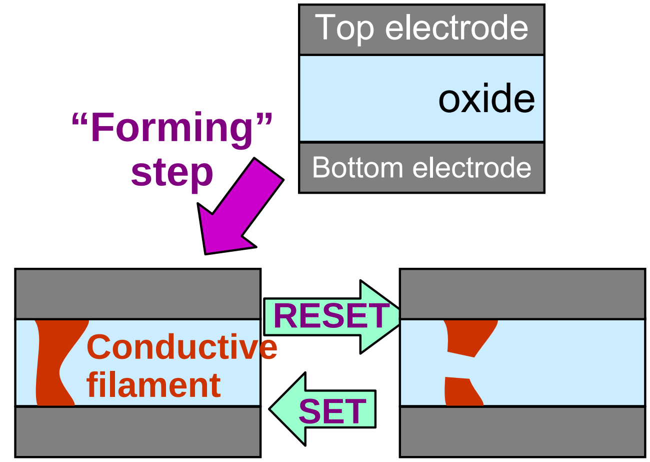

Weaknesses

- Highly immature technology – wide variation in materials hampers cross-industry learning
- Demonstrated endurance is slightly better than Flash, but lower than PCM or STT-RAM
- **Switching reliability** an issue, even within single devices, and read disturb can be an issue
- An initial high-voltage **"forming"** step is often required
- To attain low RESET switching currents, circuit must constrain current during previous SET
- Unipolar and bipolar versions – bipolar typically better in both write margins & endurance, but then requires an unconventional bipolar-capable Access Device (transistor or diode is out)
- High array yield with minimal "outlier" devices **not yet demonstrated**
- Tradeoff between switching speed, long-term retention, and reliability **not yet explored**

Outlook: Outlook is unclear. Emergence of a strong material candidate offering high array yield & reliability could focus industry efforts considerably. Absent that, many uncertainties remain about prospects for reliable storage & memory products.

Phase-change RAM

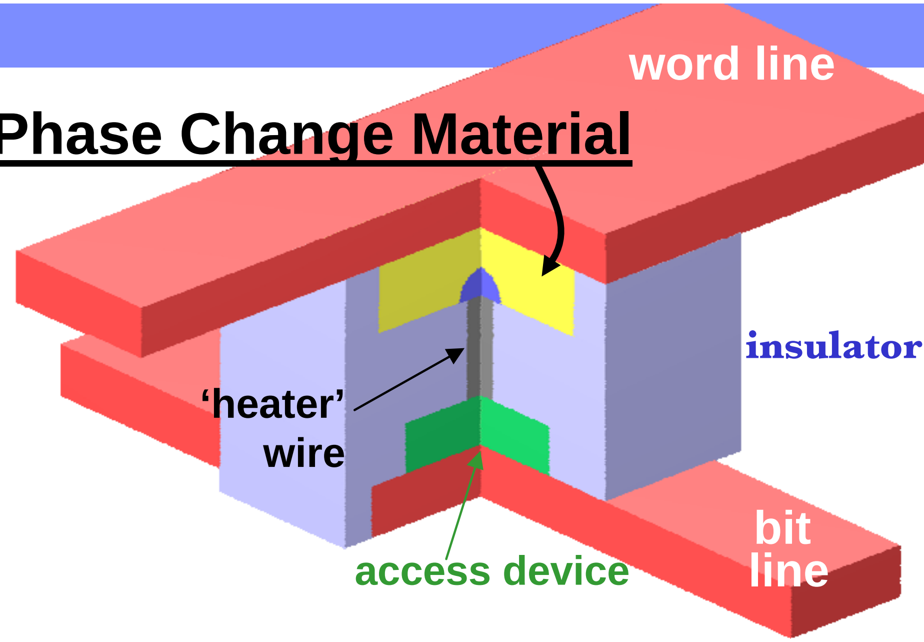

• Switching between **low-resistance** crystalline, and **high-resistance** amorphous phases, controlled through power & duration of electrical pulses

Strengths

- **Very mature** (large-scale demos & products)
- Industry consensus on material → GeSbTe or GST
- Large resistance contrast → analog states for **MLC**
- Offers **much better endurance** than Flash
- Shown to be **highly scalable** (still works at ultra-small **F**) and Back-End-Of-the-Line compatible
- Can be very fast (depending on material & doping)

Weaknesses

- RESET step to high resistance requires melting → **power-hungry**, thermal crosstalk?
  - To keep switching power down → sub-lithographic feature and high-current Access Device
    - To fill small feature → ALD or CVD → difficult now to replace GST with a better material
      - Variability in small features broadens resistance distributions
- 10-year **retention at elevated temperatures** can be an issue → recrystallization
- Device characteristics change over time due to elemental segregation → device failure
- **MLC** strongly affected by relaxation of amorphous phase → "resistance drift"

Outlook: NOR-replacement products now shipping → if yield-learning successful and MLC drift-mitigation and/or 3-D Access Devices can offer high-density (=low-cost), then opportunity for NAND replacement, S-type, and then finally M-type SCM may follow

PCM @ IBM Almaden

- **Prototype memory devices**

3nm * 20nm → 60nm^2

≈ Flash roadmap for **2013**

→ **phase-change *scales***

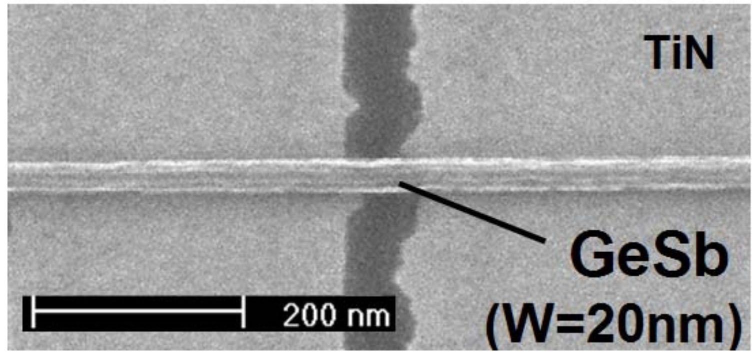

[Chen IEDM 2006]

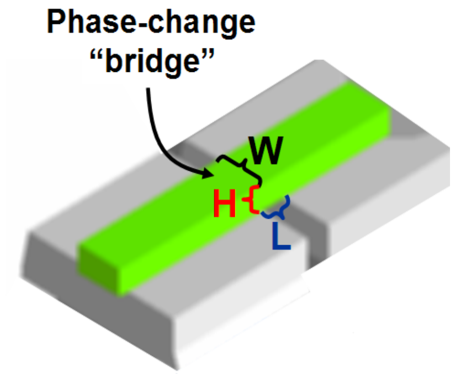

W defined by **lithography**

H by **thin-film deposition**

- **Phase-change materials work**

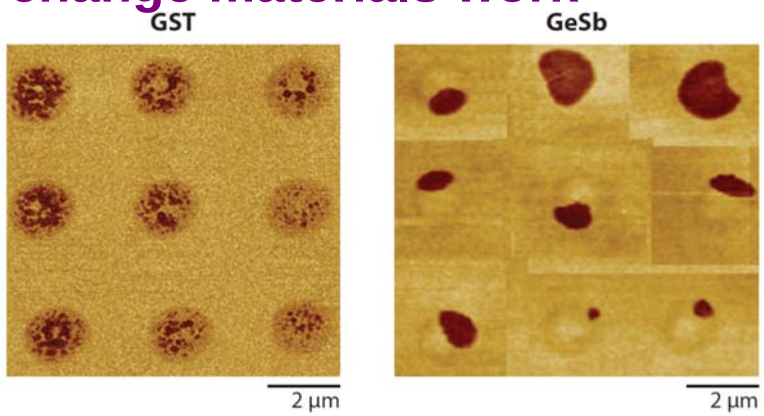

- **Improvements to PCM**
  - Drift of intermediate resistances
  - Interplay between stoichiometry changes & endurance (+ other effects)

- **Device & materials modeling**

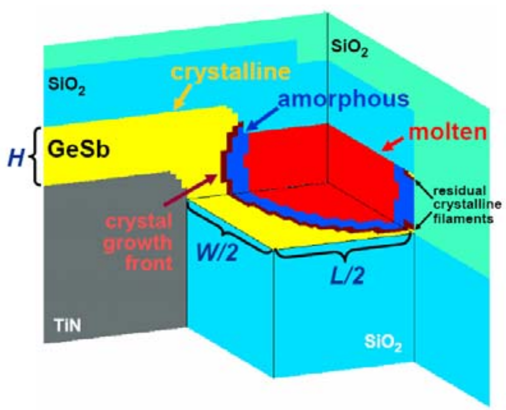

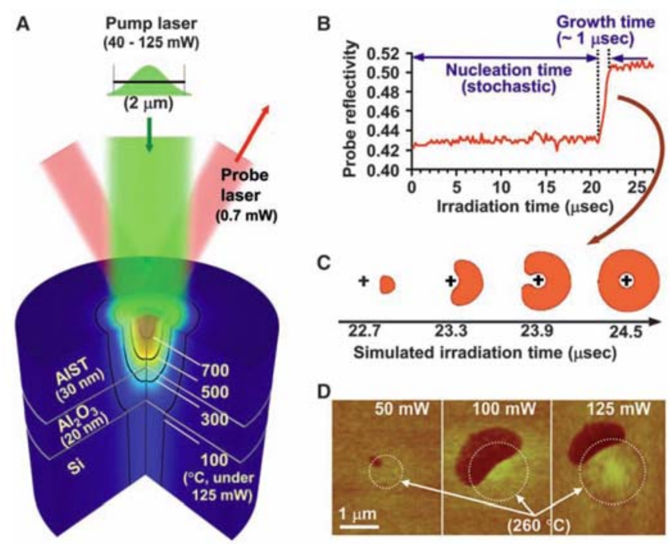

[B. S. Lee Science 2009]

PCM polarity: electromigration + crystallization-induced segregation

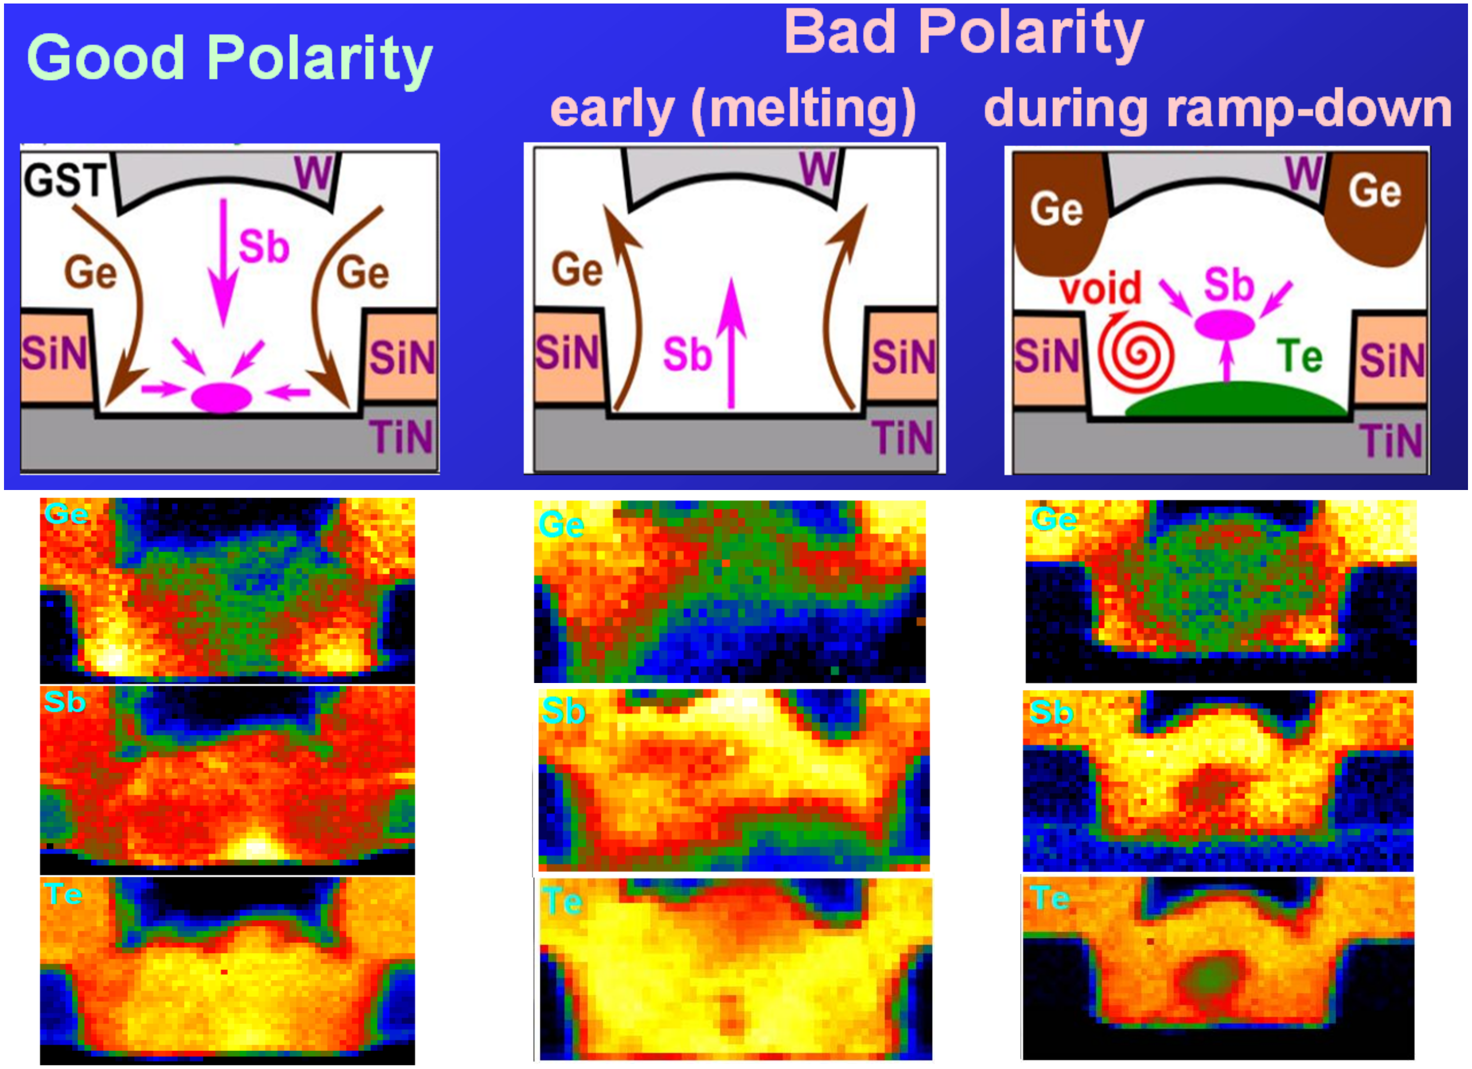

- **Better understanding of…**
  - how PCM devices fail
  - how to improve PCM endurance through better materials science, cell design, cell programming, and creative use of the "bad" polarity

Access Devices for 3-D memory in the Back-End-Of-the-Line

	Strengths	Weaknesses
Silicon transistor or diode	• Well understood – diode offers higher current in a small footprint	• Not an option for 3-D BEOL
Polysilicon diode	• Well understood, with small footprint	• High ON current & low OFF current require processing at temperatures too high for BEOL
Conductive oxide tunnel barrier	• Demonstrated by Oxide Semiconductor (now Micron)	• Critically dependent on thickness • Poor ON-OFF contrast
OTS (Ovonic Threshold Switch)	• Essentially a permanently-amorphous PCM cell, undergoes electrical breakdown • Similar to PCM materials	• OFF state leakage not specified • Poor control over read operation – margin & non-destructive reads?
MIEC device (Mixed Ionic-Electronic Conduction)	• High enough ON currents for PCM • Low enough OFF current for large arrays	• Large arrays would require even higher voltage margin • Novel, Cu-containing materials survive BEOL with high yield?

Novel Mixed-Ionic-Electronic-Conduction Access Device

Strengths

- **High** enough **ON currents** for PCM – cycling of PCM has been demonstrated
- **Low** enough **OFF current** for large arrays
- Very large (>>1e10) endurance for typical 5uA read currents
- Voltage margins > 1.0V with tight distributions
- CMP process demonstrated

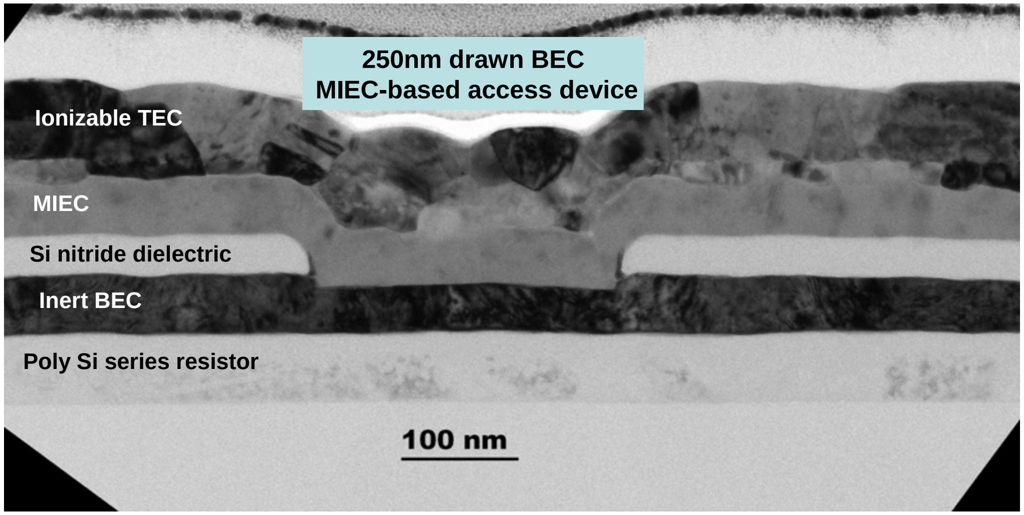

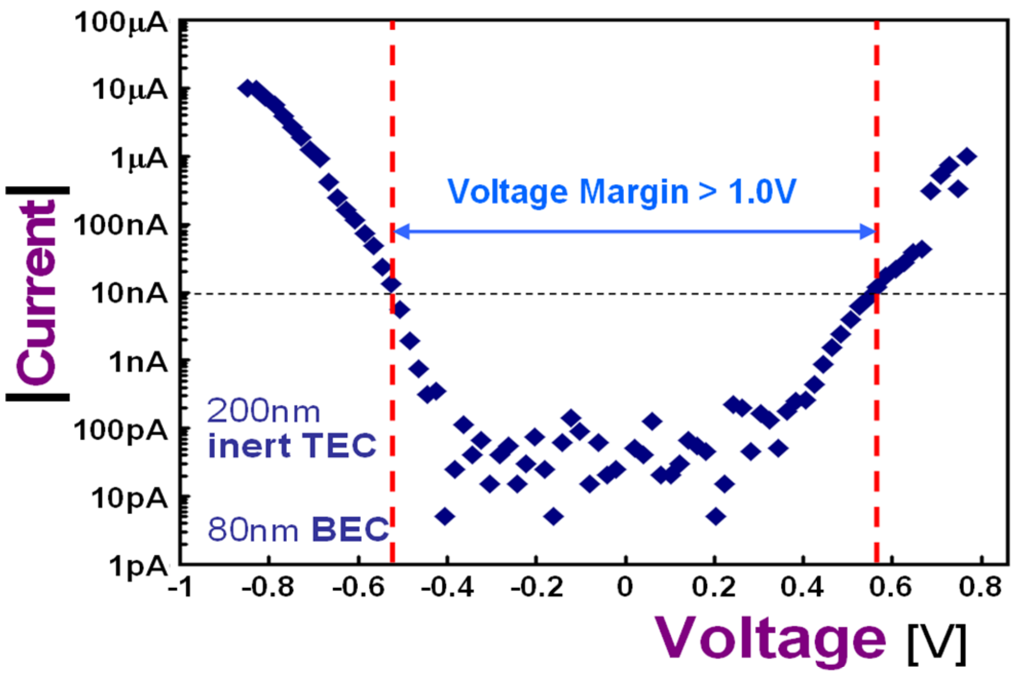

Gopalakrishnan, VLSI 2010

Weaknesses

- Large arrays will likely require voltage margin >1.5V
- Do these novel, Cu-containing materials survive BEOL with **high yield**?
- Maximum voltage across companion NVM during switching must be low (~1V) → influences half-select condition and thus achievable array size
- **Endurance during NVM programming** is strongly dependent on programming current

Competitive Outlook among emerging NVMs

Future NOR applications
(program code, etc.)
1) **PCM** (but market disappearing)

Future NAND applications
(consumer devices, etc.)
1) **3-D NAND** (but crossover to succeed 20nm conventional NAND may require >50 layers!)
2) **PCM / RRAM**

Embedded Storage
(low density, slower ON-chip storage)
1) **NAND** (but complicated process)
2) **RRAM / PCM**

S-type **Storage Class Memory**
(**high-density**, very-near-ON-line storage)
1) **PCM / RRAM**
2) **Racetrack** (future)

M-type **Storage Class Memory**
(**high-density**, fast OFF- (or ON*)-chip NVM)
1) **PCM / RRAM**
2) **Racetrack** (future)

Embedded Non-Volatile Memory
(low-density, fast ON-chip NVM)
1) **STT-RAM**

High Speed (↓)

Low co$t (→)

* ON-chip using 3-D packaging

Industry SCM activities

- Numonyx (FLASH & PCM)
  … acquired by Micron
  … shipping 128Mb **PCM chip** with 10^6 write cycles.
- Samsung shipping 512Mb **PCM chip** for smart phone applications.
- Over 30 companies work on SCM including all major IT players
- Numerous **RRAM** efforts at NAND Flash companies
- Numerous **STT-RAM** efforts, including Qualcomm & Samsung

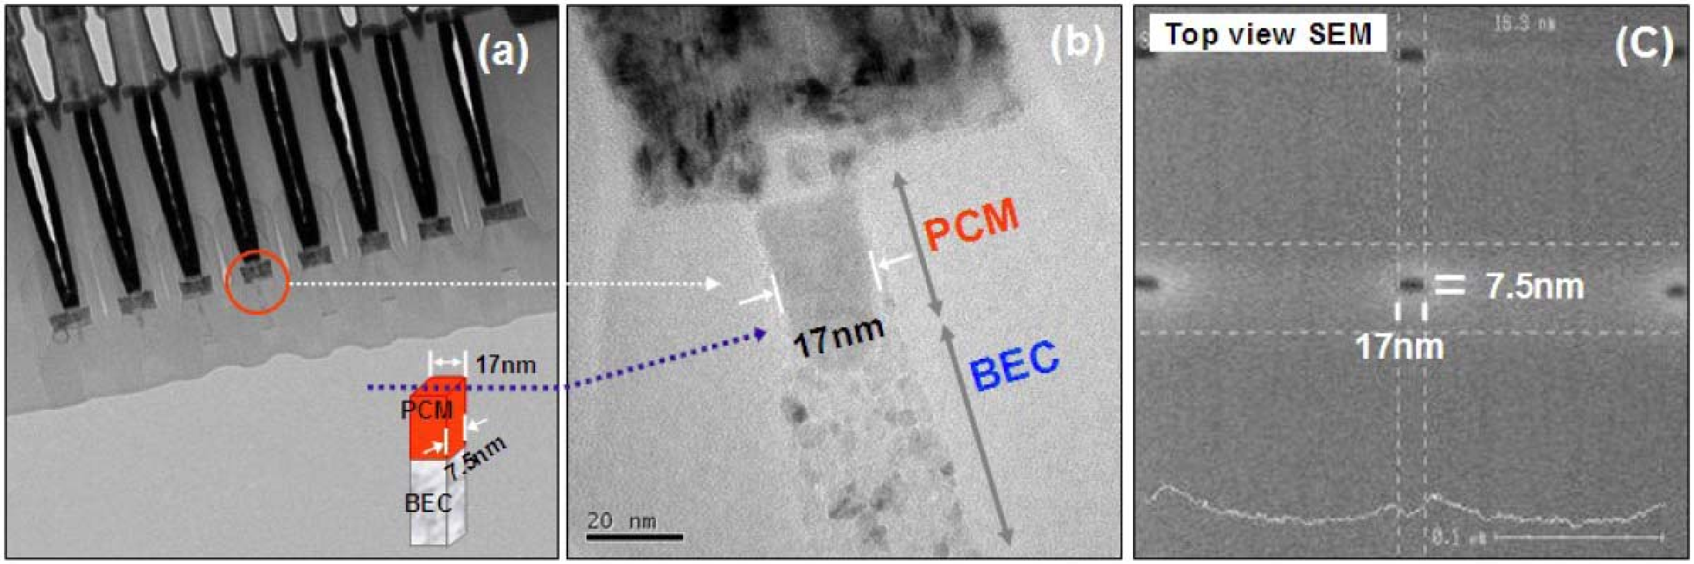

Kim (Samsung), VLSI 2010

Input from the device cost crystal ball

- Projections assume:
  - ... 40% CAGR
  - ... that Flash scaling will falter
- To replace flash, an SCM technology must:
  - ... have a device cost that falls in the shaded region
  - ... have a significant performance and endurance **advantages**
  - ... be manufacturable with **little or no disruption** to current fabs

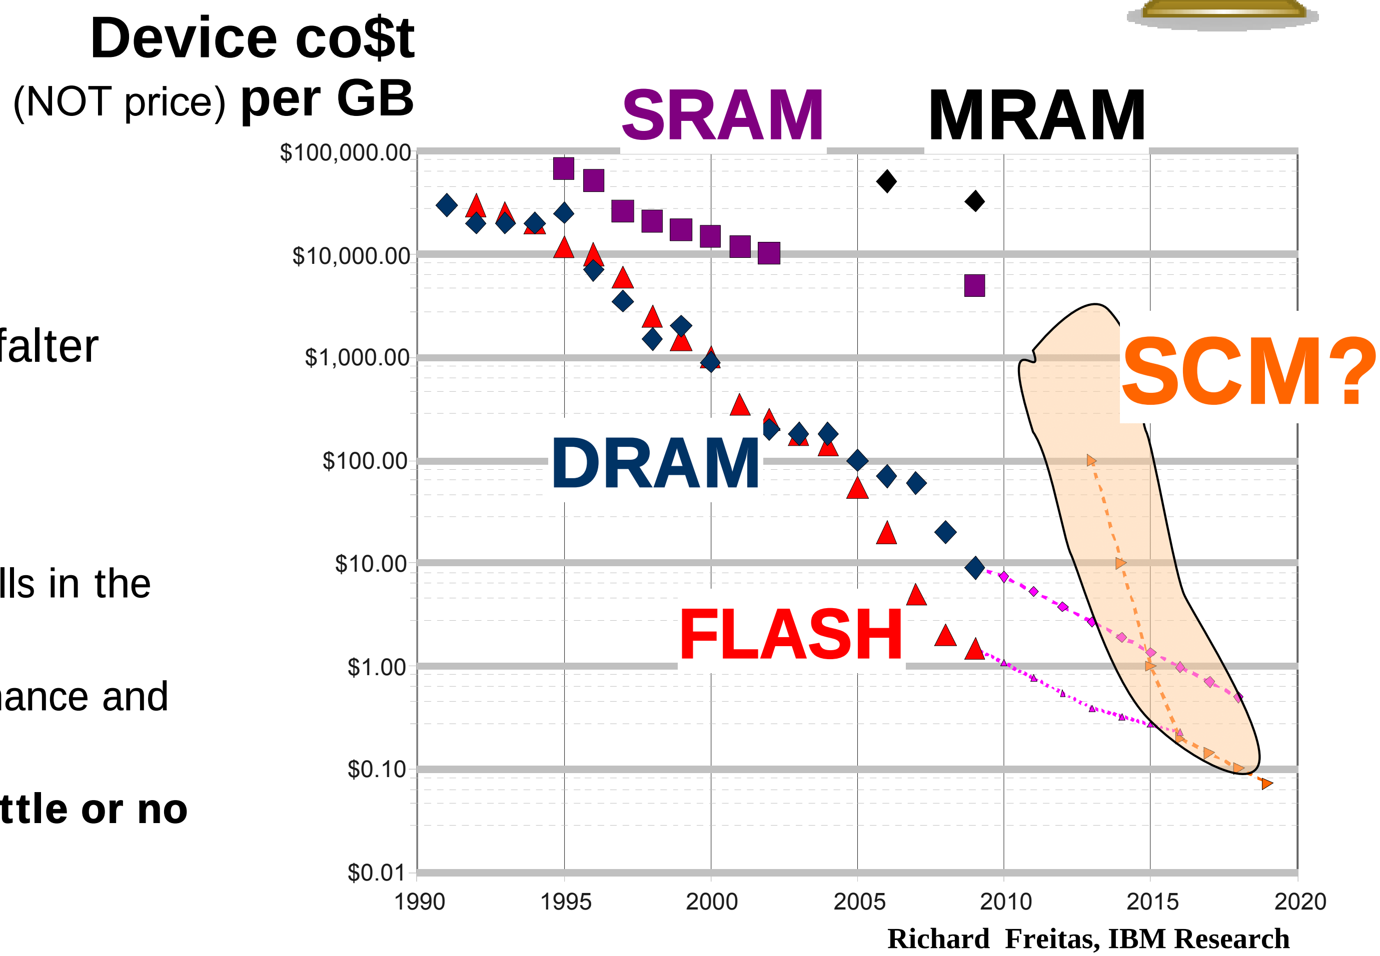

Price Trends for Magnetic and Solid State Disks

- A fair price comparison includes controllers, packaging, etc.
- SSDs will not replace consumer disks any time soon
- SSDs may drive out Enterprise disks
  - Performance
  - Power
- An SCM technology (PCM?) will replace Flash
  - Write endurance
  - Performance

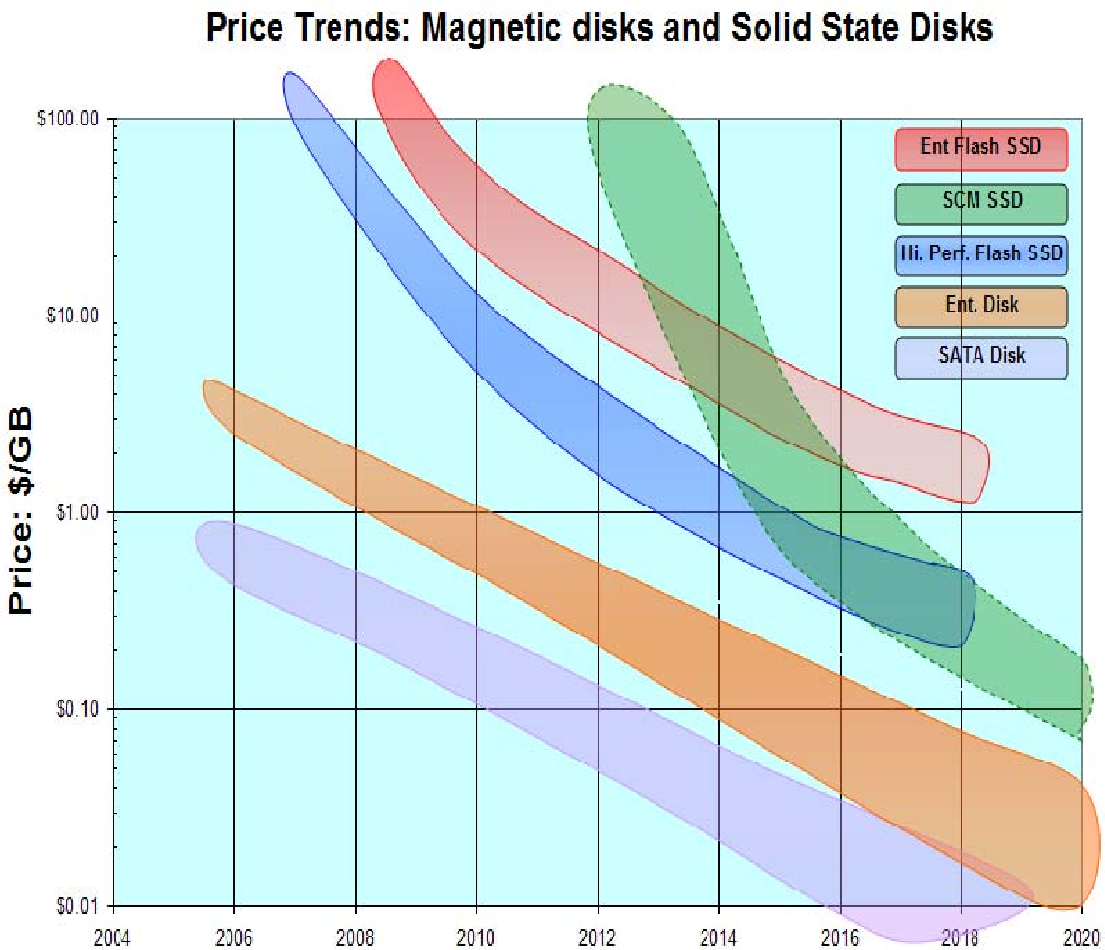

What does the future hold?

- Consumer disk and enterprise tape will persist for the foreseeable future
- Flash will come into its own
- Flash may drive out enterprise disk, and if it doesn't, SCM will
- When will SCM arrive?

  That will depend on the path the NAND industry takes after the 16-20nm node...

  - 3-D NAND succeeds → new NVMs (such as PCM, RRAM, STT-RAM) will develop slowly, driven only by SCM/embedded market
  - 3-D NAND **fails** or is **late** → one new NVM will be driven rapidly by NAND market
- If the latter, SCM will become the dominant storage technology by 2020
- The application software stack will be redesigned to utilize SCM-enabled persistent memory

For more information & *acknowledgements*

• R. S. Shenoy, K. Gopalakrishnan, **Bryan Jackson**, K. Virwani, G. W. Burr, C. T. Rettner, A. Padilla, **Don S. Bethune**, R. M. Shelby, A. J. Kellock, M. Breitwisch, E. A. Joseph, R. Dasaka, R. S. King, K. Nguyen, A. N. Bowers, M. Jurich, A. M. Friz, T. Topuria, P. M. Rice, and B. N. Kurdi, "Endurance and Scaling Trends of Novel Access-Devices for Multi-Layer Crosspoint Memory based on Mixed Ionic Electronic Conduction (MIEC) Materials," *Symposium on VLSI Technology*, T5B-1, (2011).

• **Alvaro Padilla**, G. W. Burr, K. Virwani, A. Debunne, C. T. Rettner, T. Topuria, P. M. Rice, B. Jackson, D. Dupouy, A. J. Kellock, R. M. Shelby, K. Gopalakrishnan, R. S. Shenoy and B. N. Kurdi, "Voltage polarity effects in {GST}-based phase change memory: physical origins and implications," IEDM Technical Digest, paper 29.4, (2010).

• K. Gopalakrishnan, R. S. Shenoy, C. T. Rettner, **Kumar Virwani**, Don S. Bethune, R. M. Shelby, G. W. Burr, A. J. Kellock, R. S. King, K. Nguyen, A. N. Bowers, M. Jurich, B. Jackson, A. M. Friz, T. Topuria, P. M. Rice, and B. N. Kurdi, "Highly-Scalable Novel Access Device based on Mixed Ionic Electronic Conduction (MIEC) Materials for High Density Phase Change Memory (PCM) Arrays," *Symposium on VLSI Technology*, 19.4, (2010).

• G. W. Burr, **Matt J. Breitwisch**, **Michele Franceschini**, **Davide Garetto**, K. Gopalakrishnan, B. Jackson, B. Kurdi, C. Lam, **Luis A. Lastras**, A. Padilla, **Bipin Rajendran**, S. Raoux, and R. Shenoy, "Phase change memory technology," Journal of Vacuum Science & Technology **B**, **28**(2), 223-262, (2010).

• G. W. Burr, **Bulent Kurdi**, J. C. Scott, C. H. Lam, **Kailash Gopalakrishnan**, and **Rohit Shenoy**, "An overview of candidate device technologies for Storage-Class Memory," *IBM Journal of Research and Development*, **52**(4/5), 449 (2008).

• S. Raoux, G. W. Burr, M. J. Breitwisch, C. T. Rettner, Y. Chen, R. M. Shelby, M. Salinga, D. Krebs, S. Chen, H. L. Lung, and C. H. Lam, "Phase-change random access memory — a scalable technology," **52**(4/5), 465, *IBM Journal of Research and Development*, (2008).

• **Rich Freitas** and **Winfried Wilcke**, "Storage Class Memory, the next storage system technology," **52**(4/5), 439, *IBM Journal of Research and Development*, (2008).

• **Yi-Chou Chen**, **Charlie T. Rettner**, **Simone Raoux**, G. W. Burr, S. H. Chen, **R. M. (Bob) Shelby**, M. Salinga, W. P. Risk, T. D. Happ, G. M. McClelland, M. Breitwisch, A. Schrott, J. B. Philipp, M. H. Lee, R. Cheek, T. Nirschl, M. Lamorey, C. F. Chen, E. Joseph, S. Zaidi, B. Yee, H. L. Lung, R. Bergmann, and **Chung Lam**, "Ultra-Thin Phase-Change Bridge Memory Device Using GeSb," *IEDM Technical Digest*, paper S30P3, (2006).

