

RETROSPECTIVE: Use ECP, not ECC, for Hard Failures in Resistive Memories

Stuart Schechter[†] Gabriel H. Loh[‡] Karin Strauss* Doug Burger*
[†]Harvard University [‡]Advanced Micro Devices, Inc. *Microsoft Research

I. THE PROMISE OF RESISTIVE MEMORIES

This work grew out of an investigation into new memory technologies at Microsoft Research starting in 2008. This research was driven by concerns about emerging challenges with DRAM: their ability to continue improving cost/bit, while maintaining reliability while keeping static power dissipation under control. One appeal of resistive memories, which store state as via atomic positioning and material arrangement, is that its scale is not limited by leakage concerns and it does not require refreshing. With the founding of Numonyx to commercialize Phase Change Memory (PCM) in 2008, PCM appeared to be the first resistive memory that might reach scale production. (Micron would announce its acquisition of Numonyx on February 9, 2010, coincidentally one week after the ISCA 2010 rebuttal deadline.)

Resistive memories had their own unique limitations that must be overcome to displace incumbents: rapid cell wear-out, high write latency, high write energy consumption. They also started decades behind DRAM and Flash in manufacturing processes to scale density and performance. Prior work suggested that re-architecting PCM arrays might allow their density and performance to catch up with commercial DRAM DIMMs [1].

Given the performance results, we embarked on a research effort to address reliability. The goal of this research was to ensure that resistive arrays could survive progressive failure of worn-out cells (up to a point), given the significant variance in cell lifetimes. In PCM arrays, no individual cell has a guaranteed minimum lifetime; there is no way to predict which cells would fail until the failures occurred. At the time, one obvious way to increase cell lifetimes was to perform a read/compare/write cycle upon every write, only writing bits that were changing. That approach shifts the lifetime of a resistive cell to be more proportional to the cell datum's entropy. Low-entropy cells would last longer. High entropy cells would fail much more quickly.

This lifetime disparity made it important to balance entropy across the memory cells, to maximize the average time before a critical threshold of cell failures occurred. We quickly realized that (a) error correction would become even more important given permanent cell failures, and (b) that ECC (as typically implemented) was a poor fit as the ECC cells themselves would have the maximum possible entropy.

After working on this problem, it became clear to us that it might be better to *replace* worn cells when they failed than

to *correct* the errors that would result when failed cells were read. The address of a failed cell would need to be written only once, and so replacement would not increase the entropy of writes in the way that correction would. Replacement held promise as a general solution for memory cells that failed permanently (and were detectable), which was true not just of PCM but of resistive memories in general.

With the acquisition of Numonyx in 2010 by Micron, and the subsequent partnership between Intel and Micron to develop resistive memory products (later called 3D crosspoint, or 3DXP, and branded as Optane™ memory), it seemed that resistive memories were on track to become an important part of the computing stack.

II. HIGH-VOLUME COMMODITIES STRIKE BACK

Displacing extremely-high-volume incumbent commodities is incredibly hard because that volume drives massive investment into improved efficiency; manufacturers chase any possible efficiency improvement because, in a large commodity market with single-digit margins, a few percentage points can be the difference between massive profits or losses. Resistive memories faced a steep challenge, and thirteen years later, they have neither displaced DRAM or NAND, nor have they been widely adopted to augment those technologies.

In addition to the continued advance in DRAM density, solid-state storage capacity took a leap with the arrival of 3D NAND. Resistive memories, caught between DRAM and NAND, failed to significantly outperform either technology at their core functions (i.e., they were neither fish nor fowl). As a result, from a system-level perspective, it made little sense to integrate them as a new level in the memory hierarchy. Consequently, although they were released into production as Intel Optane™ technology, resistive memories have largely remained a niche solution, except for very specific use cases. Today, DRAM and NAND still seem dominant for the foreseeable future.

III. ANCILLARY INFLUENCES OF ECP

When we were invited to write this retrospective, we went back and examined the topical categories of research papers citing this paper. Unsurprisingly, the majority of them described subsequent research into resistive memory architectures and technologies. ECP remains an important and influential concept in resistive memory research, which continues to this day.

We were surprised, however, to find that researchers investigating problems in DRAM, GPU register files, and SRAM found some utility in the ECP work. For example, the bit granularity repair method we proposed in the context of resistive memories has been studied recently in the context of systems using DRAM with on-die ECC [2]. Another example is an ICCD’20 paper that referenced an Intel patent that developed a technique called PCLS, or Partial Cache Line Sparing [3]. The paper referred to it as “an extension of ECP” and implied that the technique was shipping in production microprocessors (we have not verified that claim independently). We would not have expected ECP to have any utility for microprocessors’ SRAM caches, so seeing a derivation in that domain was surprising.

Finally, when researchers develop new ideas, they intend to influence others, but they also influence themselves. Working on this paper brought us awareness of new types of materials and media, and their application to memory and storage. This awareness has motivated at least one of us (Karin Strauss) to further examine this space by exploring the use of synthetic DNA to store digital data [4], [5] (also see the retrospective for DNA-based Molecular Architecture with Spatially Localized Components).

IV. LOOKING FORWARD

Thirteen years after the publication of the ECP paper, the computing landscape has seen monumental shifts, throughout cloud and mobile, with generative AI being one of the newest major drivers. This transformation places computation and storage at the core of societal advancement, necessitating rapid technological changes to balance to the enormous computing requirements and costs of cutting-edge AI models with the capabilities they offer. The demands on the memory system are enormous, with the requirement for multiple high bandwidth memory (HBM) stacks per accelerator and terabytes of bandwidth per second pushing the boundaries of what may be possible.

Given these trends, investment in computing and memory technologies will likely continue to accelerate, with the pace of innovation matching or exceeding that urgency. As we continue to drive down to atomic scales, the technologies and solutions may shift, although we caution that today’s technologies may remain dominant with relatively minor shifts. It is unclear to us whether repair (ECP) or correction (ECC) will grow in importance as memory technologies advance. Or, as AI grows as a fraction of aggregate compute, it may be that errors become less important for the AI storage, as AI models have inherent noise tolerance. This relaxed requirement may permit more aggressive memory cell designs where some correctness can be sacrificed for cost, speed, and density. If that occurs, we may find ourselves reading a paper with the title “Use neither ECP nor ECC in AI Memories”. The intersection between density, speed, errors, and cost will continue to be important and interesting.

REFERENCES

[1] Benjamin Lee, Engin Ipek, Onur Mutlu, and Doug Burger. Architecting phase change memory as a scalable DRAM alternative. In *Proceedings*

of the 36th ACM/IEEE Annual International Symposium on Computer Architecture (ISCA), 2009.

[2] Minesh Patel, Geraldo F. Oliveira, and Onur Mutlu. HARP: Practically and effectively identifying uncorrectable errors in memory chips that use on-die error-correcting codes. In *Proceedings of the 54th IEEE/ACM Annual International Symposium on Microarchitecture (MICRO)*, 2021.

[3] Xiaoming Du and Cong Li. DPCLS: Improving partial cache line sparing with dynamics for memory error prevention. In *Proceedings of the 38th IEEE International Conference on Computer Design (ICCD)*, 2020.

[4] James Bornholt, Randolph Lopez, Douglas M. Carmean, Luis Ceze, Georg Seelig, and Karin Strauss. A DNA-based archival storage system. In *Proceedings of the 21st ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2016.

[5] Lee Organick, Siena Dumas Ang, Yuan-Jyue Chen, Randolph Lopex, Sergey Yekhanin, Konstantin Makarychev, Miklos Z. Racz, Govinda Kamath, Parikshit Gopalan, Bichlien Nguyen, Christopher N. Takahashi, Sharon Newman, Hsing-Yeh Parker, Cyrus Rashtchian, Kendall Stewart, Gagan Gupta, Robert Carlson, John Mulligan, Douglas Carmean, Georg Seelig, Luis Ceze, and Karin Strauss. Random access in large-scale dna data storage. *Nature Biotechnology*, 36:242–248, February 2018.