CS/ECE 6780/5780

Al Davis

Today's topics:

- Non-volatile RAM
  - a survey
    - current and future technologies
    - a bit of detail on how the technology works
  - much of this material was developed with
    - Christopher Hoover, HP Labs

Generic Taxonomy: V & NV

- **Volatile (last lecture had the details)**
  - **SRAM - 5 or 6 transistors per cell**
    - » fast but costly & power hungry
    - » usage
      - on chip - caches, register files, buffers, queues, etc.
      - off chip usage now rare except in embedded space
  - **DRAM - 1 T & 1 C per cell (lots of details later in the term)**
    - » focus on density and cost/bit
      - too bad both power and delay properties are problematic
    - » usage - main memory
      - EDRAM now moving on chip for large "last cache" duties
    - » specialty parts for mobile systems
      - low-power
      - self-refresh
      - takes advantage of light usage
    - » battery backed DRAM - common in data-center

NV

- **Traditional non-volatile**
  - **Magnetic Disk**
    - » cheap
    - » mixed use: file system and scratch
  - **CD, DVD**
    - » even cheaper per unit but less capacity
    - » media and SW distribution, personal archival
  - **Tape**
    - » cheapest
    - » archival storage
  - **Solid state**
    - » more spendy but faster
      - PROM in various flavors - now primarily masked on chip
      - FLASH has essentially taken over at the component level
      - new contenders are on the horizon however
        - – focus of today's discussion

NVRAM Alternatives

	Flash		FeRAM	MRAM	PCM	Probe Storage
Cell Type	NOR 1T	NAND 1T	1T/1C	1T/1R	1T/1R	AFM-based
Cell Size (F^2)	10	4 or 5	30-100	30-50	8-16	0.4 (no litho)
Endurance W/R	10^6/inf		10^12/10^12	>10^14/inf	10^12/inf	10^5-10^12/10^7-inf
Read Time (random)	60 ns	60 ns / serial	40 + 80 ns (read + write destructive read)	30 ns	60 ns	2-20ms
Write time (byte)	1 us	200 us / page		30 ns	10 ns	0.1-1 ms for each tip
Erase time (byte)	1 s / sector	2 ms / block		30 ns	150 ns	< 1 us /bit
Scalability	Fair	Fair	Poor	Poor	Good	Very Good
Scalability Limits	Tunnel oxide, HV		Capacitor	Current Density	Litho	None
Multi-bit capability	Yes		No	No	Yes	No
Relative cost/bit	Medium	Low	High	High	Medium	Very low
Maturity	Very high		Medium	Low	Low	Very low

Source: Pirovano ICMTD-2005

Commercial Aspects

- **Recent reports a bit more gloomy**
  - **due to world economy issues**
- **2004 $16B - predicted $72B by 2012**
  - **CAGR = combined annual growth rate**
    - » **critical metric from a business perspective**
  - **NOR - 30% CAGR in '04, similar now but reports vary**
    - » **1 Gb and 2 Gb packages**
  - **NAND - 70% CAGR in '04 but now down to ~20%**
    - » **8 - 64 Gb packages (3D)**
    - » **needs a write controller**
      - **today it's on the chip**

NOR vs. NAND Geometry

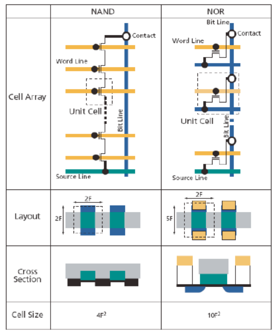

Source: Micron

NAND: $4F^2$
NOR: $10F^2$
DRAM: 6-$8F^2$

NAND vs. NOR Properties

Source: Micron

	NAND	NOR
Advantages	Fast writes Fast erases	Random access Word writes Read-while-write Read-while-erase
Disadvantages	Slow random access No word writes	Slow writes Slow erases
Random read	25 us first byte, 0.03 us for remaining 2,111 bytes	0.12 us
Sustained read (sector basis)	23 MB/s (x8) or 37 MB/s (x16)	20.5 MB/s (x8) or 41 MB/s (x16)
Random write	~300 us/2112 bytes	180 us/32 bytes
Sustained write (sector basis)	5 MB/s	0.178 MB/s
Erase block size	128 KiB	128 KiB
Erase time (typ)	2 ms	750 ms
Part Number	MT29F2G08A	MT28F128J3

Flash Component

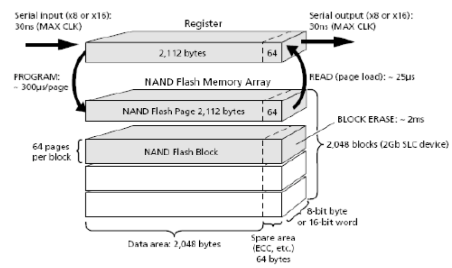

Source: Micron

NAND Trends

Source: Shin, 2005 Symp. VLSI Ckts

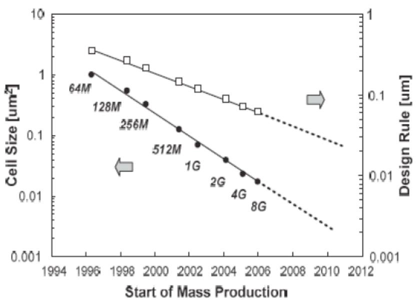

NAND vs. DRAM 2007

- DRAM
  - 65 nm process
  - 2 Gb on 100 mm² die
  - 1.94 Gb/cm²
- NAND SLC
  - 56.7 nm process
  - 4 Gb on 80.8 mm² die
  - 4.3 Gb/cm²
- NAND MLC (2 bits/cell)
  - 56.7 nm process
  - 8 Gb on 80.8 mm² die
  - 11 Gb/cm²

What's Wrong with FLASH?

No problem unless
- You care about speed, power
  - » looks good when compared to disk except for price
- OR operate in write rarely land
  - » exactly the case for embedded system code
- There are some alternatives BUT
  - They all have some downsides
    - » maturity, expense, density, market & investment, etc.
    - » scaling claims - just how real are they
- Worth tracking since FLASH futures may not be bright
  - IEDM 2005 Panel ==> run out of gas in 2010 likely?
    - » note it's now 2010 and FLASH is still running strong
    - » beware of predictions
- Question
  - obvious market niche: thumb drives, cameras, micro-controllers, etc.
  - SSD and checkpoint storage role might be in doubt
    - » but currently SSD's are gaining ground in the "cloud"

What's Next?

- Talk about likely future NVRAM candidates
  - Ignore quantum and DNA soup like structures
    - » Distant future maybe - near future unlikely
    - » Note: fab ramp is as important as the devices
  - Many have been around for a long time
    - » Development to deployment is a long and rocky road
- How they work focus
  - Maybe more technology than a user cares about
  - Hopefully aid awareness of what to look for as the technologies progress
  - Architects must track technology trends
- Try and assess where their future might lie
  - Memory shapes the systems around it
    - » A fact most architects have ignored to date
      - a few architects have always been memory centric
        - e.g. Burton Smith & somebody you know
    - » Von Neumann's corollary
      - memory is the bottleneck → so focus on getting it right

Flash (Hot Chips '04)

	NOR Flash	NAND Flash
Applications	Code, data	Mass storage
Future applications	MLC: mass storage	Code and data
Density range	Up to 512Kb	Up to 4Gb
READ latency	60ns-120ns	25µs
Max Read bandwidth	41 MB/s-112 MB/s (16b)	40 MB/s (16b bus)
Max Write bandwidth	0.25 MB/s	5MB/s
Erase time	400ms (128KB blk)	2ms (128KB block)
Read device current	1.6x	1x
Write device current	3x	1x

Source: Micron tutorial

Note - NAND read times haven't changed in years

Density improvement is excellent

Known FLASH issues

- **Speed - slow writes OK, but 25 usec reads??**
  - High voltage on both read and write create problems
    - » Charge pump takes time
    - » Jitter on bit lines requires lengthy settle margin
  - Conclusion is that reads are unlikely to get much faster
- **Retention**
  - Thicker tunnel oxide (7-12nm) provides good retention, but
    - » High voltage requirements create reliability issue.
      - Channel punch through, junction breakdown, etc.
      - Also increases the read and write energies
- **Scaling**
  - Concern over single defect memory loss limits vertical scaling
  - High voltage also limits lateral scaling to some extent
  - Rad hard arrays are difficult to achieve
  - Support circuitry doesn't scale as well as the arrays

More Issues

- **Retention**
  - 10^6 block erase wear out
    - » Gets considerably worse for multi-bit cells
  - Density/Retention trade-off
  - Wear leveling a must for computer systems
    - » Who cares for iPods, cameras, ES code, etc.
      - there just aren't that many writes
      - both reads are sequential, writes are block oriented
- **Use model**
  - Somewhat goofy
    - » Write once cells or block erase
    - » Complex controller
  - Not much worse than DRAM however
    - » with all of it's timing complexity

SONOS/MONOS

- **ONOS - oxide nitride oxide semiconductor**
  - M=metal gate - common outside US
  - S= silicon - more common in US
- **Varying views**
  - some view as a FLASH evolution
  - others view as a fundamentally different technology
  - both views are credible but who cares?

Why should we be interested

- **Relatively mature**
  - already in production
    - » SONY is basing their SoC strategy on this
    - » TSMC, Grumman, Hitachi, Philips & Toshiba also have the process
    - » Compatible with CMOS fab
  - density
    - » $6F^2$ cell (same as DRAM)
  - lower than FLASH program voltage 5-8V
    - » reduces the wearout and lateral scaling problem
  - scales better
    - » working @ 20 nm, 1ms program and erase
    - » reported IEDM '05 by TSMC (J. R. Hwang et al)

Not a new technology

- **Current usage**
  - Satellite and space craft
    - » Inherently rad-hard
      - important at small size & enables cheap packaging
- **Why haven't we seen it until recently**
  - Concerns about data retention
  - Density not as good as FLASH
    - » cost is critical in high volume markets
- **What's changed**
  - 2 bit per cell ==> density better than FLASH
    - » Possible for FLASH too but much harder to control
  - Retention now at 10 years after 10^7 write/erase
    - » Primarily due to anneal w/ deuterium rather than hydrogen
    - » Promise of hi-K dielectrics - viz. HfO & HfO_2

MONOS/SONOS vs. Floating Gate (FLASH)

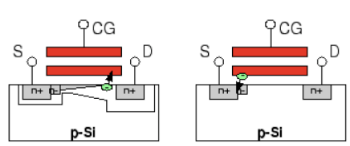

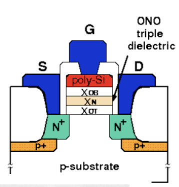

Sources: Bu & White (IEDM '05) & Sony Corporation

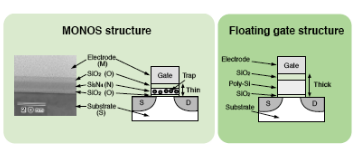

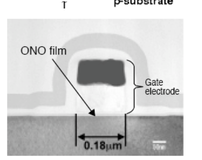

■ Figure 9 TEM Photograph for MONOS Memory Transistor

SONOS Operation

- **Write - positive gate bias 5-8V**
  - Electrons tunnel through thin top layer
  - Trapped in cavities in the nitride layer
    - » Due to thicker bottom layer oxide
  - Current thickness: 2, 5-10, 5 nm
- **Read @ 4.5V**
  - Vds forward bias
  - If Ids current then 0, else 1
- **Block Erase**
  - Similar to FLASH but @ 2V

SONOS Pro's and Con's

- **Pros**
  - Scaling and wear-out much improved over FLASH
    - » Wear out due to electrons trapped in Nitride layer
    - » FLASH - oxide deterioration and single point of failure
  - Reduced Energy due to lower voltage operation
    - » Philips has a 2T version which decreases energy/op by 3-5x
- **Cons**
  - Write and erase currently slower than FLASH
    - » Promise to be faster in 65 nm - but I can't find a report to confirm
- **Bizarre**
  - No report found in the literature on read access times

Phase Change RAM

- **Tower of Babel naming**
  - PCRAM, PRAM, PCM, OUM, CRAM
- **Basis**
  - Chalcogenide material
    - » 2 states - crystalline and amorphous
      - actually lots of states in between
    - » 0 = Amorphous - quench after heating to > 619 C
      - high resistive, high refractive index
    - » 1 = Crystalline - heat > 223 C
      - low resistive, low refractive index
    - » Quench must cool to < 100 C
  - NOTE
    - » Properties and temps vary slightly w/ specific material

Also Not a New Technology

- **Timeline**
  - '66 Stanford Ovshinsky (ECD) first patent
  - '69 ECD patent and working device
  - '99 Ovonyx joint venture starts as license source
  - '04 64 Mb Samsung part
  - '05 256 Mb Samsung plus w/ 100 uA programming
    - » Hitachi 100 uA @ 1.5v programming current
  - '06 BAE puts rad-hard parts in space
    - » 1st commercially available part
  - '06 STM 128 Mb commercial
  - '07 IDF demo by Justin Rattner of Intel version
  - today – multiple vendors and higher capacity parts

We use this stuff now - differently

- **CD-RW and DVD-RW**
  - Chalcogenide based
  - Laser to do the heating
  - Read based on refraction differences - not resistance

Basic Device

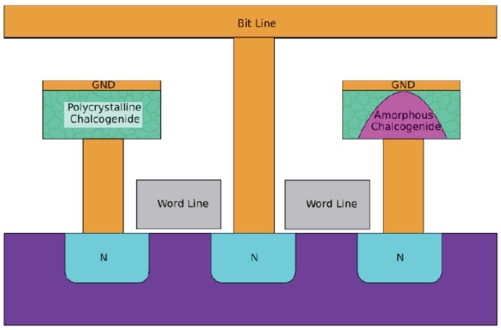

Lot's of Chalcogenides

Binary	Ternary	Quaternary
Ga Sb	$Ge_2Sb_2Te_5$	Ag In Sb Te
In Sb	In Sb Te	(Ge Sn)Sb Te
In Se	Ga Se Te	Ge Sb (Se Te)
Sb_2 Te_3	Sn Sb_2 Te_4	$Te_{81}Ge_{15}Sb_2S_2$
Ge Te	In Sb Ge	

Most commonly used is GST

Source: Ovonyx

Assymetric Properties

Amorphous Phase **Crystalline Phase**

TEM Images

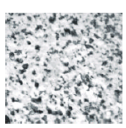

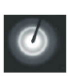

Electron Diffraction Patterns

Material Characteristics

- Short-range atomic order
- Low free electron density
- High activation energy
- High resistivity

- Long-range atomic order
- High free electron density
- Low activation energy
- Low resistivity

Source: Ovonyx

Large R diff & Wide operating range

R(I) vs. Programming Current at Room Temperature and 77K

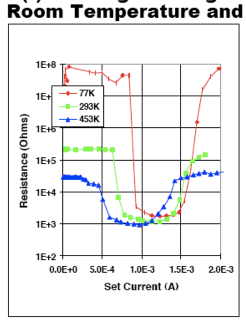

Data demonstrating wide operating temperatures of OUM technology.

Multi-bit/cell option is obvious

16-bit/cell demonstrated

K=C+273.15

Source: Ovonyx

Excellent Retention & Durability

Cycle Life > 10^{13} Write/Erase Cycles

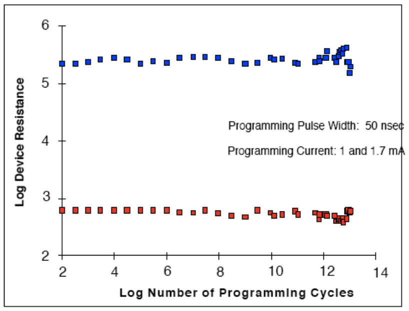

10 year retention at 130 C

Retention reduced with higher temps

Source: Ovonyx

Multi-bit requires Multi-pulse

Multi-State Storage

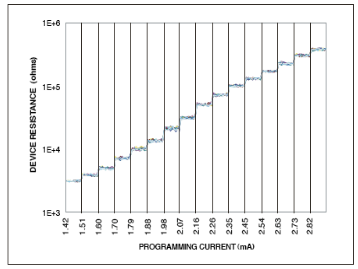

Easier control regime than a single pulse w/ varying duration

Source: Ovonyx

- Multiple-bit storage in each memory cell (10 pulses per step, repeated ten times.)

Basically a very cheap material

Cost/Bit Reduction

- Small active storage medium
- Small cell size – small die size
- Simple manufacturing process – low step count
- Simple planar device structure
- Low voltage – single supply
- Reduced assembly and test costs

Source: Ovonyx

Ovonyx claimed advantages

Near-Ideal Memory Qualities

- Non-volatile
- High endurance – >10^{13} demonstrated
- Long data retention – >10 years
- Static – no refresh overhead penalty
- Random accessible – read and write
- High switching speed
- Non-destructive read
- Direct overwrite capability
- Low standby current (<1μA)
- Large dynamic range for data (>40X)
- Actively driven digit-line during read
- Good array efficiency expected
- No memory SER – RAD hard
- No charge loss failure mechanisms

Other Advantages

- Scalability
  - Primarily limited by lithography
    - » caveat - thermal isolation bands may not scale as well
      - claim is quaternary materials are the solution here
        - several lab demos but hasn't hit commercial space yet
        - why? material is more expensive and fabs may not be tuned for it
  - Performance improves linearly w/ feature size
- What we care about in a write mostly environment
  - E.g. check point memory
    - » Where the ideal is "read never" since nothing bad happened
  - write time is short
  - low write energy
- 3D possible w/ epitaxial thin films
  - Claimed but not demonstrated as far as I can tell

OK where's the downside

- Based on the Ovonyx spin
  - Everybody should use this stuff and FLASH should be dead
  - It isn't so what's up?
- HEAT
  - Semi-conductors give off ~50% of their power as heat
    - » The rest is returned to the power supply
  - In PCram write operations - ~100% of the power is given off as heat
  - Longer quench time if writes to same neighborhood - control problem
- Issues
  - Retention tracks ambient temps
  - Good cooling means higher write currents
  - BIG ONE: material defect issues currently have yield issues
    - » manifests itself in wide range of resistance values
      - e.g. all 0's don't look the same
      - hence wide 0- and 1-band margins
    - » It's a long way from the lab to profitable product

FeRAM/FRAM

- Ferro-electric basis
  - 1 T and 1 C currently
    - » Like DRAM but the C is a ferro-electric device
  - Behavior is similar to the old core memories
    - » But voltage rather than current based
    - » Magnetic polarity is used to determine the state
- Also not a new technology
  - Research
    - » Samsung, Matsushita, Oki, Toshiba, Infineon, Hynix, Symetrix, Cambridge University, University of Toronto and the Interuniversity Microelectronics Centre (IMEC, Belgium).
  - Production
    - » RAMTRON - most of the development
    - » Licensed to Fujitsu with the largest capacity production line

Dwarfed by FLASH

- Gartner Group 2005 reports
  - 18.6 B$ FLASH
  - 23 M$ for Ramtron
    - » probably the largest supplier
      - Fujitsu fab
- Promise (conflicting reports)
  - When compared to FLASH
  - FeRAM offers
    - » lower power
    - » faster write speed
    - » much greater maximum number (exceeding 10^{16} for 3.3 V devices) of write-erase cycles.

FeRAM Device Basics

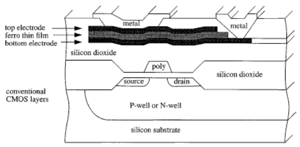

Basic Planar Design

Looks a lot like DRAM using planar C's

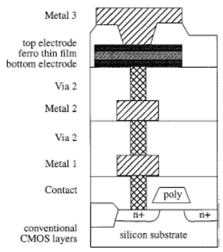

Smaller Stacked Via Variant

Source: Proc IEEE, V. 88, No. 5, May 2000

Ali Sheikholeslami, MEMBER, IEEE, AND P. Glenn Gulak, SENIOR MEMBER, IEEE

Compared w/ Flash and EEPROM

Nonvolatile Memory	Area/Cell (normalized)	Read Access-Time	Write (prog.) Access-Time	Energy* per 32b Write	Energy* per 32b Read
EEPROM	2	50ns	10µs	1µJ	150pJ
Flash Memory	1	50ns	100ns	2µJ	150pJ
Ferroelectric Memory	5 (†)	100ns	100ns	1nJ	1nJ

Note: Flash access times are not correct - makes one wonder about the rest

-- the stacked version area is 2x bigger than Flash

-- Larger size is due to old process

* 2005 Fujitsu line used 350 nm for FeRAM
* 2006 Toshiba Flash process in 60 nm

-- Scalability of the Fe Cap is not discussed

Source: Proc IEEE, V. 88, No. 5, May 2000

FeCap Hysteresis Issues

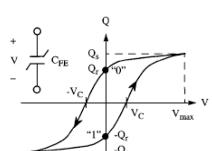

Fig. 6. Hysteresis loop characteristic of a ferroelectric capacitor. Remanent charge (Q_r), saturation charge (Q_s), and coercive voltage (V_C) are the three important parameters that characterize the loop. The + and − signs beside the capacitor symbol represent the applied voltage polarity.

2 Options:

- 1T/1C
  * access transistor compensates for soft hysteresis
- Square hysteresis loop
  * different materials under investigation
  * intersecting wires rather than 1T
  * Given wire scaling it's not clear if this is a win

Source: Proc IEEE, V. 88, No. 5, May 2000

Operation & Issues

- **Destructive read (like DRAM but w/o refresh)**
  - Write a 1: if 0 the reversal generates a small current
  - Detected by sense amp
- **Wear out mechanism**
  - Imprinting - tendency to prefer one state if held there for a long time + neighborhood issue
- **Scaling**
  - Has scaled with Moore's Law as feature size shrinks
- **Issues**
  - Less dense than FLASH
  - But with a longer future? TBD
  - Need for a constant voltage reference ==> column overhead
    - » Potential problem due to future increasing process variation

23 M$ Sold - for What?

- Ramtron shows increases in which segments
  - Automotive air bags and black boxes
    - » Seems odd given lots of magnetics - starters and alternators
  - RFID tags
  - Smart cards
  - Medical
  - Printers
  - RAID controllers
    - » due to better wearout

MRAM - Magneto-Resistive RAM

- Basics
  - 2 Ferromagnetic plates separated by an insulator
- Not a new technology once again
  - ‘55 cores used a similar principle
  - ‘00 IBM/Infineon joint development partnership
  - ‘04 16 Mb Infineon prototype
    - » TSMC, NEC, Toshiba announce MRAM cells
  - ‘05 2 GHz MRAM cell demonstrated
    - » Renesas & Grandis show 65 nm MRAM cell
    - » Freescale enters fray with spin torque technology or transistor (STT)
  - 06 Freescale markets 4 Mb STT chip
    - » NEC markets 250 MHz SRAM compatible MRAM

Device

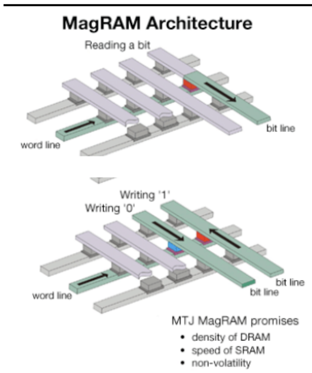

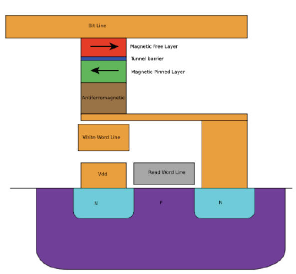

Source: IBM

3 Operation Modes

- “Classic”
  - Read
    - » Two plates same polarity ==> lower R = 0
    - » Opposite polarity ==> higher R = 1
  - Write
    - » Crossing wires as in previous figure
  - Problems
    - » Neighborhood problem at small size
      - False writes to neighboring cells
      - Limits density to >= 180 nm
    - » Only a problem for write

Toggle Mode

- Multi-step write and multi-layer cell
  - More complex process
  - Read
    - » Same as classic
  - Write
    - » Timed write current offsets in the 2 wires to rotate field
    - » Reduces neighborhood effect
      - Scales well to 90 nm

STT

- The current focus of all research
  - Also a multi-layer cell
- Operation
  - Read as usual
  - Write
    - » Inject polarized (spin) electrons
      - As they enter a layer if spin state changes it exerts a "torque" on nearby layer
    - » Advantage
      - Much reduced neighborhood effect
        - Much lower current requirements on bit and word lines
        - Scales below 65nm (haven't seen a limit projection)
          - currently in fab on a 55 nm line
      - Reduces write energy to near read energy

Properties

- Power
  - Read energy =~ DRAM but w/ no refresh
    - » Claim 99% less in normal operation
  - Write energy 3-8x > DRAM for classic
    - » STT solves this as Rd and Wr energy ~ same
- Longevity
  - Indefinite
- Density
  - Until market adopts non-critical (a.k.a. large) fabs used
    - » B$+ fab is the key barrier
  - Hence nowhere near DRAM or FLASH in maturity
    - » but gaining ground rapidly

Properties (cont'd)

- Speed
  - Fast reads and writes < 2ns observed
- Overall
  - Speed similar to SRAM
  - Density similar to DRAM
    - » But not as good as FLASH
  - No degradation
  - No block erase - true random access
- Synopsis
  - It's one to watch closely – as in VERY CLOSELY
    - » this could easily be a DRAM replacement
  - Freescale is probably the best focus

Probe Storage

- **Pioneered by IBM Zurich**
  - Leverages AFM (atomic force microscope) technology
  - Micro-machined cantilever to read and write indentations in a polymer substrate
- **Current demonstration density**
  - 641 Gb/in^2
- **Interestingly**
  - One of the current drivers of this technology is HP
  - QSR currently leads the AFM race

Simple Concept - Hard to Build

- **Idea**
  - Read
    - » Use a cold probe to see if there is a dimple or not
  - Write
    - » Use a hot probe
      - Write 1 - touch probe and a dimple is formed
      - Write 0 - put probe close to surface but not touching
        - If it's already a 1 the dimple goes away
        - If it's a zero nothing happens
    - » VIOLA!
  - Probes fab'd in an array and physically move
    - » Mechanical nature limits speed
    - » Z-axis vibrations are an issue given the small dimensions
    - » Scaling properties are excellent
      - Fundamental limitation is molecular size

IBM calls it Millipede

Read & Array Illustration

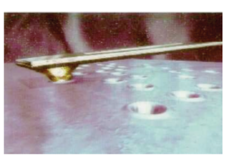

Writing a 1

Source: IBM

Problems

- **Mechanical motion**
  - Small makes it good BUT
    - » Need to move the array likely slower than electrical approach
      - Even at the scalable limit
- **Yields**
  - Still experimental so device yield is off the chart low
- **Role**
  - More likely a disk replacement than anything else

Carbon Nanotube - NRAM

- Least mature of the lot so far
- Nantero owns most of the IP
  - Information more of a marketing blurb than anything else
  - Have not found real publication data to date
    - » Hence no quantification or scaling properties
    - » Numerous press releases which say the same thing
- Nantero claims
  - Faster and denser than DRAM or FLASH
  - Portable as FLASH
  - Resistant to environment: temperature, magnetism

Idea Basis

- Sprinkle nanotubes over a silicon substrate
- Pattern to create a bridge over a 13nm channel
- Then
  - Read
    - » Resistance based - usual sense amps etc.
  - Write
    - » Bend the nanotube down to touch or not
      - Van der Waals forces keep it bent

Structure

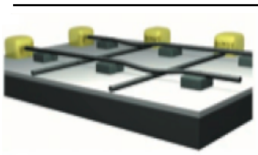

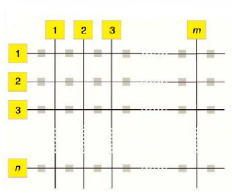

- Nonconductive spacers keep the higher nanotubes flat and raised above the lower level. These spacers can be between five and ten nanometers in height to separate the layers of nanotubes.
- These spacers must be tall enough to separate two layers of nanotubes from each other when both are at rest, yet short enough to allow small charges to attract and cause bends in the nanotubes.

Source: *Thomas Rueckes, et al.,"Carbon Nanotube Based Nonvolatile Random Access Memory for Molecular Computing", SCIENCE, VOL 289, 7 JULY 2000.*

Nantero NRAM

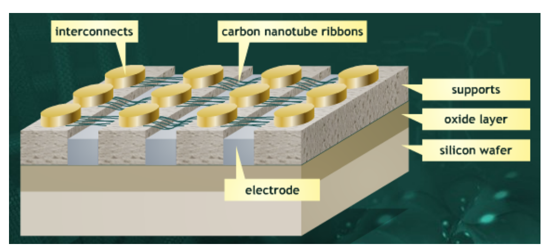

Source: Nantero

Structure

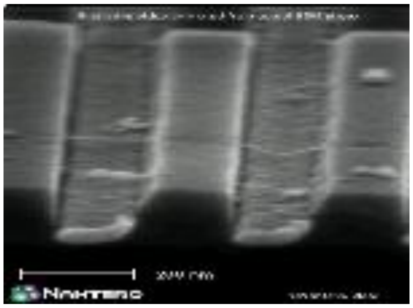

- Fabricated on a silicon wafer, CNT ribbons are suspended 100 nanometers above a carbon substrate layer.

Source: Nantero

NRAM Jury is Still Out

- Concept is good - fab is problematic
  - 5 nm gap between nano-tubes and channel hard to achieve
  - Patterning must be very precise
    - » Tubes have to be thin enough and long enough to bend to create a contact
- Potential for universal memory
  - Fast: 3 ns access demonstrated in 2006 by Nantero
  - Scales: 22 nm demo in 2006
- But
  - Commercial fab and a 1 cell lab test are miles apart

RRAM - Resistive RAM

- Missing Link (so far)
  - Lots of companies claim to be working on it
    - » NTT, Sharp, Samsung, Fujitsu
    - » Have yet to find performance and power numbers
      - Obvious claims - low power, fast, high endurance
- Materials vary
  - Perovskites (PCMO = $Pr_{1-x}Ca_xMnO_3$)
    - » Supply problem: Praseodymium is a rare earth metal
      - hence very expensive
  - Various transition metal oxides (groups 3-12)
  - Chalcogenides (already covered in PCRAM part)

Mechanism

- PCMO
  - Electron concentration at cathode
    - » Due to correct pulse width at low voltage
    - » High resistance
  - Field collapse under negative pulse
    - » Low resistance
  - Problem
    - » 2-5x resistance change - multibit cells problematic
- Transition metal films
  - High resistance change 10-100x
  - Ion migration (similar to electrolytes)

Literature so far

- **Limited to claims and process technology**
  - **all demonstrated cells are relatively large**
    - » **100's of nm**
  - **claim is that they can be as small as 10 nm**
- **Patents refer to single cell properties**
- **Future**
  - **I'll report more if I find it**

Synopsis

	NOR Flash	Nand Flash	SONOS	FeRAM	MRAM	PCRAM	Probe	NRAM
Cell Type	1T	1T	1T	1T/1C	1T/1R	1T/1R	AFM-base	1 channel
Cell Size F^2	10	4-5	6	30-100	30-50	8-16	0.4 (no litho)	?
Endurance W/R	10^6/inf	10^6/inf	10^7-10^8/inf	10^12/10^12	>10^14/inf	10^12/inf	0^5-10^12/10^7-ir	?
Read Time (random)	60 ns	60 ns/serial	?	40+80ns destructive read	30 ns	60 ns	2-20 ms	?
Write Time (byte)	1 us	200 us/page	250 us	80 ns	30 ns	10 ns	.1-1 ms seek <1ms/bit	?
Erase time (byte)	1s/sector	2 ms/block	9 ms	NA	30 ns	150ns	.1-1 ms seek <1ms/bit	?
Scalability	fair	fair	good	poor	poor	good	very good	?
Scalability Limit	tunnel oxide high voltage	tunnel oxide high voltage	ONO oxide	Fe-Cap	Current Density	Lithograpy	None	?
Multi-bit capable	Yes	Yes	?	No	No	Yes	No	No
Relative cost/bit	Medium	Low	Low	High	High	Medium	Very Low	?
Maturity	Very High	Very High	Medium	Medium	High	Low	Very Low	Lowest

Note - values are extrapolated from the varying reports/claims

Source: HP Exascale Memory Report - Al Davis & Christopher Hoover

Memristor

- **Newest contender in the NVRAM space**
  - **Mayday 2008 first report**
    - » **article in Nature**
    - » **developed over many years by a group of people at HP Labs**
      - **principals: Duncan Stewart, Stan Williams, Phil Kuekes**
        - **lots of others involved as well**
  - **billed as the "4th passive element"**
    - » **R (1827), L (1831), & C (1745) are the other 3**
    - » **there's a reason that ~2 centuries passed before this discovery**
  - **mathematically predicted earlier**
    - » **obscure 1971 paper by Leon Chua, UC Berkeley**
      - **proved that memristor behavior could not be duplicated**
        - **by any RLC circuit**
      - **hence it had to be a fundamentally different element**
- **Note passive elements are not enough**
  - **no gain**
    - » **hence use them w/ transistors to do something interesting**

Typical Hydraulic Analogy

- **Ohm's law for water**
  - **R=V/I**
    - » **V is pressure**
    - » **I is still current**
    - » **R is the diameter of the pipe**
  - **For R's – pipe diameter stays the same**
    - » **independent of direction of current flow**
  - **Memristors are weird pipes**
    - » **current flows in one direction the pipe gets bigger**
      - **R goes down**
    - » **other direction pipe gets smaller**
      - **R goes up**
    - » **and the really weird bit**
      - **turn current off and the pipe doesn't change shape**
        - **EVER!**
- **Hmm looks like a Memory Resistor**
  - **hence the name & an NVRAM candidate**

Explaining the 2 Century Gap

- **It's all about scale**
  - **memristive properties are everywhere**
    - » but like quantum physics
    - » at large scale the properties are very hard to observe
    - » at the nanoscale they become apparent
  - **Kirchoff and Maxwell were close**
    - » but in the 1800's it's unlikely that "nano" was part of the language
  - **memristive effect obeys an inverse square law**
    - » 10^6x more pronounced at nm scale than μm scale
    - » 1800's things were at mm scale
      - memristive properties were damped by 10^{12}
      - hence impossible to measure with the tools of the day

Memristor Devices

- **Simple array of crossing nanowires**
  - **with some goop in between at the intersection points**
    - » it took a long time to find the right goop

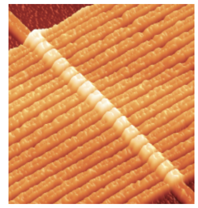

Source: HP Labs, Stan Williams
STM image

Note: Stan's lab was a key player in the development of STM

Another View

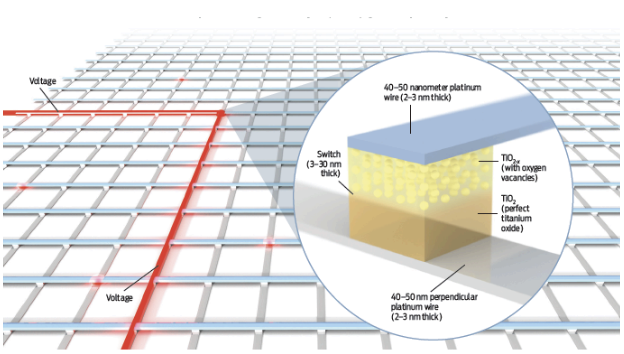

How Does it Work

- **Goal: switch resistance variation of 1K:1**
  - **note transistors typically have a 10K:1 ratio**
- **Materials**
  - **TiO_2 – semiconductor but typically viewed as an insulator**
    - » common – e.g. it's in sunscreen and white paint
  - **$Ti0_{2-x}$ - conductive**
    - » this one was a serendipitous find
      - appeared under STM inspection near the top electrode
      - missing 2-3% of it's oxygen
        - – call these oxygen voids and the have a positive charge
      - Write: key is that these voids MOVE under an applied voltage
        - – travel when current is moving
      - AND in the absence of current flow they stay where the were
    - » applying a small voltage doesn't significantly move the voids
      - so that's how read works

Visually

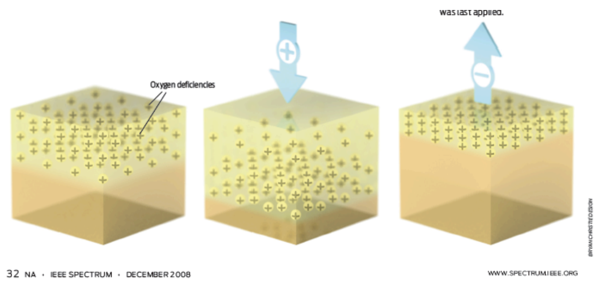

Weird Hysteresis

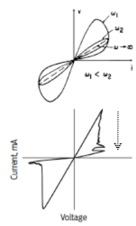

Sharp corners ➔ very fast switching times

in fact too fast to measure even with exotic instruments

Memristor Comments

- **Chua's finding**
  - **it takes 15 transistors to mimic a memristors behavior**
- **Stan's speculation**
  - **memristor will revolutionize circuit design**
  - **final comment is worth reading**
    - » **killer app for transistors was the hearing aid**
      - **and space craft (although he didn't mention that)**
      - **the rest is history**
    - » **killer app for memristors**
      - **"will be invented by a curious student who is now deciding what EE courses to take next year"**
- **Bottom line**
  - **there is some secret sauce but the devices are easily manufactured on today's equipment**
    - » **properties even more important as technology scales**
    - » **ramp to date has been faster than expected**
      - **ability to logic functions recently "leaked"**
- **Definitely one to watch**

Final Remarks

- **NVRAM comes in many flavors**
  - **important and diverse applications demonstrated**
    - » **today by FLASH**
  - **other technologies have different properties**
    - » **some may replace disk and DRAM**
    - » **since memory is the key**
      - **processor architectures may change with new structures**
- **Exciting field**
  - **hopefully this material gives you a foundation**
  - **to understand the literature**
    - » **and more importantly what comes next**