




IBM Research

Storage Class Memory

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

Geoffrey W. Burr

IBM Almaden Research Center



April 12, 2010

Outline

▪ Motivation

- by 2020, server-room power & space demands will be too high
- evolution of hard-disk drive (HDD) storage and Flash cannot help
- need a new technology – **Storage Class Memory (SCM)** – that combines
 - ❖ the benefits of a solid-state memory (**high performance** and **robustness**)
 - ❖ the **archival capabilities** and **low cost** of conventional HDD

▪ How could we build an SCM?

- combine a scalable non-volatile memory
STT-RAM? RRAM? PCM?
- with **ultra-high density** integration, using
 - ❖ micro-to-nano addressing
 - ❖ multi-level cells
 - ❖ 3-D stacking


▪ Conclusion

- With its combination of **low-cost** and **high-performance**,
SCM could impact much more than just the server-room...

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

...yet critical applications are also undergoing a paradigm shift

Compute-centric paradigm

Main Focus: Solve differential equations

Bottleneck: *CPU / Memory*

Typical Examples: Computational Fluid Dynamics
Finite Element Analysis
Multi-body Simulations

Data-centric paradigm

Analyze petabytes of data

Storage & I/O

Search and Mining
Analyses of social/terrorist networks
Sensor network processing
Digital media creation/transmission
Environmental & economic modeling



(at 90% CGR → need 1.7 PB/sec)

• **5.6 million HDD**

- **19,000** sq. ft. !!
- **25 Megawatts**




(at 90% CGR → need 8.4G SIO/sec)

• **21 million HDD**

- **70,000** sq. ft. !!
- **93 Megawatts**

**Extrapolation
to 2020**

Problem: The access-time gap between memory & storage




- Modern computer systems have to be carefully designed to **hide this gap**
- Lots of server-room **power** spent to keep disk-drives spinning
- With no on-chip storage, sensor (& other) data must be passed **deep into the network**
- **Slow searching** through huge databases because all the data must come off of disk first


*...what if we could build a technology that could **bridge** this gap?*

Storage Class Memory

A solid-state memory that **blurs the boundaries** between storage and memory by being **low-cost, fast, and non-volatile**.




Storage-type vs. memory-type Storage Class Memory



Storage Class Memory

A solid-state memory that **blurs the boundaries** between storage and memory by being **low-cost, fast, and non-volatile.**



■ SCM requirements for **Memory (Storage) apps**

- No more than 3-5x the **Cost** of enterprise HDD ($< \$1$ per GB in 2012)
- **$< 200\text{nsec}$ ($< 1\mu\text{sec}$)** Read/Write/Erase time
- $> 100,000$ **Read I/O operations** per second
- **$> 1\text{GB/sec}$ ($> 100\text{MB/sec}$)**
- **Lifetime** of $10^9 - 10^{12}$ write/erase cycles
- 10x lower **power** than enterprise HDD

Can HDD & Flash improve enough to help?



■ Magnetic hard-disk drives (HDD)

- **bandwidth** issues (hidden with parallelism, but at power/space cost)
- **slow access** time (not improving, hard to hide with caching tricks)
- **reliability** (newest drives are *less reliable* → data losses inevitable)
- **power** consumption (must keep drives spinning to avoid even longer access times)

■ Flash

- slow read/write **access time** (yet processors keep getting faster)
- low write **endurance** ($<10^6$) (need $>10^9$ for continuously streaming data)
- block architecture
- **scalability** beyond the end of this decade? (yes, at least until the 1X node)

NAND Scaling Road Map




Evolution??

- Migrating to Semi-spherical TANOS memory cell 2009
- Migrating to 3-bit cell in 2010
- Migrating to 4-bit cell in 2013
- Migrating to 450mm wafer size in 2015
- **Migrating to 3D Surround-Gate Cell in 2017**

Source: Chung Lam, IBM

How does SCM compare to existing technologies?



Candidate device technologies

- **Improved Flash**

- improvements would be required in write endurance & speed

- **FeRAM** – commercial product but difficult to scale!

- **MRAM** – commercial product, also difficult to scale!

- **STT-RAM** – sacrifice infinite endurance for lower currents

- **Racetrack memory** – new concept w/ promise, still at point of early basic physics research



- **RRAM (Resistive RAM)**

- Organic & polymer memory
- **Memristor**
- **Solid Electrolyte**




- **PC-RAM (Phase-change RAM)**

RRAM (Resistive RAM)

- Numerous examples of materials showing hysteretic behavior in their I-V curves
- Mechanisms not completely understood, but major materials classes include
 - metal nanoparticles(?) in **organics**
 - could they survive high processing temperatures?
 - metallic filaments in **solid electrolytes**
 - oxygen vacancies in **transition-metal oxides**
 - tremendous progress over last few years
 - Integrated demos with <1us write, 1e6 endurance in ~50nm diameter devices
 - things I'd worry about if I worked on RRAM:
 - need for higher-voltage forming step
 - “switching reproducibility”
 - bipolar operation/read-disturb vs. cross-point diode operation
 - voltage-time dilemma (switch in 100ns, can it still retain for 10yrs?)



Phase-change RAM



Potential headache:
High power/current
 → affects scaling!

Potential headache:
If crystallization is slow
 → affects performance!


PCM @ IBM Almaden

■ Prototype memory devices


$3\text{nm} * 20\text{nm} \rightarrow 60\text{nm}^2$

$\approx$ Flash roadmap for **2013**

$\rightarrow$ phase-change scales




[Chen
IEDM 2006]





W defined by **lithography**

H by **thin-film deposition**

■ Phase-change materials work



■ Device & materials modeling




[B. S. Lee
Science 2009]

PCM @ IBM Almaden

■ Ultra-scaled PCM pore cells

→ PCM as a CMOS-compatible synapse




■ Improvements to PCM

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

Cost structure of silicon-based technology

Co\$t determined by


- cost per wafer
- # of dies/wafer
- memory area per die [sq. μm]
- **memory density** [bits per $4F^2$]
- **patterning density** [sq. μm per $4F^2$]





Device	Critical feature-size F	Area (F^2)	Density (Gbit /sq. in)
Hard Disk	50 nm (MR width)	1.0	250
DRAM	45 nm (half pitch)	6.0	50
NAND (2 bit)	43 nm (half pitch)	2.0	175
NAND (1 bit)	43 nm (half pitch)	4.0	87
Blue Ray	210 nm ($\lambda/2$)	1.5	10

[Fontana:2004,
web searches]


Need a 10x boost in density BEYOND Flash!




Paths to ultra-high density memory



...store **M** bits/cell
with 2^M multiple levels



...go to 3-D with
L layers




MLC with PCM


- Direct write NOT sufficiently accurate
- Various schemes for iterative write
- Iterative write works, but
 - latency, endurance impact?
 - errors despite fixed # of iterations?
- Drift of resistance states

2:30pm **Luis Lastras Montano**

2:50pm **Michele Franceschini**






Single write



[Nirschl:2007 IEDM]

Iterative write



[Kang:2008 VLSI]

3-D stacking of PCM

Intel/Numonyx: use 2nd amorphous PCM-like material as threshold switch

- similar material to PCM
- simply stack on PCM

but...

- readout by breakdown, so
→ **read-reinforce SET**, but destructive read of MLC?
- high read-currents → bandwidth?
- half-select leakage? (linear plots?)







Fig. 7 One layer of the PCMS array fully integrated with a CMOS technology. The memory cell stack, including rows and columns, is shown sandwiched between M2 and M3. M3 is not shown.

[Kau
IEDM 2009]

IBM Almaden: PCM-capable access device

→ to appear in VLSI Technology Symposium, June 2010

If you could have SCM, why would you need anything else?



Conclusions

▪ Motivation

- by 2020, server-room power & space demands will be too high
- evolution of hard-disk drive (HDD) storage and Flash cannot help
- need a new technology – **Storage Class Memory (SCM)** – that combines
 - ❖ the benefits of a solid-state memory (**high performance** and **robustness**)
 - ❖ the **archival capabilities** and **low cost** of conventional HDD

▪ How could we build an SCM?

- combine a scalable non-volatile memory
STT-RAM? RRAM? PCM?
- with **ultra-high density** integration, using
 - ❖ micro-to-nano addressing
 - ❖ multi-level cells
 - ❖ 3-D stacking

▪ Conclusion

- With its combination of **low-cost** and **high-performance**,
SCM could impact much more than just the server-room...

For more information...

- G. W. Burr, M. J. Breitwisch, **Michele Franceschini**, **Davide Garetto**, K. Gopalakrishnan, **Bryan Jackson**, B. Kurdi, C. Lam, **Luis A. Lastras**, **Alvaro Padilla**, B. 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).
- **Simone Raoux**, G. W. Burr, M. J. Breitwisch, **Charlie Rettner**, Y. Chen, **R. M. (Bob) Shelby**, M. Salinga, D. Krebs, S. Chen, H. 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).
- T. Nirschl, J. B. Philipp, T. D. Happ, G. W. Burr, **Bipin Rajendran**, M.-H. Lee, A. Schrott, M. Yang, **Matt Breitwisch**, C.-F. Chen, E. Joseph, M. Lamorey, R. Cheek, S.-H. Chen, S. Zaidi, S. Raoux, Y.C. Chen, Y. Zhu, R. Bergmann, H.-L. Lung, and **Chung Lam**, "Write strategies for 2 and 4-bit multi-Level phase-Change memory," *IEDM Technical Digest*, paper 17.5, (2007).
- **Yi-Chou Chen**, C. T. Rettner, S. Raoux, G. W. Burr, S. H. Chen, R. M. 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 C. Lam, "Ultra-Thin Phase-Change Bridge Memory Device Using GeSb," *IEDM Technical Digest*, paper S30P3, (2006).
- M. Breitwisch, T. Nirschl, C.F. Chen, Y. Zhu, M.H. Lee, M. Lamorey, G.W. Burr, E. Joseph, A. Schrott, J.B. Philipp, R. Cheek, T.D. Happ, S.H. Chen, S. Zaidi, P. Flaitz, J. Bruley, R. Dasaka, B. Rajendran, S. Rossnagel, M. Yang, Y.C. Chen, R. Bergmann, H.L. Lung and C. Lam, "Novel Lithography-Independent Pore Phase Change Memory," *Symposium on VLSI Technology*, pages 100-101, (2007).



IBM Research

Thank you!

