Maximizing solid-state storage capacity in small form factors

Kent Smith, Senior Director of Marketing, Flash Components Division, LSI

Users want ever-smaller and lighter devices but also demand ever-increasing storage capacity to keep more apps and data loaded on their mobile computing platforms. To accommodate these two competing objectives, solid-state storage form factors will need to get smaller, while NAND flash memory geometries will be shrinking and storing more bits per cell. The combination is having an impact on the way flash memory is being designed into ultrabooks, netbooks and other mobile computing devices.

The first consideration in designing for maximum capacity is the form factor of the printed circuit board (PCB) for the storage components. The latest storage form factors being standardized are known as M.2 (previously called the next generation form factor or NGFF). As shown in Figure 1, the most popular M.2 form factor among system manufacturers is 40 percent smaller than the mSATA card. In addition to being more compact, the M.2 specification has been optimized for solid state storage and includes connector keys for SATA, 2x or 4x PCI Express.

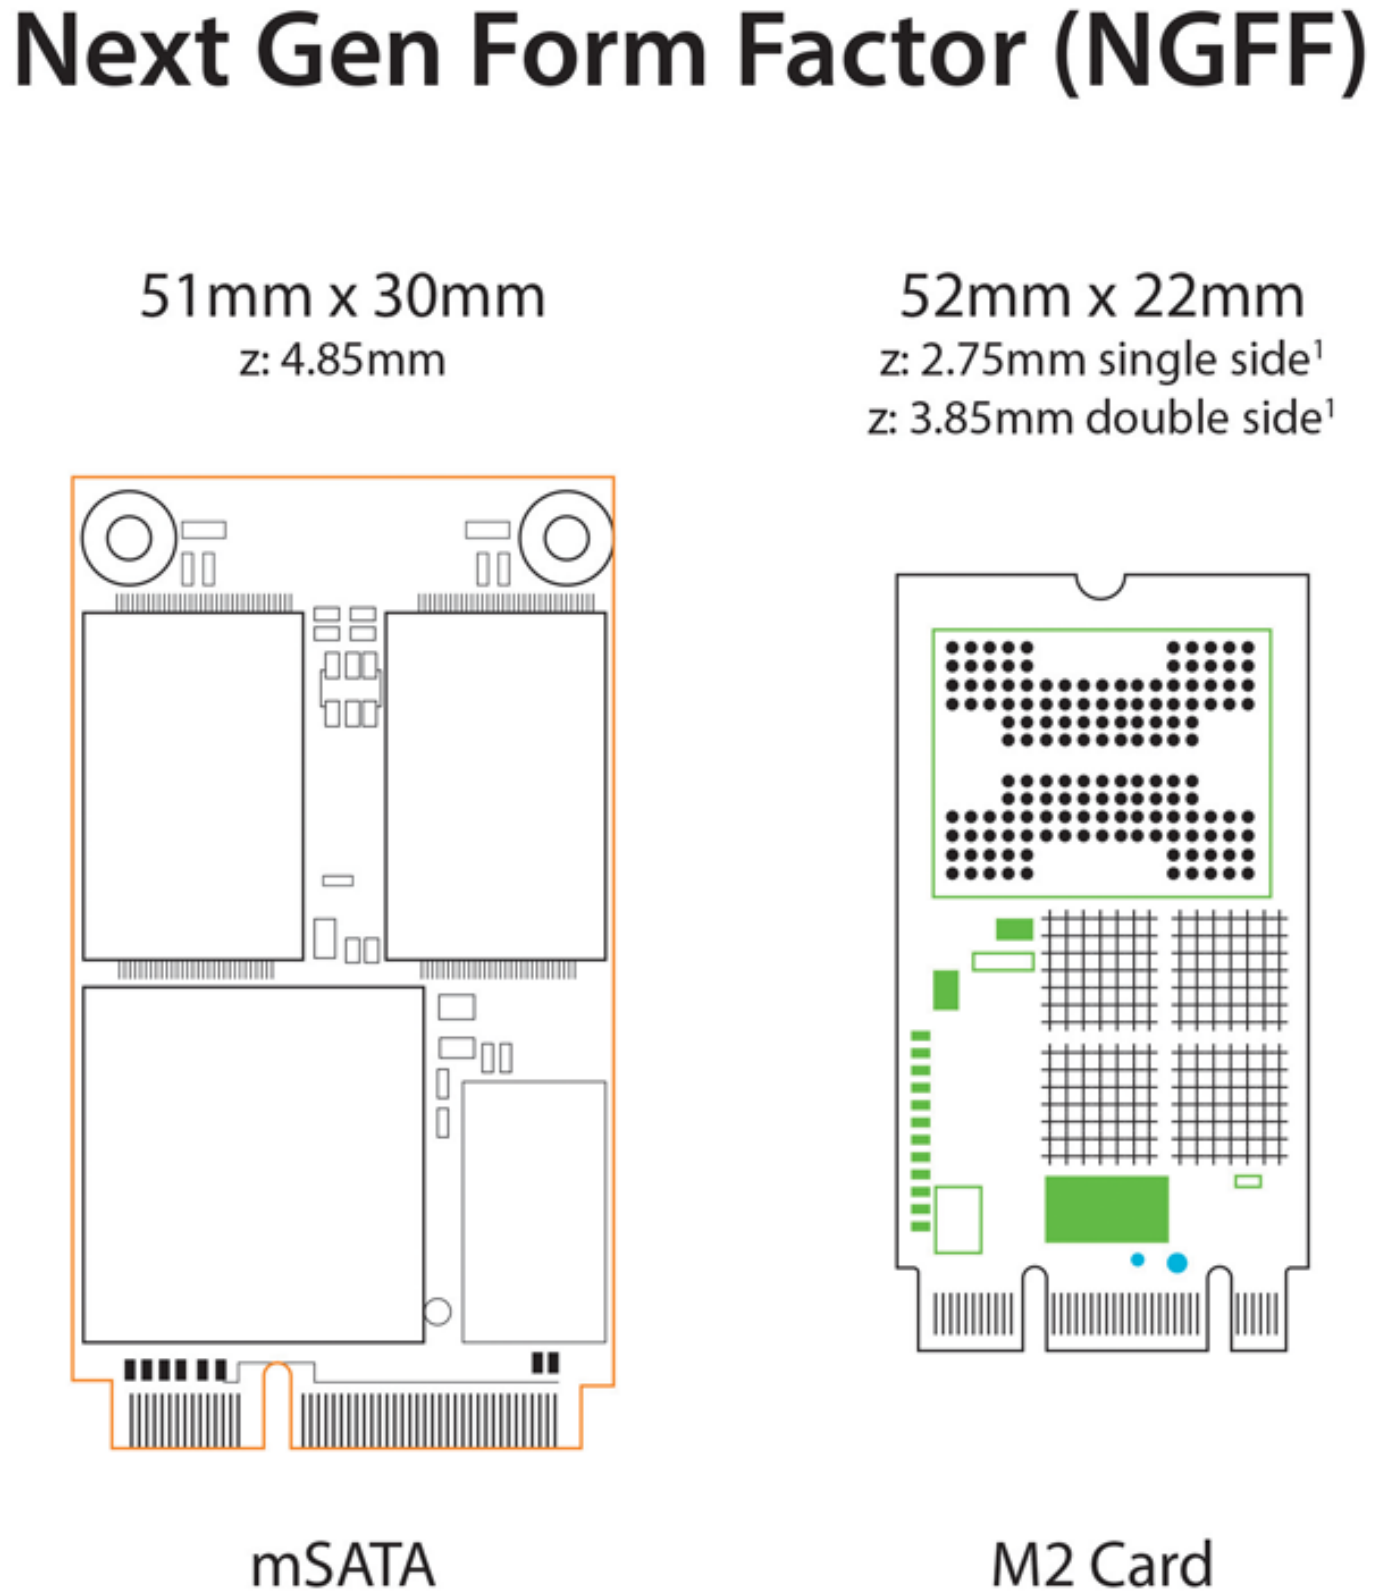

For applications where additional capacity is required (and space is available), the M.2 specification supports other card dimensions, including some with lengths up to 110 mm, providing nearly 60 percent more area than mSATA. There are also other custom and proprietary designs that include stacking multiple flash memory packages or using multiple PCBs that are much taller in the z-height dimension of the base PCB, but reduce the overall footprint by decreasing the aggregate cubic volume.

The smaller area available on the M.2 card is driving the need for using smaller

flash memory geometries and/or more bits per cell. As shown in Figure 2, the combination has dramatically increased the density of storage possible. For example, in the same footprint, 50 nm flash using single-level cells (SLC) can store only 2 Gigabytes (GB), while 19 nm flash using multi-level cells (MLC) can store 32 GB—16 times the density for approximately the same cost. With triple-level cells (TLC), also at 19 nm, the same footprint could have a capacity as high as 48 GB.

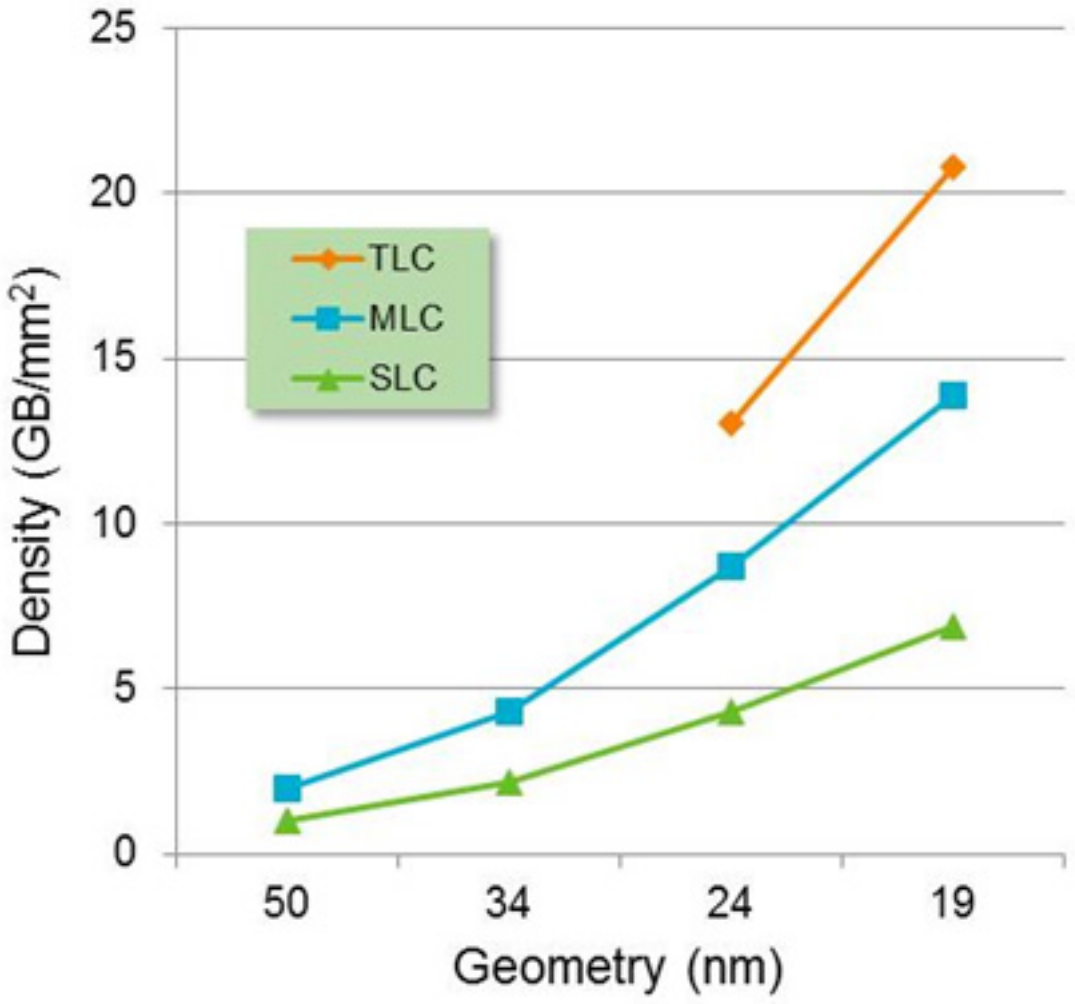

Next-generation flash storage processors

Taking full advantage of shrinking geometries and higher bit densities of NAND flash memory requires some changes to flash storage processors (FSP). The FSP is responsible for managing the pages and blocks of flash memory, and also provides the input/output (I/O) interface with the system. Two of the biggest challenges for FSPs today involve error correction and endurance.

As flash memory geometries shrink, cells become smaller and, therefore, hold less of a charge for the one, two or three bits they store. For illustrative purposes imagine a 50 nm cell storing a single bit, which might hold about 1000 electrons, and a 20 nm cell storing two bits, which might hold only 100 electrons—an order of magnitude fewer. While the number of electrons cited here does not reflect actual measurements, the comparison does demonstrate that the lower charge available with fewer electrons increases the potential for read errors from the flash, which must be corrected by the FSP.

Traditional approaches to error correction, such as Reed-Solomon (RS) or BCH (also named for its co-inventors Bose, Ray-Chaudhuri and Hocquenghem), are giving way to the Low-Density Parity Check (LDPC) in next-generation FSPs. LDPC can provide error correction performance close to the theoretical limits of any technique. Adding sophisticated digital signal processing enables detection and correction of even more errors. The few errors that cannot be corrected could then be handled by an integral data protection technology, much like the RAID (redundant array of independent disks) technology used in direct-attached storage and storage area network controllers.

Higher density flash cells with higher error rates wear out sooner. For this reason, the garbage collection and wear-leveling capabilities of the FSP have become increasingly important. The need for garbage collection and wear-leveling in NAND flash causes the amount of data being physically written to flash memory to be a multiple of the logical data intended to be written. This phenomenon is expressed as a simple ratio called “write amplification,” which ideally would approach 1.0. Because these “unnecessary” writes wear out cells prematurely, next-generation FSPs will benefit greatly from some type of data reduction technology to minimize write amplification and, thereby, maximize the flash memory’s useful life.

Another technique for increasing capacity is to eliminate the need for a separate DRAM buffer, which is required in solid state storage solutions to maintain the “map” consisting of a combination of the flash memory file index and logical block addresses (LBAs). But the DRAM chip consumes precious space and power that could (and should) be used for more flash memory. DRAM-less chip designs, such as the LSI SandForce FSP, are also key to enabling SSD manufacturers to develop higher capacity drives for today’s growing class of thin-and-light ultrabook platforms. By creating designs that do not require an external DRAM buffer, these next-generation single-chip FSPs are what will make it possible to maximize solid state storage capacity in small form factors.

About the author

Kent Smith is senior director of Marketing for the Flash Components Division of LSI Corporation, where he is responsible for all outbound marketing and performance analysis. Prior to LSI, Smith was the senior director of Corporate Marketing at SandForce, which was acquired by LSI in 2012, his second company to be sold to LSI. He has over 25 years of marketing and management experience in the storage and high-tech industry, holding senior management positions at companies including SiliconStor, Polycom, Adaptec, Acer and Quantum. Smith holds an MBA from the University of Phoenix.