

FLASH MEMORY'S PER-MEGABYTE COST IS PLUMMETING, AND ITS PER-CHIP DENSITY IS CLIMBING. YOUR CUSTOMERS' DEMANDS FOR SYSTEM PERFORMANCE, LOW POWER CONSUMPTION, COMPACT FORM FACTORS, LIGHT WEIGHT, RUGGEDNESS, AND RELIABILITY ARE UNRELENTING. IS IT TIME TO CONSIDER AN UPTICK IN YOUR PER-SYSTEM SOLID-STATE-CAPACITY BUDGET?




Figure 1 USB add-on "sticks" (a), embedded mini-PCI Express modules (b), and solid-state drives (c) are three of the many ways that high-density flash memory is muscling into computing platforms (courtesy Kingston, Intel, and Samsung, respectively).

SOLID-STATE STORAGE: FEASIBLE PLAN OR FLASH IN THE PAN?

BY BRIAN DIPERT • SENIOR TECHNICAL EDITOR

Flash memory gained its first foothold in the PC architecture in 1991 with manufacturers' adoption of user-upgradable BIOS as a successor to ROM- and EPROM-based predecessors. Since then, incremental flash-memory-usage success has come through the use of solid-state-storage modules in a variety of form factors and system interfaces: USB, PCMCIA, CompactFlash, Secure Digital, and others. However, unlike BIOS, these additional flash-memory beachheads have been predominantly transitory, with users employing the solid-state storage as backup for the primary system hard-disk drive, for example, or as a means of interchanging data between two PCs or between a PC and another system, such as a PDA or a digital camera.

Over the past several years, however, industry interest has blossomed and accelerated in boosting the amount of flash memory in a system from multiple megabits to multiple megabytes—and even to tens of gigabytes (Reference 1). The flash-memory suppliers have to some degree cultivated this density-expansion attention; with the flattening out of consumer demand for traditional, high-density-flash-memory applications, such as cell phones, digital cameras, and PDAs, and with substantial incremental flash-memory supply becoming available, the memory vendors are eager to cultivate new markets (Reference 2). Hard-drive suppliers, similarly,

are looking to augment—but not necessarily replace—their products with flash memory. They see this augmentation as a means of differentiating—and thereby garnering incremental revenue and, they hope, profits from—what has otherwise become a highly commodity-based rotating-magnetic-storage business.

Microsoft, too, is leveraging its Windows Vista's embrace of high-capacity flash memory in multiple implementation variations as a means of encouraging consumers to adopt the company's latest operating-system version. Supply-side cynicism aside, computer users can garner a number of potential advantages by migrating to high-density-flash-memory-inclusive computers. Those of you who are designing and using embedded systems may already be aware of—and exploiting—these advantages (see sidebar “Non-PC applicability”).

READYBOOST ME?

From Microsoft's earliest DOS days with the SmartDrive utility, the company has striven to employ semiconductor memory as a fast cache for the slower hard-disk drive behind it. Until the emergence of Windows Vista, DRAM was the only applicable cache-technology candidate, and each incremental

AT A GLANCE

Operating-system enhancements enable flash memory to act as a hard-disk-drive cache, an attractive scenario when you cannot upgrade the system's DRAM allocation.

More encompassing caching support, supporting system-boot and return-from-hibernation scenarios, comes both integrated within the hard-disk drive and in external-to-hard-drive flash memory.

Replacing magnetic rotating storage with a solid-state alternative is cost-effective only if your required amount of flash-memory storage is below the price “floor” of the hard-drive alternative.

Regardless of the high-density flash-memory implementation, availability and cost concerns bear careful consideration before you proceed.

Embedded systems have in many cases already blazed the hard-disk-drive-to-solid-state-drive trail that computers are now beginning to traverse.

Windows version improved the effectiveness of that cache. Vista's SuperFetch is the most recent iteration of Windows'

NON-PC APPLICABILITY

The magnetic-to-solid-state-storage trend is old news to those of you in noncomputer applications who have lived through the transition by virtue of the fact that, for example, your equipment's storage requirements are far less substantial than those of a PC. A dual-flash-memory-technology combination is one common system-implementation approach, employing a NOR device for initial boot code and NAND-flash memory as a solid-state drive. Either a dedicated hardware controller or flash file-system software from Datalight or another vendor manages these NAND devices.

If you prefer to boot directly from a NAND-flash memory, you must comprehend potentially “known-bad” blocks in your system memory-mapping scheme. Designers

tolerate these blocks because they boost the effective device yield. You also must compensate for the fact that NAND-flash memory's random-access latency is longer than that of NOR devices. Some NAND-flash memories embed a small SRAM array along with a NAND-to-SRAM autotransfer controller and system-microprocessor hold-off-control signals. Alternatively, you can implement a NAND-to-NOR control-signal translator and wait-state generator in a discrete PLD, an FPGA, or another logic device (Reference A).

REFERENCE

A “QuickLogic Storage Solution Portfolio Runs the Gamut with Addition of NAND Flash Controller,” QuickLogic, Feb 5, 2007, <http://ir.quicklogic.com/releasedetail.cfm?ReleaseID=228315>.

precaching. Its improvements over the Windows XP-based predecessor include built-in intelligence that suppresses the effects of background tasks, such as hard-disk-drive-backup and virus-scanning sessions, that might otherwise result in unnecessary and undesirable cache “thrashes.” SuperFetch, according to Microsoft, can also “learn” from and adapt its behavior to a user's PC-operating patterns, not only for global-program preferences, but also for more finely grained time-of-day and day-of-week tendencies.

However, as Microsoft's Web site notes, “Adding system memory (typically referred to as RAM) is often the best way to improve a PC's performance, since more memory means more applications are ready to run without accessing the hard drive. However, upgrading memory can be difficult and costly, and some machines have limited memory-expansion capabilities, making it impossible to add RAM” (Reference 3). As a result, Windows Vista's SuperFetch supports hard-disk-drive caching not only to traditional DRAM, but also to supplemental flash memory, which requires no power-consuming constant refreshing as DRAM does. The company has tagged this approach with the ReadyBoost marketing moniker. This memory may reside on the motherboard. Examples of this approach include Intel's (formerly, Robson Technology's) Turbo Memory mini-PCI Express module and AMD's HyperFlash array. Alternatively, ReadyBoost can take the form of an add-in module. Thanks to SuperFlash's built-in, lossless 2-to-1 average compression and AES (Advanced Encryption Standard)-128 algorithm, Microsoft's hard-disk-drive-caching scheme makes secure and efficient use of the available flash-memory storage.

If you're thinking that all this news sounds too good to be true, here are a few caveats. Windows Vista discards information cached in a removable flash module upon exit from system hibernation, because the operating system can't tell whether the user has ejected the module and, therefore, potentially altered its contents during the system's hibernation interval. Flash memory's read and, especially, write performance is slower than the DRAM-implemented cache alternative. Random reads from

flash memory, by virtue of their avoidance of rotating storage's rotational and seek latencies, are faster than those from a hard-disk drive. However, sequential-read-performance comparisons, along with both random- and sequential-write-performance evaluations, don't obviously favor flash memory. These assessments may, in fact, tip the scales in the hard-disk drive's favor, depending on the two technologies' system interfaces and the amount and effectiveness of the hard drive's built-in buffer RAM. As a result, Microsoft requires flash memory for ReadyBoost to satisfy stringent capacity and speed criteria, and Windows Vista won't activate a flash-memory module for ReadyBoost use unless the OS detects that the flash memory satisfies the following conditions (**Reference 4**):

- USB flash drives themselves and the host controllers into which you insert them must both use the USB 2.0 standard.

- USB flash drives must have at least 230 Mbytes of free space.

- Flash-storage devices must meet minimum performance requirements to support ReadyBoost, including 2.5-Mbyte/sec throughput for random 4-kbyte reads and 1.75 Mbytes/sec for random 512-kbyte writes.

- Devices must meet higher performance requirements to achieve the "enhanced-for-ReadyBoost" designation. These requirements are 5 Mbytes/sec for random 4-kbyte reads and 3 Mbytes/sec for random 512-kbyte writes. Microsoft recommends that users employ en-

TABLE 1 READYBOOST-PERFORMANCE REQUIREMENTS

Scenario	Sustained throughput (Mbytes/sec)
4-kbyte random reads to a file pinned in the nonvolatile cache	More than 4
4-kbyte random writes to a file pinned in the nonvolatile cache	More than 4
64-kbyte sequential reads to a file pinned in the nonvolatile cache	More than 16
64-kbyte sequential writes to a file pinned in the nonvolatile cache	More than 8

hanced-for-ReadyBoost devices to take full advantage of the benefits ReadyBoost provides.

- ReadyBoost can support devices with cache sizes ranging from 230 Mbytes to 4 Gbytes. Microsoft recommends at least a 1-to-1 ratio of ReadyBoost cache size to main memory (RAM) capacity. Ratios as high as 3-to-1 realize the optimal performance benefits.

Although many flash-memory "sticks" meet Microsoft's ReadyBoost requirements, they're inconvenient and unaesthetic for constantly in-use hard-disk-drive-caching purposes because they protrude from the front, back, or side of the computer, depending on where its USB ports reside. Alternatively, SecureDigital, ExpressCard, and other memory-card formats provide a means of safely cradling the ReadyBoost device within a computer's integrated card-reader slots. For desktop PCs, motherboard-resident USB connectors provide self-contained-system routes to implementing ReadyBoost. And mini-PCI Express connec-

tors are also ReadyBoost candidates for both mobile and desktop computers. However, by populating these connectors with flash memory, you preclude their use for wireless connectivity or other functions. This issue is of particular concern with limited-I/O chip sets or small-form-factor motherboards, because both factors preclude the inclusion of more than one mini-PCI Express slot.

Theoretical capabilities aside, the near-term prognosis on ReadyBoost is mixed at best, although Microsoft released Windows Vista

patches just before this article went to press. Those patches may more definitively deliver the performance boost that Microsoft has been promising for several years now. The ReadyBoost algorithms strive to take maximum advantage of flash memory's strengths and circumvent nonvolatile semiconductors' limitations. To do so, Microsoft's engineers tuned the ReadyBoost code using research-defined assumptions about how consumers typically use computers. If a given tester's synthetic or "real-life" benchmark radically detracts from Microsoft's vision by extensively writing to the hard drive, for example, ReadyBoost's improvements will be less evident than they would be otherwise, and, in the worst case, the system may even run more slowly with ReadyBoost turned on (see sidebar "Additional materials and a hands-on test-drive"). From the numerous ReadyBoost-testing results I've perused so far, four key conclusions regularly appear:

- ReadyBoost-performance improvements are most evident in mobile-computing environments, which tend to use slower hard drives than their desktop counterparts as a means of reducing power consumption.

- ReadyBoost particularly shines from a power standpoint with mobile PCs, because the scheme enables the hard drive to remain in its lowest power state more often. Conversely, with ac-powered desktop PCs, power consumption is generally less of a concern (see sidebar "Booting beyond BIOS?").

- By virtue of a SuperFetch cache, ReadyBoost "parks" the hard-drive head and spins down the platters for a higher

ADDITIONAL MATERIALS AND A HANDS-ON TEST-DRIVE

For more information, visit the Brian's Brain blog (www.edn.com/briansbrain) and peruse the solid-state-storage entries. You'll find online and print resources that can further expand your education on the topic, as well as an in-depth discussion on the supply-versus-demand influence on the market's evolution. I also showcase com-

PELLING commentary from industry representatives, including suppliers, users, and observers. And, as time allows, I also hope to include my real-life benchmarking results of solid-state drives versus hard-disk-drive competitors. To date, Samsung, Sandisk, and Super Talent have all promised me solid-state-disk-evaluation units.

percentage of the time, and the average operating temperature of the hard drive is therefore lower than it might be if a cache wasn't present. These factors tend to improve a hard-disk drive's overall ruggedness and reliability.

- If you can easily and cost-effectively upgrade your system with more DRAM, do so. DRAM delivers superior SuperFetch results to the supplemental-flash-memory alternative.

WILL READYDRIVE THRIVE?

Add to ReadyBoost the ability for the information stored in the flash-memory cache to survive system shutdown and hibernation cycles. Simplistically speaking, this augmented result is the concept behind ReadyDrive, another first-time Windows Vista capability. As Microsoft's documentation points out, ReadyDrive also offers hardware and software suppliers, as well as the SuperFetch algorithm itself, the ability to "pin" a given program or set of files such that the flash-memory array never flushes the pinned information. And, speaking of flushing, an operating system must support an expanded ATA (advanced-technology-attachment) command set to fully exploit a ReadyDrive-supportive storage subsystem. In Microsoft's words, "ReadyDrive uses a new NV [nonvolatile] cache command set standardized in the ATA8-ACS

command specification. This new command set allows ReadyDrive to do two important things: control what data is pinned for retention in the NV cache at any given time and control whether the drive is put into the NV cache power mode."

The ReadyDrive-supportive storage subsystem can take one of two primary forms. The members of the Hybrid Storage Alliance support the inclusion of flash memory within the hard drive, supplementing the RAM buffers in today's drives. Conversely, Intel with Turbo Memory and AMD with HyperFlash advocate a flash-memory array outside the hard-disk drive. This external implementation works with a software driver that intercepts commands that would normally go to the hard drive and appropriately directs and otherwise handles them, cognizant of the nonvolatile-semiconductor-memory alternative. Both sides of this contentious platform tug of war make compelling arguments for their preferred techniques.

Hybrid-drive advocates suggest that the company that implements the hard drive can best manage the flash-memory cache, so that the two technologies can operate as synergistically as possible. They also point to the all-in-one simplicity of the hybrid-drive approach, which, among other things, negates the necessity of installing and maintaining

BOOTING BEYOND BIOS?

Solid-state storage consumes less power than rotating magnetic storage, making the solid-state approach a compelling fit in battery-inclusive systems. Solid-state drives' advantages are less obvious in ac-powered gear, because power consumption is less of a concern in such equipment—except, perhaps, in servers—and such systems also often embed higher performance hard-disk drives, in many cases including multidrive RAID (redundant arrays of inexpensive disks). Desktop, workstation, and server PCs also tend to have significantly higher storage-capacity requirements than their mobile brethren, thereby notably delaying the point

at which solid-state disks will become cost-competitive on a dollar-per-gigabyte basis.

However, the emergence of operating-system virtualization in the computing platform creates another compelling opportunity for high-density flash memory. In this case, it would act as a "super BIOS," storing an in-system-upgradable "hypervisor" that would subsequently launch and act as an arbiter between multiple virtual-machine images on a hard-disk drive. Rumors are circulating that Dell is working on just such a configuration, and it will be interesting to see how widespread the implementation becomes in the future.



Go to www.edn.com/070927df to find this article's associated references and vendor box. At this URL, you can also post a comment on this article.

a custom-software driver. It also precludes the potential for remnants of valuable stored information to remain accessible within the system after removal of the hard drive. Intel Turbo Memory advocate Rick Coulson admitted at May's WinHEC (Windows Hardware Engineering Conference) that such a scenario was possible but suggested that you could deal with this issue by automatically reformatting or otherwise "wiping" the hard drive, which also scrambles the Turbo Memory contents, before removing it. Intel plans to mitigate this problem in the future by encrypting the information stored in the Turbo Memory array.

INSTEAD OF SUPPLEMENTING THE HARD DRIVE WITH FLASH MEMORY, YOU MIGHT INSTEAD CHOOSE TO REPLACE IT WITH A SOLID-STATE DRIVE.

Supporters of implementing ReadyDrive in the external flash-memory array focus on the approach's flexibility in system design and supply sourcing, a compelling pitch in the current era of uncertainty about flash-memory cost and availability. A computer manufacturer can keep buying the same few hard-drive variants it always has without further complicating its supply chain with hybrid-drive options in multiple capacity, performance, and other variations. And the manufacturer can, on a system-by-system basis, choose whether to populate a motherboard-resident connector with a flash-memory module for ReadyDrive use.

Turbo Memory and HyperFlash advocates also suggest that the system-side software for managing the flash memory and the hard drive and acting as arbiter between them is conceptually no

different from today's already-mature software-RAID (redundant-array-of-inexpensive-disks) utilities. And, as a Microsoft engineer pointed out to me at WinHEC, by virtue of the fact that they are both system-based, a Turbo Memory or a HyperFlash array could in the future cache not only the hard drive, but also other system-installed and network-connected storage peripherals, such as optical drives and NAS (network-attached storage, **Reference 5**). Regardless of the ReadyDrive array's location, it must, like a ReadyBoost subsystem, meet certain performance and other criteria to obtain Windows Vista Logo certification (**Table 1**).

"Microsoft recommends an NV cache size of at least 128 Mbytes to realize performance benefits for Windows Vista. To take advantage of the NV cache for the purposes of launching OEM-specified programs and to realize further performance benefits, an NV cache as large as 256 Mbytes to 1 Gbyte would be more effective. The larger the NV cache size, the more ReadyDrive can take advantage of the NV cache to deliver user benefits. However, NV cache sizes as small as 50 Mbytes are supported by ReadyDrive in Windows Vista. It is anticipated that future price reductions in the cost of flash memory and the increasing benefits from better performing flash technology will soon make NV cache sizes of 1 Gbyte or more practical and desirable. ReadyDrive can support NV cache sizes up to 2 Tbytes."

MAGNETIC: ANACHRONISTIC?

Instead of supplementing the hard drive with flash memory, you might instead choose to replace it with a solid-state drive. Suppliers such as Apacer, Mtron, PNY, PQI, Samsung, Sandisk (by virtue of its acquisition of M-Systems), STEC, Super Talent, and Transcend are vigorously promoting 1.8- to 3.5-in. solid-state drives as hard-disk-drive replacements for both general-purpose and embedded-computing applications. Vendors such as Bitmicro Networks and Texas Memory Systems focus on nearer-term solid-state-drive niches. And hard-disk-drive suppliers have already largely ceded the market for 1-in. and smaller form factors to solid-state-drive manufacturers.

You might ask: Aren't solid-state drives expensive? The answer depends on various factors. The solid-state drives' prices—at least to a first-order approximation—linearly scale with capacity, whereas hard-disk drives have a capacity-independent price “floor.” The cost of the platters, chassis, pickup head, and other mechanical elements determines this floor price. Commensurate with that fact, the smallest hard-disk capacity, which each platter's per-side areal density determines, is on a seemingly constant upward climb, which the recent transition from longitudinal- to perpendicular-recording technology has further extended.

Therefore, if your storage needs are moderate and increasing at a slower rate than hard-drive capacities are growing, you'll eventually reach a point at which the constructed cost of a sufficiently dense solid-state drive will be less than the floor cost of a single-platter hard-disk drive. In non-PC applications, this hard-disk-drive-to-solid-state-drive transition has in many cases already occurred; one notable example is the obsolescence of Apple's 1-in.-hard-drive-based iPod mini after only 18 months of production and its replacement with the flash-memory-based iPod nano.

The Apple iPod case is equally useful in highlighting other advantages of a solid-state drive over a hard-disk drive. Highly mobile computing platforms, such as Windows Tablet PCs and UMPCs (ultramobile personal computers), and portable multimedia players value the inherent ruggedness of a solid-state drive. For example, mobile-computer supplier Alienware offers the m9750, m9700, and m5550 laptops, which optionally employ two 32-Gbyte solid-state drives in a RAID 0 striped arrangement. Asus, Dell, Fujitsu, Palm, Samsung, Sony, and Toshiba are also producing numerous solid-state-drive-based designs, despite these systems' substantial near-term price increment

and capacity decrement over hard-drive alternatives. These manufacturers' decisions lead me to two encouraging conclusions: First, the vendors wholeheartedly buy into solid-state drives' inherent benefits. Second, they believe that industry participants will fairly quickly resolve solid-state-drive capacity and cost issues in target markets, therefore justifying today's market cultivation.

The lack of a motor-driven rotating platter may also provide these drives with lower average power and energy consumption, although hard-drive suppliers dispute this point. These objections are valid in some situations, such as when RAM- and flash-memory-caching techniques enable the hard drive to remain in spin-down for a large percentage of its operating interval and when other subsystems, such as a high-performance CPU and the LCD's backlight, also consume a notable portion of the battery's stored capacity.

Do you want other reasons to consider using a solid-state drive? Ponder, for example, Seagate's recent announcement that, by early 2008, it plans to no longer produce PATA (parallel-ATA)-interface hard drives, instead exclusively focusing its attention on SATA (serial-ATA) drives (**references 6 and 7**). Solid-state drives conversely come in a variety of interface and form-factor options, and you'll likely be able to continue purchasing them far into the future, given the suppliers' greater focus on non-PC applications. Consider, too, that moderate-sized flash-memory arrays take up less system-board surface area and overall system volume than a hard drive does. And don't forget about solid-state drives' potentially better performance than that of hard-disk drives. This superiority is unquestionably true with read accesses and also possibly true for writes, depending on the application's average data-set size and write frequency and how aggressively—at a power-consumption penalty—the solid-state-drive controller implements multiple parallel updates to multiple array components and multiple blocks with each component. According to Microsoft, you can now buy NAND-flash memory with faster sequential-data access and much faster random data access than mobile-PC hard drives offer. **EDN**

You can reach
Senior Technical Editor
Brian Dipert
at 1-916-760-0159,
bdipert@edn.com,
and www.bdipert.com.

