E-NVM Architectural Integration

Survey paper on non-volatile memories. Memory design issues -- we need more volume, higher performance, better energy efficiency, better scalability.
Zero static power with NVM!
Horizontal integration (supplementing existing memory), vertical integration (replacing existing memory)

Flash memory: “storage accelerator”, having no mechanical movement allows for better random read performance.
In the past, over time, applications using HDD have been written to maximize sequential reads (in order to work around the limitations).
Explicit erase before write <- someone mentioned if this can be taken advantage of? Will discuss later.

Integrating NVM:
- as storage system
- as main memory (involves hardware management or OS support)
- as processor cache (most appropriate for LLC)

Jiwon: Confused about the different opinions on similar numbers, depending on the technology

Iris: Probably was referring to the difference between “numbers achieved in particular settings” and “numbers that are more realistic”

Phase Change Memory (PCM)
- temperatures for SET, RESET
- can store multiple levels within a cell!
- slow writes, fast wearout

Karpur: PCM - can store as many bit levels as you can distinguish. This can be applied to MRAM as well. The spins can point in however many directions your reading would allow, making that also a multi-bit storage.

Iris: advantage would be very high density, but there would be challenges, and the challenges are not discussed in detail in this paper.

Semanti: Multi-bit question -- is it three of 0-1 bits in a cell, or a single bit with 3 different values?

Iris: In this paper, it refers to the latter. So not binary, but a different representation.

Sam: For now, wouldn’t this be usable only as a 0-1 device though? But by using the 0, 1 settings that would consume least power.

(I think Iris agreed to this)

Workaround for PCM write issues:

- not really used for cache
- mixed DRAM + eNVM (frequently written data put in DRAM, others in NVM. but this would put a lot of burden on OS or application writer. This would be good

Karpur: Thermodynamic equilibrium on these devices -- has this been studied before? For example, heating up the PCM cell (to write a value) by applying current -- this operation would depend a lot on the temperature that this operates on.

Qishen: I don't think it depends really on the temperature -- the reviewers don't mention the temperature as a constraint.

Karpur: I think my confusion was coming from the fact that the reviewers didn't mention it. Maybe I'm thinking too theoretically, but I think it would impact the operation.

Iris: It might be because the authors of the paper are focused on the architectural side of these new devices. They might not have been aware of the impact that temperature can have on the devices.

Using DRAM as a "write buffer" to PCM seemed vague -- it wasn't clear what writes were being "absorbed"

Sam: If you used DRAM as a cache for PCM, shouldn't the spatial locality be captured well, more so than temporal locality? (I might have understood this the other way around…) Wouldn't it make sense to batch the writes in the DRAM, and then write it all at once to the DRAM?

Yash: If the reads and writes are interleaved, wouldn't the batching not work?

Iris: I think the assumption is that you have a bunch of writes happening at once, so the writes are batched and written to the PCM at the end. But this all seems very write-pattern dependent.

another PCM write workaround--writing only modified bits <- this seems like that could mess up error correcting codes within memory.

It seems like the techniques applied in PCM are similar to the techniques used in regular volatile main memory (for example, preferring reads over writes, write buffers), but problems are exaggerated.

Favor SET operation over RESET in PCM, since this takes less energy.

Can you reduce the RESET current to optimize endurance? Don't have a clear answer.

How do you avoid writing "useless" data? How do you know if data is useless?

Qishen: On endurance, the NVM technologies still have much higher endurance than at least NAND flash. I think it would be good if the industry uses the new NVM devices as a replacement for SSD flash.

Magnetic RAM -- as Karpur mentioned, the paper didn't mention the capability to have varying phases, not just parallel and antiparallel binary.

Writes are not that slow for MRAM, but high current is required.

Karpur: THere is also the possibility that if the temperature goes above some temperature, the spins can be disoriented and become energetically unfavorable. I think that's what they meant by thermal instability.

Iris: How does the read current disturb reads? Reads would apply very small voltage.

Karpur: Naively, I would guess that if the read current is also polarized, it could flip a cluster of spins. This would become a bigger issue as you reduce the size of devices (the spins that get flipped will become a larger portion of the cell's spins).

Qishen: Which NAND flash was used for discussion in this particular paper? there are many different NAND flash devices. Single vs. multi-layer NAND flash, for example.
(I think Iris's answer was that this was unclear, and we could discuss in more detail the different materials used for emerging devices in later lectures.)

Jiwon: The paper mentions that the write consumes a lot of dynamic energy, but wouldn't this just be fine, since you have nearly zero static energy?

Iris: agree, and you would want the writes to NVM to happen infrequently.
Collectively, the great advantage of the new NVMs is that the static energy is nearly zero. And if some data is being written frequently, that would belong in the existing SRAM/DRAM.

Sam: Relaxing non-volatility in MRAM -- isn't the point of introducing these new devices to have non-volatility, so that you reduce the static power? If you relax the non-volatility, wouldn't this re-introduce the problems that you tried to eliminate?
(Iris commented that relaxing the non-volatility wouldn't necessarily increase static power, but I didn't get the rest of it..)

Meng-Ju: what's the definition of retention?

Iris: Doesn't lose the value. In non-volatile memory, even if you disconnect the power, the value stored in that memory will be retained.