

Storage Class Memory

David Hession, Nigel Mc Kelvey, Kevin Curran¹

Computing Department, Letterkenny Institute of Technology, Donegal, Ireland

¹University of Ulster, School of Computing and Intelligent Systems, Derry, Northern Ireland

¹kj.curran@ulster.ac.uk

Abstract—Advances in processor technology have created myriads of related problems for system designers, chief amongst which is the need to improve on the latency between storage, memory and the processor itself. Significant data transfer rates between the processor, L1 cache, L2 cache and main memory are possible but it is at this point that the scenario changes and it is this problem which must be addressed in the years to come if the gains made in processor technology are to be further translated into overall system performance improvements. Storage Class Memory (SCM) offers just the bridge needed to get over the ever widening gap between CPU processing speeds, the need to move large blocks of data quickly and the read/write speeds offered by traditional disk reliant systems. In terms of their predicted cost, speed, retentive abilities (data persistence) and power requirements, the technologies that can be classed as SCM are the most likely candidates to keep Moore's Law true well into the next decade and beyond. This paper provides an overview of SCM and its role in the 'greening' of data centres around the world. The makeup of data centres with their vast banks of HDD's all-consuming power whether idle or in use is not a sustainable model as the world faces into a future of increasing energy costs and decreasing supply of traditional fuel sources.

Keywords – Storage Class Memory (SCM); Data Storage; Memory; Solid State Devices

I. INTRODUCTION

Storage Class Memory (SCM) offers just the bridge needed to get over the ever widening gap between CPU processing speeds, the need to move large blocks of data quickly and the read/write speeds offered by traditional disk reliant systems. Existing solid state devices such as NAND Flash have gone some way in acting as a bridging technology when used in tandem with traditional HDD's [1], and there are increased performance benefits so that some are believed to be the most usable solution until SCM is fully implemented in a truly stable, reliable form. Inroads into power consumption in both idle and active states are amongst the promised benefits of a full switch to SCM [2]. HDD manufacturers such as Seagate have developed solid state hybrid drives which incorporate both technologies and see this as the medium term solution until the issues that plague SCM can be overcome [3]. This type of development is sure to extend the shelf life of the HDD further but will it be enough to save it from the benefits promised by the emerging SCM technologies? Further improvements to NAND Flash are also touted as a possible SCM and this is discussed in the main body of this document.

The broadest definition of what can be classed as storage class memory is a technology that is non-volatile, cheap in a per bit cost, has fast access times for both read and writes and is solid state (has no moving parts). Some of the technologies discussed in this document are still primarily in the research and development stage of their existence. A few of the technologies such as Phase Change Memory (PCM), Magnetic Random Access Memory (MRAM) and Ferroelectric RAM (FRAM) already have working solutions in production but are by no means widely accepted and should be researched and developed further. Of the technologies still on the drawing board, some are at a more advanced stage than others; those that are in the earliest stages of development have theoretical possibilities which could not be understood at the moment clearly, if realised, it means a quantum leap forward in how computers are designed not only from the hardware side but also from the software point of view.

All types of SCM discussed use different techniques to manipulate their state to produce a readable binary bit. The advances in the use of nanotechnologies have contributed to SCM in boosting the available storage capacity possible by factors way beyond the scale that HDD's can offer. It is when the cost of this bit/cell ratio drops to the levels even close to that of the HDD that SCM will start to emerge as the preeminent storage technology. The technology that makes this leap first and can be proved to be as consistent in its performance as those existing technologies such as the HDD and NAND Flash has the greatest chance of widespread adoption. Widespread adoption will mean more funding directed towards further development and may see the other SCM as a consequence fall by the wayside. Huge efforts and significant funding have been directed at SCM by many giants of the computing world (IBM, Toshiba, HP) and it is still not clear which technology is most likely to be the winner. The following section briefly outlines each of the technologies vying to become the accepted SCM.

II. STORAGE CLASS MEMORY TYPES

A. Improved Flash

SCM is not a type of Flash memory although all proposed types of SCM are solid state devices in common with Flash. The potential of Flash technologies seems to be reaching their limits in terms of scalability although new Flash/HDD combinations are on the market which will see the technology survive in the medium term [4]. The viability of a greatly scaled NAND Flash

is questionable, whilst it is technically possible to improve the overall scale of the available Flash modules, but this can only be done by compensating in terms of speed and data persistence. The attrition caused by read/write operations also makes Flash increasingly error prone as it ages, not a good solution for the storing of potentially valuable data. Whilst Flash technologies are widespread (USB pen drives, digital camera and mobile phone memory), their flaws are also widely known and it is these flaws which means that its future seems to be as a technology of the past rather than of the future.

B. Phase Change Memory

Phase Change Memory (PCM) takes advantage of the properties of chalcogenide when heated [5]. Chalcogenide goes through distinct phases which can be used to represent the binary 1 and 0; indeed IBM has identified 4 distinct states which can be used for that representation thereby doubling the potential storage capacity of a given memory cell [6]. Distinguishing between these states without actually disturbing them was a major advancement for those researching PCM. Its advantages over NAND Flash lie in its bit-alterability (there is no erase phase involved in the write operation), its fast access times and its potential scalability. There are issues surrounding the temperature sensitivity for PCM but these haven't stopped PCM going into production. Whether or not it will be widely adopted may hinge on just how scalable the technology proves to be and the cost of overcoming issues around its heat sensitivity.

C. Magnetic RAM

Magnetic Random Access Memory (MRAM) is a method of storing data bits using magnetic charges instead of the electrical charges that are used by conventional memory types such as Dynamic RAM (DRAM). Unlike DRAM, MRAM is non-volatile, but its magnetic properties are retained when the power is switched off and are remembered again when the power supply is switched back on. This remembered state could mean computers that start instantly with no bootup required [7]. Operational power consumption is close in many respects to that of DRAM for read operations but is significantly higher when it comes to writes, which is one of MRAM's main drawbacks. Overall achievable speeds for read/write operations are close to that of DRAM but even given its many apparent qualities uptake of MRAM has not been as widespread as may be expected [8]. Its scalability is questionable and it is on this characteristic that its future may depend. MRAM though is highly durable and it may see uptake in areas where this quality is essential. Racetrack memory falls into the category of a magnetic manipulative memory type and it also promises huge advances if its potential is realised. Racetrack memory is still largely in theory and has a long way to go from being implemented to any great degree.

D. Solid Electrolyte RAM /Nano-Ionic RAM

The memory cells that make up a Nano-Ionic RAM chip contain a solid glasslike electrolyte layer compressed between two metal electrodes. When a voltage is applied to the electrolyte layer via the electrodes, it causes electrons to bind to the metal ions the electrolyte layer contains, then these ions are fused together to form metal atoms which in turn join to form a thin wire (nanoscale) which bridges the gap between the two electrodes allowing the current to pass. Reversing the current has the reverse effect causing the wire to break up. These states can be used to represent the binary 1 and 0 [9]. This type of RAM generates less heat in its operation, uses significantly less power than traditional RAM modules and has significant gains available in terms of speed and endurance [10]. Its drawbacks have been identified as being around data persistence and the difficulties that arise in its fabrication where temperature sensitivity is a problem.

E. Ferroelectric RAM

Each of the cells that make up the structure of an Ferroelectric RAM (FRAM) chip contains a ferroelectric material commonly known as PZT (lead zirconate titanate), it is the manipulation of the elements in this layer that gives FRAM its retentive qualities. FRAM relies on a scheme of switching the polarity of elements contained in its cell structure to produce what is termed as a binary switch [11]. Proponents of FRAM claim significant gains in terms of power consumption, endurance and write speeds. The technology is already in production and is used in devices such as printers, automotive brains [12].

F. The Memristor

The memristor [13] (memory resistor) was first proposed by Leon Chua in 1971[14] but it is only with the real advances in nanoelectronics made in the last decade when the potential of memristance has come to the fore. Memristance at its simplest form is the ability of an electric component to be in one state if electrical current is applied in one direction and to change its state if the current is reversed; more importantly, it is the ability to 'remember' that state when the current is switched off that makes memristance a possible candidate as a type of SCM. HP announced in August 2010 that it was moving the memristor off the drawing board and onto the production floor and hoped to have a working model of what it terms Resistive Random Access Memory (ReRAM) for general release by 2015[15]. It hopes this ReRAM will become the choice of SCM and touts its low power consumption, data persistence and versatility. HP claims their memristive technology can also be used to perform logic calculations in almost the same manner as a processor and this memristive processor has the potential to change the face of computing entirely. Greater processing speeds(beyond what is envisaged in Moore's Law), less power consumption and the state saving capacity of the technology(meaning computers can be switched on and off like light switches)along with the

scalability of the technology make the memristor one of the most exciting SCM technologies outlined [16].

Organic and polymeric memory also offers theoretical possibilities as SCM. There are several approaches being researched at the present time but these technologies are very much on the drawing board and have a long way from being realised as a working solution.

III. METRICS FOR ASSESSING SCM

A. Cost

Cost at this stage of the evolution of SCM is probably not a fair criterion to judge upon but is the chief driver in deciding whether SCM will flourish in the short or medium term. The pre-eminent storage technology in use today is the HDD; it is cheap in relative terms and offers large amounts of storage capability even in the most basic of computer systems. Performance is the hidden cost for HDD's, which is where SCM should be able to make significant breakthroughs in the years to come. It is predicted that the production costs of SCM and HDD will be broadly similar by the end of this decade and once this state of parity is reached the decline of HDD as a frontline storage medium is inevitable. Overall performance is the cost metric which will be the chief reason for that decline at that stage instead of the cost in terms of euro and cent.

B. Speed

Speed is a metric upon which SCM is streets ahead of even the very best HDD available today, the promised performance benefits of SCM can be assessed in factors of 10 over the HDD and that is in a worst case scenario, at optimal performance levels many of the SCM candidate technologies promise read/write performance which is in factors of 100s better than HDD and even beyond that again. SCM offers data transfer speeds similar to DRAM and perhaps better in some cases.

C. Persistence

Data persistence is an unproven factor for SCM at this stage of its development; theoretically many of the technologies offer the promise of good data persistence at low power. From this point of view, the HDD has proven capabilities and the issues around its capabilities and limitations are well understood. NAND Flash is also hamstrung from this point of view with data persistence being one of its major drawbacks, and degradation of storage cells in read/write operations means that consistent long term persistence of data cannot be promised. Clearing this particular hurdle is seen as a key in the development of SCM and the technology that grasps this nettle first will have made a major leap forward.

D. Power

Power consumption and the heat generated by the use of power are major areas of concern in the development of SCM and indeed all computing components. The efficient use of power in both active and idle states will also be a major criterion for judging SCM by; as alluded to in the introduction the 'greening' of computing especially, the vast data centres that make up the growing cloud is seen as essential. Any adopted SCM will have to operate within the constraints that this 'greening' enforces. The cost of power consumption is too great a prerequisite to be ignored for any viable storage solution in the years to come. Other metrics which may be used to judge SCM will include its read/write speeds, its overall data persistence and its durability.

IV. CONCLUSION

Storage Class Memory is the ultimate future of both memory and storage in computing systems; the technologies that fall into the SCM fold have blurred to such a degree the apparent differences between memory and storage that they are causing a complete rethink in how systems should be designed. The possibilities offered in cost savings, performance improvements, data persistence and overall endurance will see a paradigm shift in computing terms upon the widespread implementation of SCM. Each of the technologies is vying to become the next best thing and it is unclear at this moment in time as to which will actually become the SCM of choice. The huge costs involved in the setup of chip production frameworks are holding companies back from plumping one way or the other; many companies such as IBM and Samsung are hedging their bets by pouring significant sums into the R&D of a number of SCM. What is clear is that SCM will eventually come into mainstream production and that computing will change fundamentally because of it.

The Cloud infrastructure that is so heavily reliant on huge data centres will not be sustainable in its current form, the need to downscale in terms of running costs (e.g. power, cooling, floorspace) whilst up-scaling in terms of efficiency and computing power will be a huge driver in its future direction. SCM offers the chance to bridge the gap between cost and performance; massive savings will be possible in overall running costs whilst at the same time taking advantage of the benefits that SCM offers in terms of data access speeds and data endurance. Moore's Law has held true for many decades now but SCM offers the potential if fully utilised to see Moore's Law rethought and perhaps rewritten in the decades to come.

REFERENCES

- [1] Coughlin, T; Handy, J. (2011). *Two May Be Better Than One: Why Hard Disk Drives and Flash Belong Together*. <https://www.snia.org/sites/default/files/Storage%20Pairing%20WP%20FEB%202011.pdf>.
- [2] Philhower, E. (2009). *Using Storage Class Memory for Extreme Performance In-Memory Computing*. http://www.percona.com/files/presentations/ppc2009/PPC2009_using_storage_class_memory.pdf.
- [3] Virwani, K., Burr, G., Shenoy, R., Rettner, C., Padilla, A., Topuria, T., Rice, P., Ho, G. King, R., Nguyen, K., Bowers, A., Jurich, M., BrightSky, M., Joseph, E., Kellock, A., Arellano, N., Kurdi, B. and Gopalakrishnan, K. (2012) "Sub-30nm scaling and high-speed operation of fully-confined Access-Devices for 3-D crosspoint memory based on Mixed-Ionic-Electronic-Conduction (MIEC) Materials," IEDM Technical Digest, vol. 2, no. 7, pp. 22-34.
- [4] Burr, G., Virwani, K., Shenoy, R., Rettner, C., Padilla, A., Topuria, T., Rice, P., Ho, G. King, R., Nguyen, K., Bowers, A., Jurich, M., BrightSky, M., Joseph, E., Kellock, A., Arellano, N., Kurdi, B. and Gopalakrishnan, K (2012) "Large-scale (512kbit) integration of Multilayer-ready Access-Devices based on Mixed-Ionic Electronic-Conduction (MIEC) at 100% yield," Symposium on VLSI Technology, vol. 5, no. 4, pp. 44-52.
- [5] Mills, P. (2009). *Storage Class Memory - The future of Solid State Storage*. https://www.snia.org/sites/default/education/tutorials/2009/fall/solid/PhilMills_The_Future_of_Solid_State_Storage.pdf.
- [6] Numonyx. (2010). *The basics of phase change memory (PCM) technology*. Available: <http://www.signallake.com/innovation/PhaseChangeMemory.pdf>.
- [7] Shenoy, R., Virwani, K., Burr, G., Rettner, C., Padilla, A., Topuria, T., Rice, P., Ho, G. King, R., Nguyen, K., Bowers, A., Jurich, M., BrightSky, M., Joseph, E., Kellock, A., Arellano, N., Kurdi, B. and Gopalakrishnan, K. (2012) "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, Paris, France, May 3-5 2012, pp. 12-22.
- [8] Raoux, S., Burr, G., Breitwisch, M., Rettner, C., Chen, Y., Shelby, R., Salinga, M., Krebs, D., Chen, S., Lung, H., and Lam, C. (2008) "Phase-change random access memory — a scalable technology," IBM Journal of Research and Development, vol. 52, no. 4, pp. 465-475.
- [9] Wang, B. (2007). *PMC memory*. Available: <http://nextbigfuture.com/2007/10/programmable-metallization-cell-pmc.html>.
- [10] Byrne, S. (2007). *University develops PMC memory, a potential Flash killer*. Available: <http://www.myce.com/news/University-develops-PMC-memory-a-potential-Flash-killer-13837/>.
- [11] Burr, G., Kurdi, B., Scott, J., Lam, C., Gopalakrishnan, K., and Shenoy, R. (2008) "An overview of candidate device technologies for Storage-Class Memory," IBM Journal of Research and Development, vol. 52, no. 5, pp. 449-457.
- [12] Freitasand, R., Wilcke, W. (2008) "Storage Class Memory, the next storage system technology," IBM Journal of Research and Development, vol. 52, no: 4, pp. 439-448.
- [13] Williams, S. (2011). *HP Memristor - FAQ*. Available: http://www.hpl.hp.com/news/2008/apr-jun/memristor_faq.html?jumpid=reg_R1002_USEN.
- [14] Chua, L. (1971). "Memristor - The missing circuit element." IEEE Trans. Circuit Theory, vol. 18, no. 3, pp. 507-519.
- [15] HP. (2010). *HP Collaborates with Hynix to Bring the Memristor to Market in Next-generation Memory*. Available: http://www.hp.com/hpinfo/newsroom/press/2010/100831c.html?jumpid=reg_R1002_USEN.
- [16] 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 (2006) "Ultra-Thin Phase-Change Bridge Memory Device Using GeSb," IEDM Technical Digest, vol. 3, no. 3, pp. 102-114.

David Hession (BSc) is a graduate of computing from the Letterkenny Institute of Technology. He is currently working in the IT sector. His research interests include disk technologies, computer networks and software.

Nigel McKelvey (MSc., BSc., PGCE, MICS) is a lecturer in Computing at the Letterkenny Institute of Technology and specialises in teaching secure programming techniques at both undergraduate and postgraduate level. Other areas of interest include ethical gaming, performance based programming and digital forensics. Nigel is currently completing a Doctorate in Education focusing on adopting heuristic programming techniques within the classroom.

Kevin Curran BSc (Hons), PhD, SMIEEE, FBCS CITP, SMACM, FHEA is a Reader in Computer Science and group leader for the Ambient Intelligence Research Group. Dr Curran has made significant contributions to advancing the knowledge of computer networking evidenced by over 800 published works. He is a regular contributor to BBC radio & TV news in the UK and quoted in trade and consumer IT magazines on a regular basis. He is an IEEE Technical Expert for Security and a member of the EPSRC Peer Review College.