EIPBN 2008—The 52nd International Conference
on Electron, Ion, and Photon Beam Technology and Nanofabrication

invited talk

Storage Class Memory

Geoffrey W. Burr, Kailash Gopalakrishnan,
Rohit S. Shenoy, Charles T. Rettner, and Bülent N. Kurdi

IBM Almaden Research Center
650 Harry Road, San Jose, California 95120
Tel: (408) 927–1512 FAX: (408) 927–2100 *burr@almaden.ibm.com*

Storage-Class Memory (SCM) combines the benefits of a solid-state memory — such as high performance and robustness — with the archival capabilities and low cost of conventional hard-disk magnetic storage. Such a device would require a solid-state non-volatile memory technology that could be manufactured at an extremely high effective areal density, using some combination of sub-lithographic patterning techniques, multiple bits per cell, and/or multiple layers of devices.

In a sense, initial SCM products are already appearing on the market, as solid-state disks (SSD) built around NAND Flash arrive in products ranging from laptops to enterprise storage solutions. That this is occurring — despite the disadvantages of Flash in terms of write endurance and cost-per-bit with respect to hard-disk drives (HDD) — is a testament both to the cleverness of system designers as well as to the increasing importance of read bandwidth and latency. At the same time, significant challenges are looming for the Flash industry. The relentless march of scaling (shrinking of memory elements in near-future device generations) is projected to reduce the number of stored-electrons-per-bit in Flash devices to untenably low values.

These trends represent a tremendous opportunity for an alternative memory technology. The bar is high — any new SCM technology would have to offer an attractive combination of further scalability (to smaller device generations), low cost-per-bit, high endurance, and high performance (e.g., memory speed). That said, the size of this opportunity goes well beyond simply providing a potential successor for Flash technology, despite the vast number of applications Flash currently addresses. In fact, the emergence of a non-volatile solid-state "Storage Class" memory technology that combined high performance, high density and low-cost could usher in seminal changes in the memory/storage hierarchy throughout all computing platforms, from tiny mobile devices up to high-performance computing.

Any one of a number of solid-state non-volatile memory technologies could potentially be used to construct such a Storage-Class Memory [1]. Candidates range from evolutionary extensions of conventional Flash memory, to revolutionary new memory technologies such as ferroelectric random-access memory (FeRAM), magnetic RAM (MRAM), organic and polymeric memory, resistive RAM, and solid-electrolytes. However, recent developments would seem to have tilted the balance, in terms of potential for practical scaling to ultra-high effective areal density, towards phase-change (PCRAM) random-access memory [2].

[1] G. W. Burr, B. N. Kurdi, J. C. Scott, C. H. Lam, K. Gopalakrishnan, and R. S. Shenoy, "An overview of candidate device technologies for Storage-Class Memory," to appear in *IBM J. Res. Dev.* 2008.

[2] S. Raoux, G. W. Burr, M. J. Breitwisch, C. T. Rettner, Y. Chen, R. M. Shelby, M. Salinga, D. Krebs, S. Chen, H. Lung, and C. H. Lam, "Phase-change random access memory – a scalable technology," to appear in *IBM J. Res. Dev.* 2008.