

SEMICONDUCTOR ENGINEERING

Home > Manufacturing, Design & Test > What's Next For NAND?

MANUFACTURING, DESIGN & TEST

What's Next For NAND?



Vendors pin hopes on 3D memory, raising specter of 2D NAND shortages; other memory types add to uncertainty.

MAY 19TH, 2016 - BY: MARK LAPEDUS

NAND flash memory is a key enabler in today's systems, but it's a difficult business. NAND suppliers require deep pockets and strong technology to survive in the competitive landscape. And going forward, vendors face new challenges on several fronts.

On one front, for example, the overall NAND market is currently in the doldrums, amid soft product prices and a mild capacity glut. Demand is expected to rebound in the second half of 2016, although there is still uncertainty in the market.

Then, on the technology front, today's planar NAND is reaching its physical scaling limit. And so NAND suppliers are pinning their future hopes on the successor to planar NAND—3D NAND.

3D NAND is shipping, but the technology is taking longer than expected to enter the mainstream. It is more difficult to make than previously thought. 3D NAND resembles a vertical skyscraper, in which horizontal levels or layers are stacked and then connected using tiny vertical channels.

"The early movers, such as Samsung and Micron, are ramping up 3D NAND quickly, while SK Hynix and SanDisk/Toshiba are lagging," said Greg Wong, an analyst with Forward Insights. "The technology and yield learning is taking longer."

And it isn't getting any easier. Today's leading-edge 3D NAND chips are 32- and 48-layer devices, but the technology will likely hit the ceiling at 128 layers in the 2018 timeframe or so. So to extend 3D NAND beyond 128 layers, vendors are quietly developing a technology called string stacking. Still in R&D, string stacking involves a process of stacking individual 3D NAND devices on top of each other.

For example, a vendor might stack three separate 48-layer 3D NAND devices, creating a 144-layer chip. Even with string stacking, though, 3D NAND could hit the wall at 300 layers. It will take vast resources and capital to extend current and futuristic 3D NAND devices to 128 layers and beyond. "Scaling the number of layers is not just a technical challenge, but an economic one as well," Wong said.

TRENDING ARTICLES

Why EUV Is So Difficult

One of the most complex technologies ever developed is getting closer to rollout. Here's why it took so long, and why it still isn't a sure thing.

BY MARK LAPEDUS

Silicon Photonics Comes Into Focus

Using light to move large quantities of data looks promising, but gaps remain and the adoption timeline will vary by application.

BY ED SPERLING

Will 5G Enable Connected Cars?

New standards are being considered for automobiles, but there are still a myriad of issues to be resolved.

BY ANN STEFFORA MUTSCHLER

Can We Measure Next-Gen FinFETs?

Why the semiconductor industry needs breakthroughs, and why it's getting tougher to provide them.

BY MARK LAPEDUS

Moore's Law Debate Continues

As the industry speculates about 5nm and below, questions surrounding node shrinks remain.

BY ANN STEFFORA MUTSCHLER

Knowledge Centers

Entities, people and technologies explored

LEARN MORE



SPONSORS



NEWSLETTER SIGNUP

Email:

- Interests:
- ☐ System-Level Design
- ☐ Low Power-High Performance
- ☐ Manufacturing, Design & Test
- ☐ IoT, Security & Automotive

Subscribe

POPULAR TAGS



RECENT COMMENTS

- MarkC on Manufacturing Bits: Nov. 23

YMTC
Exhibit 2025
Micron v. YMTC
IPR2025-00191
Ex. 2025-001


NAND could hit the wall at 300 layers. It will take vast resources and capital to extend current and futuristic 3D NAND devices to 128 layers and beyond. "Scaling the number of layers is not just a technical challenge, but an economic one as well," Wong said.

In any case, OEMs will need to keep close tabs on the technology. To help OEMs, Semiconductor Engineering has taken a look at the status of the following technologies—planar NAND; 3D NAND; and futuristic 3D NAND with string stacking.

Planar NAND

Invented in the 1980s, NAND flash is a **nonvolatile memory** technology that can be electrically erased and reprogrammed. Basically, NAND is used for data storage applications. Its close cousin, NOR flash, is geared for code storage. In today's systems, NAND plays a key role in the traditional memory hierarchy. Basically, **SRAM** is integrated into the processor for cache. **DRAM** is used for main memory. And disk drives and NAND-based solid-state drives (SSDs) and memory cards are used for storage.

The NAND market is dominated by large suppliers. (See chart below.)



Source: Semiconductor Engineering chart compiled from TrendForce data.

In 2015, worldwide NAND revenues increased by 10%, while bit growth grew by 52%, according to Web-Foot Research. "In 2016, revenue will be flat at a 0.4% growth rate even though bit growth will come in around 46%," said Alan Niebel, president of Web-Foot Research. "This is due to oversupply of planar NAND and price declines for mobile, consumer and some SSDs."

Today, prices for planar NAND are inexpensive and continue to fall. Not long ago, for example, hard disk drives (HDDs) were significantly cheaper than rival NAND-based SSDs. By year's end, the price difference between a 128-GB SSD and a 500-GB HDD will fall to less than \$3, according to TrendForce.

For OEMs, though, there is a potential problem. NAND flash vendors have converted some of their planar NAND fab capacity over to 3D NAND. If demand picks up, and vendors can't ramp up 3D NAND fast enough, OEMs

CAPTURED

LEARN MORE

RELATED ARTICLES

What Transistors Will Look Like At 5nm

As finFETs run out of steam after 7nm, what comes next? The debate is just beginning.

BY MARK LAPEDUS

Focus Shifts To Architectures

As chipmakers search for the biggest bang for new markets, the emphasis isn't just on process nodes anymore. But these kinds of changes also add risks.

BY ED SPERLING

Deeper Inside Intel

Top process execs shed light on 10nm and beyond, plans for longer time between nodes, and the future of EUV, finFETs and stacked die.

BY MARK LAPEDUS

Neural Net Computing Explodes

Deep-pocket companies begin customizing this approach for specific applications—and spend huge amounts of money to acquire startups.

BY GALE MORRISON

Packaging Wars Begin

OSATs and foundries begin to ramp offerings and investments in preparation for mainstream multi-chip architectures.

BY MARK LAPEDUS

Why EUV Is So Difficult

One of the most complex technologies ever developed is getting closer to rollout. Here's why it took so long, and why it still isn't a sure thing.

BY MARK LAPEDUS

The Zen Of Processor Design

AMD's CTO talks about how to achieve more performance per watt and how chip architectures are changing.

BY ED SPERLING

RECENT COMMENTS

- MarkC on Manufacturing Bits: Nov. 23
- witeken on More EUV Mask Gaps
- Mark LaPedus on More EUV Mask Gaps
- Tim on Can We Measure Next-Gen FinFETs?
- memister on Why EUV Is So Difficult
- memister on Why EUV Is So Difficult
- memister on Why EUV Is So Difficult
- memister on Why EUV Is So Difficult
- Chris on Why EUV Is So Difficult
- Ed Sperling on Making 2.5D, Fan-Outs Cheaper

For OEMs, though, there is a potential problem. NAND flash vendors have converted some of their planar NAND fab capacity over to 3D NAND. If demand picks up, and vendors can't ramp up 3D NAND fast enough, OEMs could face product shortages in the second half of 2016 and perhaps beyond.

"Our outlook is for a shortage in NAND in general," said Jim Handy, an analyst with Objective Analysis. "It all depends on whether or not people are able to overcome the technical barriers to get out 3D NAND. And it's really hard to predict when that will happen."

The other problem? Planar NAND is running out of steam.

Planar NAND is running out of steam, but it's far from dead.

Planar NAND is based on a floating gate transistor structure. A NAND device, which resembles a MOSFET, consists of a source and a drain with a channel running between them. Unlike a MOSFET, there are two gates on the top of a NAND structure. A control gate is on top, while the floating gate is on the bottom. The two gates are insulated by an oxide layer.

The data is stored in a NAND cell. In single-level cell (SLC) NAND flash, there is 1 bit of data per cell. Today's mainstream NAND makes use of multi-level cell (MLC) and triple-level cell (TLC) technology, which stores 2 bits and 3 bits of data per cell, respectively.

NAND vendors have scaled the cell size by roughly 100 times over the past decade, according to Micron Technology. In fact, thanks to advanced [lithography](#), vendors have extended planar NAND down to the mid-1xnm node regime.

Today, NAND vendors are shipping 16nm and 15nm parts. But at those nodes, NAND is running out of gas and it will no longer scale. It is becoming difficult to scale the memory cell and floating gate in planar NAND.

But planar NAND is not dead. "Planar NAND is not going away," said Kevin Kilbuck, director of NAND strategic planning at Micron. "Customers want 2D NAND for the foreseeable future for let's call it the non-high capacity storage market. Not all of the applications are moving to 3D NAND right away. And not all of the manufacturers will convert all of their capacity from 2D to 3D NAND overnight. It's not economical to manufacture low-density products in 3D NAND fabs. It's also very expensive to build new 3D NAND fabs or convert 2D NAND fabs to 3D."

3D NAND

Still, the future rests with 3D NAND. This technology first appeared in 2007, when Toshiba introduced the world's first 3D NAND technology. Later, Samsung, SK Hynix and the Micron/Intel duo introduced 3D NAND.

3D NAND represents a radical departure from planar or 2D NAND. Planar NAND involves the production of horizontal strips of polysilicon. The strips are used to make the wordlines. These, in turn, connect the control gates of the memory cells.

achieve more performance per watt and how chip architectures are changing.

BY ED SPERLING

Sorting Out Next-Gen Memory

A long list of new memory types is hitting the market, but which ones will be successful isn't clear yet.

BY MARK LAPEDUS

3D NAND represents a radical departure from planar or 2D NAND. Planar NAND involves the production of horizontal strips of polysilicon. The strips are used to make the wordlines. These, in turn, connect the control gates of the memory cells.

In 3D NAND, the polysilicon strips are stretched, folded over and stood up vertically. Basically, 3D NAND involves a stack of layers. The layers are connected with tiny vertical channels. The layers, which are horizontal, are the active wordlines. "The bitlines also run horizontally in the metal layers on the top of the chip," Objective Analysis' Handy said. "The vertical channels are the NAND 'strings' that attach to the bitlines."

There are other differences as well. Micron and its technology partner, Intel, extended the floating gate architecture to 3D NAND. In contrast, Samsung, SK Hynix and the SanDisk/Toshiba duo are not using a floating gate for 3D NAND. Instead, these vendors went to a technology called charge trap flash.

All told, 3D NAND has some advantages over planar NAND. "Compared to planar NAND, 3D NAND offers a significant bit density increase," said Yang Pan, chief technology officer for the Global Products Group at [Lam Research](#). "As such, there is increasing adoption of 3D NAND by SSDs as a storage solution for servers/datacenters as well as high-end consumer applications."

Still, 3D NAND has some challenges. "From the device side, it's the channel mobility," said Er-Xuan Ping, managing director of memory and materials within the Silicon Systems Group at [Applied Materials](#). "Polysilicon does not have good mobility."

In 3D NAND, the goal is to move the current through a polysilicon-based vertical channel. A 3D NAND device with fewer layers, and a shorter channel length, might have acceptable mobility.

Problems can arise as vendors scale their devices with more layers, meaning the channel length becomes taller. "When you go up, the polysilicon channel will be limited by its mobility," Ping said.

Mobility is affected by the number of layers.

In addition, 3D NAND involves a number of new and complex process steps in the fab. Planar NAND is dependent on advanced lithography. In contrast, 3D NAND makes use of trailing-edge geometries from 40nm to 20nm. 3D NAND requires lithography, but the challenges shift from advanced patterning to deposition and etch.

Despite the challenge, 3D NAND is set to take off. In 2015, the overall 3D NAND market reached \$4.5 billion in terms of sales, according to Web-Foot Research. "In 2016, 3D NAND will come in big time with both Toshiba and Micron producing volume shipments, thereby growing worldwide bit shipments by 350% and revenues by 230%," Web-Foot Research's Niebel said.

Each vendor, meanwhile, is ramping up 3D NAND at various stages. In 2013, Samsung shipped the world's first 3D NAND device, a 24-layer, 128-gigabit (Gb) chip. Then, last year, Samsung shipped its third-generation 3D

Each vendor, meanwhile, is ramping up 3D NAND at various stages. In 2013, Samsung shipped the world's first 3D NAND device, a 24-layer, 128-gigabit (Gb) chip. Then, last year, Samsung shipped its third-generation 3D NAND device. The device is 48 layers, based on tri-level cell (TLC) technology, resulting in a 256-Gb chip.

Samsung's previous chip was a 32-layer chip. The 48-layer device is roughly 2X to 2.2X faster in terms of a sequential read and a sequential write, according to Jim Elliott, corporate vice president of Samsung Semiconductor.

The 48-layer chip also enables Samsung to penetrate new markets, namely the enterprise SSD sector. "The big story is TLC in the enterprise space," Elliott said at the time of the product announcement, which was last year. "That's where that disruption is taking place."

Samsung's rivals, however, opted to skip the 24-layer regime. For most, a 16nm or 15nm planar NAND chip is still cheaper on a cost-per-bit basis than a 24-layer 3D NAND device.

Meanwhile, the Intel/Micron duo recently entered the 3D NAND market by rolling out a 32-layer device. Meanwhile, SK Hynix and the SanDisk/Toshiba duo are separately sampling 48-layer products.

The 32-layer device from Intel/Micron supports both multi-level cell (MLC) and TLC technology, enabling densities of 256- and 384-Gb, respectively. The 384-Gb chip, according to Micron, is the highest density 3D NAND device in the market.

To accomplish this feat, Micron integrated the logic circuitry under the layer stack. It refers to this as **CMOS** under array. "We can get the majority of logic under the array, thereby saving a lot of space. That allows more density in a given silicon area," Micron's Kilbuck said. "It also allows us to segment the array. It can increase the number of planes, for example. Our NAND pages are divided into four planes. What that does is increase your throughput and performance of an individual die."

Despite the obvious benefits of 3D NAND, there is one big question: When will 3D NAND reach price parity with 2D NAND?

Today, 3D NAND is sold at a premium. The 32- and 48-layer devices are inching closer in terms of price parity with 2D NAND. But it may take chips with 64 layers and beyond to reach the magical price parity point.

"Our TLC-based 3D NAND is getting close to what 2D MLC can do in endurance," Kilbuck said. "My gut feeling is sometime next year we would see price parity between 2D and 3D. I'm talking about price per gigabyte. But everybody needs to be in high volume (production) for that to occur."

What's next?

So what happens next? "This is the year of 48 layers going to 64," Applied's Ping said. "Whether 3D NAND can go to 96 or 128 is limited by the etching capability."

In fact, it's difficult to fabricate devices with 64 or more layers. Right now, the high aspect ratio etch tools are not ready or struggling to fabricate

Ping said. "Whether 3D NAND can go to 96 or 128 is limited by the etching capability."

In fact, it's difficult to fabricate devices with 64 or more layers. Right now, the high-aspect ratio etch tools are not ready or struggling to fabricate devices at 64 layers and beyond. For now, the aspect ratios are too complex and difficult.

And at 64 layers and beyond, channel mobility becomes an issue, as previously stated. "It will limit the device performance or device height," Ping said.

So going forward, NAND suppliers will simultaneously follow two parallel paths. The first path is to wait for the etch tools and other manufacturing techniques to arrive. And if they arrive on time, vendors could scale today's 3D NAND device from 32- and 48-layers, to 64 layers, to 96 and then to 128.

The second path is to move towards string stacking technology. This involves stacking two or more individual devices on top each other. Each device is separated by an insulating layer.

String stacking is already in the works. Recently, Micron presented a paper on a new 64-layer chip. Micron, according to multiple sources, stacked two 32-layer chips on top of each other.

In theory, string stacking could involve several different combinations. For example, a vendor could stack three 32-layer chips, enabling a 96-layer device. In addition, a vendor could stack three 96-layer chips, resulting in a 288-layer product.

The trick is to connect the various chips together. Today, vendors are looking at several different interconnect schemes. "Some of them will put a source line in the middle," Ping said. "That is just one of the many options."

Clearly, string stacking is challenging. And even with this technology, 3D NAND could hit the wall at 300 layers or so. "That's the limit," he said. "That's based on yield and stress."

All told, 3D NAND will remain viable at least to 2020, and perhaps beyond. But suddenly, 3D NAND has some new competition, which complicates the landscape.

Today, the Intel/Micron duo are sampling 3D XPoint, a [ReRAM](#)-like device that could potentially compete with 3D NAND in enterprise SSD applications. ReRAM, a nonvolatile memory technology, is attractive because it delivers fast write times with more endurance than today's flash.

There are other promising technologies in R&D, particularly vertical ReRAM. "The densities (in vertical ReRAM) won't be as high as 3D NAND," Ping said. "But the speeds can overcome some of these density limitations."

Time will tell if ReRAM will displace 3D NAND. In fact, ReRAM and the other next-generation memory types still have a lot to prove. At one time, the

Today, the industry is still sampling 3D XPoint, a new type of device that could potentially compete with 3D NAND in enterprise SSD applications. ReRAM, a nonvolatile memory technology, is attractive because it delivers fast write times with more endurance than today's flash.

There are other promising technologies in R&D, particularly vertical ReRAM. "The densities (in vertical ReRAM) won't be as high as 3D NAND," Ping said. "But the speeds can overcome some of these density limitations."

Time will tell if ReRAM will displace 3D NAND. In fact, ReRAM and the other next-generation memory types still have a lot to prove. At one time, the newfangled memory types were supposed to replace conventional memory, but they have fallen short of their promises. And conventional memory, such as DRAM and NAND, continue to chug along.

Related stories

- 3D NAND Market Heats Up
- Memory Hierarchy Shakeup
- What's Next For DRAM?
- 1xnm DRAM Challenges



TAGS: 3D NAND 3D POINT APPLIED MATERIALS CHARGER TRAP FLASH FORWARD INSIGHTS

LAM RESEARCH MICRON MOSFET NAND OBJECTIVE-ANALYSIS RERAM SAMSUNG SANDISK

SK HYNIX SOLID STATE DISK SSD STRING STACKING TOSHIBA TRENDFORCE

WEB-FEET RESEARCH



Mark LaPedus (all posts)

Mark LaPedus is Executive Editor for manufacturing at Semiconductor Engineering.

Deploying Multi-Beam Mask Writers

Mark LaPedus

< Previous Article

Next Article >

Mask Maker Worries Grow

Mark LaPedus

SEMICONDUCTORENGINEERING

ABOUT

ABOUT us

CONTACT us

Advertising on SemiEng

Newsletter SignUp

NAVIGATION

Homepage

Special Reports

System-Level Design

Low Power-High Perf

Manufacturing Des & Test

IoT, Security & Automotive

CONNECT WITH US

Facebook

Twitter @semiEngineering

LinkedIn

YouTube

Document title: Semiconductor Engineering ... What's Next For NAND?

Capture URL: <https://web.archive.org/web/20161127122539/https://semiengineering.com/whats-next-for-nand/>

Page loaded at (UTC): Fri, 15 Aug 2025 09:37:31 GMT

Capture timestamp (UTC): Fri, 15 Aug 2025 09:38:39 GMT

Capture tool: 10.59.0

Collection server IP: 54.145.42.72

Browser engine: Mozilla/5.0 (X11; Linux x86_64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/138.0.7204.97 Safari/537.36

Operating system: Linux (Node 22.17.0)

PDF length: 8

Capture ID: pL972Qg2z2269T5o9F6UQz

Display Name: Alan.Man