

3D XPoint: A New Revolution in Memory?

[Electronic Design](#)

[Jim Handy](#)

Thu, 2015-12-03 09:26



[Intel Corp.](#) and [Micron Technology](#) recently announced a new memory technology called 3D XPoint, and it's creating a lot of buzz. This buzz continued during last August's Intel Developer Forum with Intel's announcement of its Optane support products for 3D XPoint memories. So, once all the smoke clears away, what do we have?

3D XPoint, which is to roll out next year, is a new approach to alternative memories. Most alternative memory technologies have been relegated to niche markets as they patiently wait for today's leading technologies, DRAM and NAND flash (whose combined 2014 market amounted to \$75 billion), to reach a "scaling limit," where producers will lose the ability to perform any further cost reductions. Once that point is reached, these alternative technologies have the opportunity to succeed the incumbent and take over its market. Trouble is, researchers continue to find ways to scale DRAM lithographically, and NAND flash has moved to 3D, and appears to be poised to continue scaling vertically for a long time to come.

Related

[3D XPoint Technology Claims to Find Sweet Spot Between NAND and DRAM](#)

[For Designers, A New Selection of Storage Choices](#)

[Non-Volatile 3D Resistive RAM Packs 162 Gbits](#)

Intel and Micron were working on one such technology called Phase Change Memory (PCM), and it had the opportunity to take a share of the NOR flash market. However, the NOR flash market, which has shrunk to only \$2 billion in annual revenues, is not really worth pursuing.

The two companies appear to have come upon a breakthrough that should allow their new memory type, which one would assume is based on their significant research into PCM, to compete against DRAM on cost. As a result, they can create a new market for an additional memory layer between DRAM and NAND flash. This is today's market thrust for 3D XPoint memory.

How important is this technology? Well, after the original late July introduction, Intel included 3D XPoint in three of its keynotes and seven of its sessions during its August Intel Developer Forum, and Micron spoke of it prominently at the company's August 14 Summer Analyst Day. At SNIA's SDC (Storage Developer Conference) in September, a total of 15 papers showed how to use nonvolatile memory technology from both hardware and software standpoints.

This analyst was asked to present at the last moment during the Flash Memory Summit in early August. The session drew so many attendees that the show's management had to open the wall to the adjoining room to

ommodate the overflow crowd.

A Perspective on 3D XPoint's Future

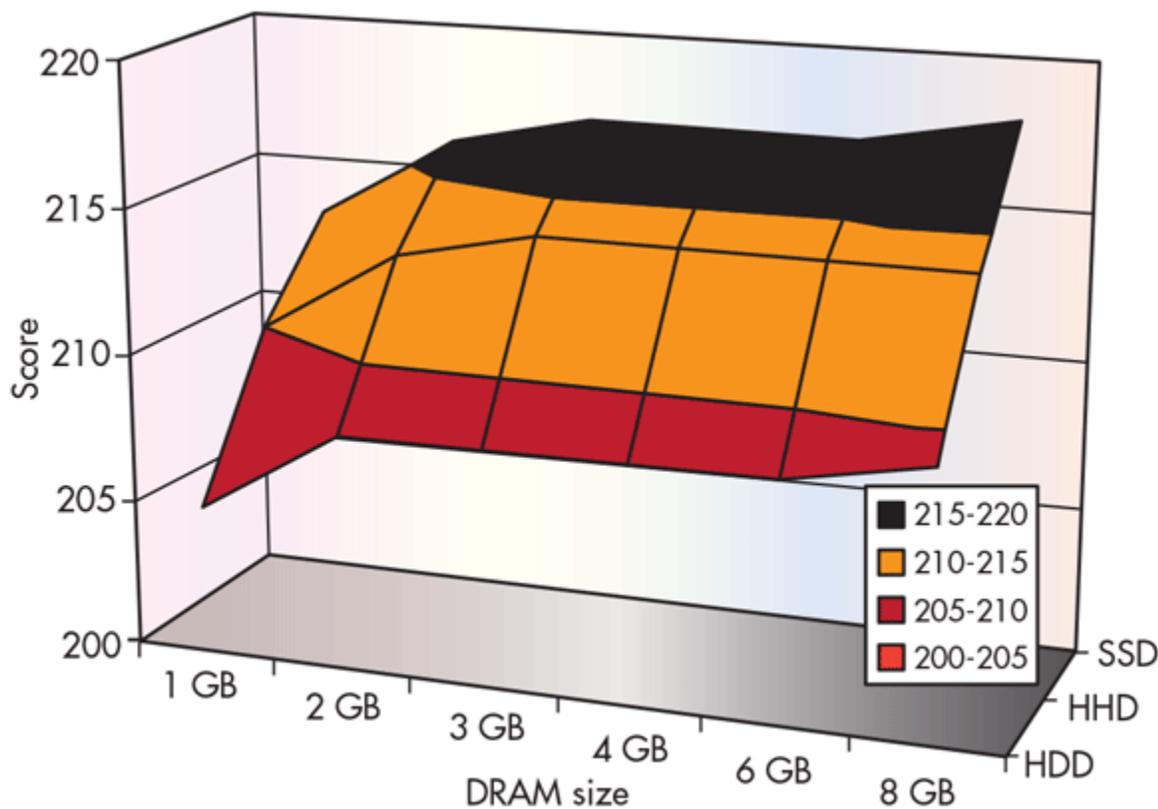
My company, [Objective Analysis](#), recently released a report covering 3D XPoint technology ([A Close Look At The Micron/Intel 3D XPoint Memory](#)). Despite the fact that Micron and Intel have not disclosed the memory technology used by this product, the companies have given ample information on the ways that they will market the product and its potential impact on computing.

We started with this information and tempered it with our extensive backgrounds in the memory business to provide a future view of the market for 3D XPoint memory; the issues the two partners will face trying to create that market; production challenges; and the impact that 3D XPoint will likely have on today's established markets, including DRAM, NAND flash, and hard-disk drives (HDDs).

We also investigated the market for persistent memory and explain what will be required to make this new technology gain traction. In a nutshell, there's no application software in place today that supports nonvolatile memory in PCs and servers. It will take some time for this software to be written and tested, but the need for standards has already been addressed very well by the Storage Networking Industry Association (SNIA) and its [NVM Programming Model](#).

Interpreting Benchmarks

Without this software support, is there even a market for 3D XPoint? There certainly is! Some time back, Objective Analysis performed nearly 300 benchmarks to determine whether NAND flash could be used as an alternative to DRAM when upgrading a PC ([How PC NAND Will Undermine DRAM](#)). We had already learned that IT managers were using solid-state disks (SSDs) to reduce their servers' DRAM requirements, and we expected the same to be true of PCs.

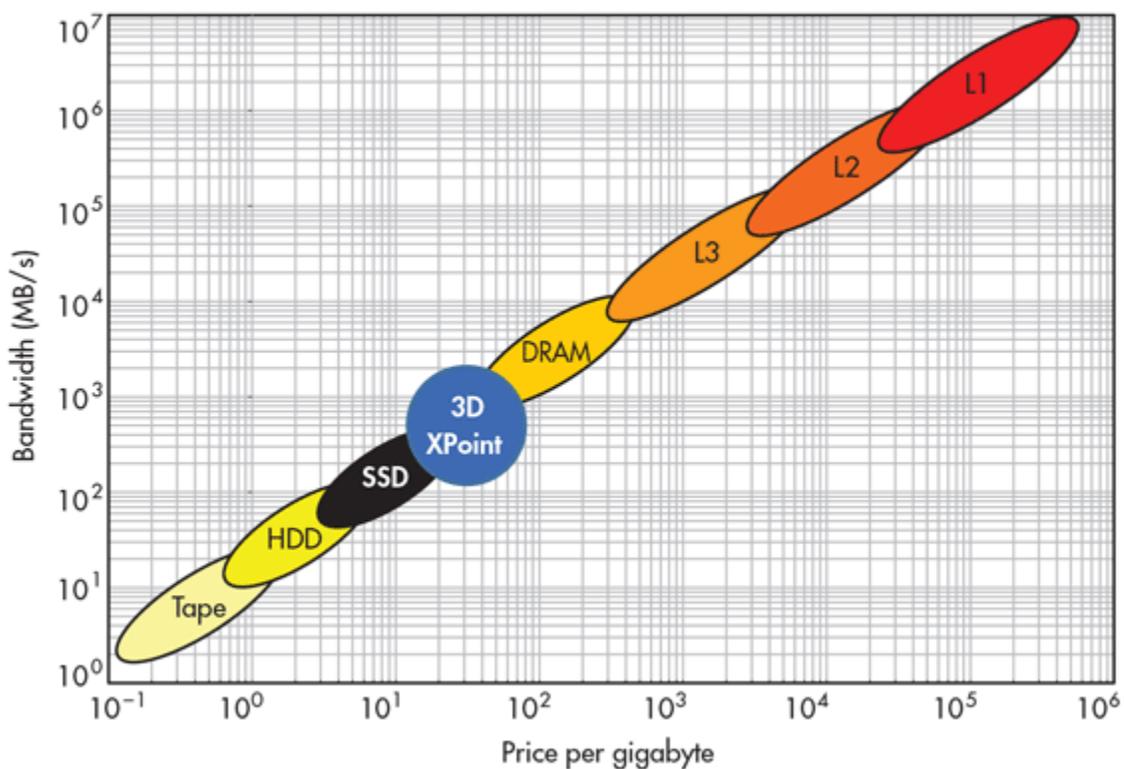


at we learned from these benchmarks was that a PC definitely needs the first 1 to 2 GB of DRAM. After that, though, an upgrade dollar is better spent by adding NAND flash (typically as a small SSD) than by adding the same dollar value of DRAM. DRAM costs 15 times as much as NAND flash, so you can add 15 GB of SSD for every 1 GB of DRAM you choose *not* to add to the system. Even though DRAM is about 1000 times faster than NAND flash, the larger size of the NAND flash matters more than the speed difference, allowing it to dramatically reduce HDD accesses to improve the speed.

To help visualize this point, the chart in Figure 1 shows the SysMark 3D benchmark's performance (vertical axis) with a number of combinations of DRAM size (horizontal axis) and NAND flash (depth axis). Increasing the DRAM size from 1 GB to 2 GB greatly increases performance, but after that, there's little benefit in adding DRAM until we get to the 6- to 8-GB transition. Throughout the chart, though, important performance enhancements result from the addition of flash.

The same rules will apply to 3D XPoint memory. Whether it's added to the memory bus or as an SSD (as was done in our benchmark study), it will allow IT managers to reduce their DRAM usage. As long as 3D XPoint memory costs less than DRAM, then it should provide more "bang for the buck" than DRAM, even though it's a slower-than-DRAM technology.

The position of 3D XPoint memory is illustrated by the chart provided in Figure 2. This conceptual diagram shows price per gigabyte across the horizontal axis and bandwidth up the vertical axis. The orbs represent different memory types, from the caches inside the processor (L1 and L2) right down to magnetic tape in the lower left corner. Since 3D XPoint is cheaper than DRAM, but more costly than NAND, and since it's faster than NAND but slower than DRAM, it should fit well into the memory/storage hierarchy.



Cost Concerns

One very tricky part of this approach must be managed very carefully by Micron and Intel. Today, 3D XPoint is experimental. The companies have produced enough working chips to demonstrate one SSD at IDF. This is not

roduction technology, and that means that it's currently far more costly than DRAM, even though the companies assert that 3D XPoint is 10 times as dense as conventional memory.

The technology will not be adopted until the companies bring the price below that of DRAM, which will require the production of millions of chips to carry the technology down the learning curve. These millions of chips will not find homes unless they are sold for sub-DRAM prices. In brief, Micron and Intel will need to sell millions of chips at a loss before it becomes rational to use 3D XPoint in a system.

Consequently, Intel and Micron plan to sell this product into a market for which there's no specific software support, and in which the companies' cost structure doesn't support XPoint's use as a new memory layer. How, then, will this market develop? Intel and Micron will need to sell 3D XPoint below its cost until high-volume production is reached. After that, the economies of scale will cause this technology's cost to fall well below that of DRAM. Unfortunately for Intel and Micron, this will probably occur well after the late 2017 DRAM price collapse that's been predicted by Objective Analysis for years.

Another very important piece of the puzzle needs to be put in place to allow this market to develop: System hardware and software will need to accept this new memory layer. Fortunately, Intel has a commanding portion of both the PC and the server processor markets, allowing the company to decide when and how to add industry-wide support for this new memory layer. This wouldn't be possible with any other memory manufacturer.

What's the Motivation?

Why does Intel want to do this? Intel sells processors, and a large part of its sales consists of upgrades. Faster or better processors must provide a compelling improvement to the end user. The company determined in the late 1980s that it needs to manage the entire platform to accomplish such a goal. If the processor speeds up, but the system can't support that additional speed, then the system won't perform any better with a fast processor than it did with a slow one.

To this end, Intel has managed the introduction of a number of increasing DRAM speeds (from synchronous DRAM through DDR, DDR2, DDR3, and now DDR4) and a number of other initiatives whose purpose is to improve the user experience. By adding 3D XPoint to a system, the apparent size of the DRAM may increase by as much as 10 times, freeing the processor to perform at its peak.

What is Micron's motivation? By teaming up with Intel to produce a memory that will undermine DRAM revenues, Micron can steal DRAM market share away from its competition (Samsung and SK hynix) while shipping a product that's likely to support higher margins over the long term, by nature of its sourcing.

Will it succeed? Objective Analysis believes that it will. We see the two partners working together to sell this product at below-DRAM prices before their costs support this approach in order to create a market. Once achieved, 3D XPoint memory is in a position to penetrate any system where performance and economy are both important—in other words, every computing system in the world.

Source URL: <http://electronicdesign.com/memory/3d-xpoint-new-revolution-memory>