Cover Story WILLIAM WONG | Technology Editor

New technologies address the universal challenge of having enough storage space.

## High-Density

ne can never have enough storage, and chip designers have been working hard to deliver as much storage as possible. Almost every avenue is being pursued, from Nantero's nanotubebased NRAM to AMD's 3D high-bandwidth memory (HBM) chip architecture (*Fig. 1*). New technologies like NRAM, MRAM, and FRAM deliver non-volatility to challenge not only flash memory, but DRAM.

Conventional single-die chip storage continues to be replaced by 3D multi-die (*Fig. 2*). It pays to limit interconnect runs as speeds and capacity increase rapidly. Moving storage closer to the processor is important to performance, and it is occurring at all levels. For example, Diablo Technologies' Memory Channel Storage moves flash memory onto the DRAM memory channel.

## **3D ARCHITECTURES**

Using multiple dies in the same package allows a designer to easily mix different chip technologies. It is also a way to increase the packing density to improve capacity. Furthermore, it can boost performance by reducing interconnect length and allowing higher-speed transfers, since designers have better control over the connections.

Multiple dies on a chip have been commonly available for many years, but they were often used in areas where their higher

cost could be justified by the added performance and reliability. A chip carrier would typically be wired internally in a fashion similar to connecting a die to external pins.

GPU

Die stacking has also become common because of the tight spaces required by smartphones and other mobile devices. Microcontrollers incorporate flash memory and RAM on-chip, but these are often separate chips for more powerful microprocessors. Stacking a flash memory and DRAM die on top of a microprocessor die cuts down on space while providing a more reliable platform. Again, the wiring between chips is similar or identical to the chip-to-carrier connection.

The use of silicon interposer layers with through-silicon vias (TSVs) lets designers drop down multiple chips, but the approach is more akin to a silicon circuit board than the multichip carriers. The latter requires wiring to connect the chips together and to the outside world.

The use of interposers has been pushed by FPGA vendors such as Xilinx and Altera. Xilinx originally called the technology 2.5D (*see "10,000 Connections Between FPGA Slices" on electronicdesign.com*),; however, these days everyone refers to it as 3D. Xilinx initially used this approach to provide more logic elements using multiple FPGA die. The company also paired it with high-speed SERDES die that employed different silicon technology.



2. Solid-state storage has moved from memory chips (left) to multicarrier packages (middle) and now to 3D chips (right). The 3D chips essentially incorporate a multilayer circuit-board approach to provide an on-chip interconnect.

1. AMD's high-bandwidth memory (HBM) stacks storage on a silicon interposer along with CPU and GPU chips.

Altera's Stratix 10 uses a similar approach, but delivers 5.5

million logic elements and a quad-core Cortex-A53 complex

on a single die. Altera's SERDES die are connected using the

interposer layer that the company calls the Embedded Multi-

die Interconnect Bridge (EMIB). The main difference is that

the EMIB is a small set of interposers that link the FPGA die to

adjacent dies referred to as Tiles (Fig. 3). Altera plans on having

A number of vendors are doing 3D on the vertical side. The

Hybrid memory cube (HMC) addresses the RAM side of stor-

age (see "Hybrid Memory Cube Shows New Direction for High-

*Performance Storage*" *on electronicdesign.com*). It stacks die

using TSV connections. The base layer is the controller that has

multiple, high-speed SERDES linking it to one or more hosts.

The HMC specification addresses a range of configurations

from a simple single-port connection to one host to multiport

connections (Fig. 4). It targets enterprise servers, where the

There's an equally unquenchable appetite for flash storage,

too. Single-level cell (SLC) has given way to 2-bit multi-level cell

(MLC) and 3-bit triple-level cell (TLC). The tradeoffs between

SLC, MLC, and TLC tend to be space and capacity versus write

different I/O Tiles for application-specific interfaces.

V-NAND. This includes an M.2 and mSATA versions. The latter has a capacity up to 1 Tbyte using a 6-Gb/s SATA interface.

AMD's HBM combines the interposer and vertical stacking. CPU and GPU dies are single-layer types, but the memory is a stack of multiple RAM die. HBM's advantage over HMC is that it does not have to transfer data using high-speed SERDES. Instead, data can be moved in parallel.

HBM will first find a home in AMD's Radeon R9 Fury X GPU card. This will feature 4 GB of HBM plus the latest Fiji GPU from AMD with 4,096 stream processors, 64 GCN (graphics core next), 128 render output units, and 256 texture mapping units.

The memory interface is 4,096 bits wide. The GPU delivers 8.6 FLOPS while only having a TDP of 300 W. This is only 10 W more than the earlier Radeon R9 290X, yet performance improves by more than 50%.

## NONVOLATILE MEMORY TECHNOLOGIES

Packaging technologies are one way to improve storage density and performance. Changing the storage technology is another that is complementary to the packaging. Quite a few technologies have been around a while and found a niche, as have emerging technologies. All look to take away all or part of the market from SRAM/DRAM and flash memory. These technologies include various resistive RAM approaches, FRAM, and MRAM.

Nantero's NRAM is a new resistive memory technology that uses nanotubes (*Fig. 6*). It meshes well with current CMOS fab technology and offers impressive retention and performance specifications. It should retain data at room temperature for over 1,000 years. It can handle extended temperature ranges with lower retention capability. It should scale to 5 nm and can handle 3D structures and MLC support. It is based on a crossed nanotube (CNT) film deposited on a silicon substrate.

Crossbar's resistive RAM is similar (*see "RRAM Challenges Flash Storage" on electronicdesign.com*). It uses a switching medium that includes metal filaments with metal ions migrated from one end to another to change state. It has a 10-year retention and support for MLC and TLC techniques as well as 3D stacking. It is 20 times faster than flash, and features endurance support of over 10,000 cycles.

FRAM has been showing up in quite a few of Texas Instruments' (TI) MSP430 microcontrollers. It has replaced SRAM and flash storage on some models. TI also has a set of routines



stacked on top of each other. Samsung's **3. High-end FPGAs, such as Altera's latest Stratix 10, are putting the interface logic on separate** SSD 850 EVO line is the initial home for **TILEs. These die are connected using an interposer layer akin to a silicon circuit board.** 

demand for RAM is insatiable.

speed and write lifetime.

Samsung V-NAND, also known as ver-

tical NAND, is a 32-layer TLC device (Fig.

5). It provides the highest storage density

at this point (see "Mass Production Under-

way for 32-Layer 3D V-NAND Flash" on

*electronicdesign.com*). Identical dies are stacked on top of each other. Samsung's



that can save and restore system settings to FRAM that are not stored in the memory. This can include system settings like status and control registers for peripherals.

The Compute Through Power Loss (CTPL) utility API allows an almost instant-on capability, since the application stack and status are kept in the nonvolatile FRAM. A sleeping system automatically restores these settings upon power-up.

Everspin MRAM is similar to FRAM functionality, but uses a much different storage mechanism for its magnetic tunnel junction (see "Q&A: Everspin Takes MRAM to the

*Mainstream*" on electronicdesign.com). It has a data retention of more than 20 years with a 35-ns read/write time. It offers unlimited endurance, making it a good choice for a number of applications like RAID controller cache. The latest products employ Spin-Torque MRAM (ST-MRAM), which allow it to be used in DDR3 systems. Everspin also has parallel and SPI MRAM products. Chip capacity is 64 Mbits.

Phase change memory (PCM) originally utilized behavior of chalcogenide glass. Micron has phased out its commercial PCM line based on technology acquired from Numonyx. Research is ongoing, so the technology may resurface in the commercial space in the future.



6. Nantero's NRAM uses nanotubes that can be modified to different resistances.

5. Samsung's V-NAND chip uses 32 layers of flash-memory die stacked on top of each other.

ELECTRODE

FLE

## **NEW FLASH CONFIGURATIONS**

DRAM and flash memory in general remain the dominant force in storage. Flash memory has pushed rotating magnetic media into more niche markets as end-user devices have moved to mobile platforms requiring more compact and power-sipping flash storage. Large hard disks are still part of enterprise solutions, set-top box recorders, and PC backup, but flash memory has moved from its niche into the primary nonvolatile storage mechanism in most areas. Flash memory's ability to fit into almost any space makes it a much better choice than a hard disk for a growing number of applications, including wearable tech (*see "Wearing Your Technology" on electronicdesign.com*).

Diablo Technologies' Memory Channel Storage (MCS) moves flash memory to the memory channel of high-performance microprocessors. It licenses its technology to companies ranging from SanDisk to IBM that deliver DDR3 and

> DDR4 modules competing with DRAM for slots on the motherboard. The DIMM modules use the same interface, but the transactions are more like flash-memory storage devices rather than DRAM since the write speeds for flash are significantly slower.

MCS storage can be utilized as a block device, essentially providing a faster transfer rate to the system

> since the memory channel has a higher bandwidth than other I/O interfaces and there's less system overhead. MCS storage has a smaller transaction size than block devices that are designed to accom-

modate blocks of data on rotating magnetic media. The C2 version of MCS has a feature called NanoCommit that moves the minimum transfer size to a CPU cache line (*see "Memory Channel Flash Storage Provides Fast RAM Mirroring" on electronicdesign.com*). This allows mirroring of key data structures, opening the door to significant performance improvements for applications like database servers.

DRAM still tends to be faster and larger than other technologies to date. It is one reason that nonvolatile DRAM has emerged in enterprise environments where reliability is critical. Viking Technology's ArxCis-NV is a hybrid DIMM with DRAM and flash memory on-board (*see "Nonvolatile DIMMs* 

> and NVMe Spice Up the Flash-Memory Summit" on electronicdesign.com). The flash memory is used to retain the DRAM contents when the system is powered down. A supercap provides power in the event of system power loss, enabling the onboard controller to copy DRAM to flash memory. DDR3 and DDR4 versions are available.

> There is never a dull moment when it comes to storage technology. The mix of new storage technologies and packaging techniques provides designers with a wide range of options.