

HOME

SPECIAL
REPORTS

SYSTEM-LEVEL
DESIGN

LOW POWER -
HIGH PERFORMANCE

MANUFACTURING &
PROCESS TECHNOLOGY

PACKAGING, TEST &
MATERIALS

IOT, SECURITY
& AUTOMOTIVE

KNOWLEDGE
CENTER

STARTUP
CORNER

INDUSTRY
RESEARCH

SEMICONDUCTOR ENGINEERING

Select Language

Home > Manufacturing & Process Technology > NAND Market Hits Speed Bump

MANUFACTURING & PROCESS TECHNOLOGY

NAND Market Hits Speed Bump



Transition from planar to 3D NAND is harder and more time-consuming than expected.

JUNE 21ST, 2017 - BY: MARK LAPEDUS



Demand for NAND flash memory remains robust due to the onslaught of data in systems, but the overall NAND flash market is stuck in the middle of a challenging period beset by product shortages, supply chain issues and a difficult technology transition.

Intel, Micron, Samsung, SK Hynix and the Toshiba/Western Digital duo continue to ship traditional planar NAND in the market, but this technology is reaching its physical limit at the current 1xnm node regime. So for some time, the same vendors have been developing and ramping up a next-generation technology called 3D NAND, which is used for storage applications like smartphones and solid-state storage drives (SSDs).

3D NAND is expected to reach the mainstream by year's end. This is about two to three years longer than expected. As it turns out, 3D NAND is more difficult to fabricate than previously thought. Unlike planar NAND, which is a 2D structure, 3D NAND resembles a vertical skyscraper, in which horizontal layers are stacked and then connected using tiny vertical channels.





Fig. 1: 2D NAND architecture. Source: Western Digital.



TECHNICAL PAPERS

Autonomous Vehicle Navigation In Rural Environments Without Detailed Prior Maps (MIT)

MAY 15, 2018 BY TECHNICAL PAPER LINK

Silicon CMOS Architecture For A Spin-Based Quantum Computer

JANUARY 24, 2018 BY TECHNICAL PAPER LINK

Making High-Capacity Data Caches More Efficient

OCTOBER 24, 2017 BY TECHNICAL PAPER LINK

How Neural Networks Think (MIT)

OCTOBER 18, 2017 BY TECHNICAL PAPER LINK

New AI Algorithm Monitors Sleep With Radio Waves (MIT & Mass General)

AUGUST 22, 2017 BY TECHNICAL PAPER LINK

Knowledge Centers

Entities, people and technologies explored

LEARN MORE

RELATED ARTICLES

Packaging Chips For Cars

Why packaging is becoming increasingly critical and difficult in the automotive market.

BY JEFF DORSCH

FPGAs Becoming More SoC-Like

Lines blur as processors are added into traditional FPGAs, and programmability is added into ASICs.

BY ANN STEFFORA MITSCHER

Overall site

Search

SPONSORS



NEWSLETTER SIGNUP

Email:

Interests:

- ☐ System-Level Design
- ☐ Low Power-High Performance
- ☐ Manufacturing & Process Tech
- ☐ Packaging, Test & Materials
- ☐ IoT, Security & Automotive

Subscribe

POPULAR TAGS



RECENT COMMENTS

- Taylor Bishop on What's What In Advanced Packaging
- John Schneider on The Battle To Embed The FPGA
- Tom on CEO Outlook On Chip Industry
- Gary Huang on Chip Dis-Integration
- Trelliant Fang on Advanced Packaging

YMTC
Exhibit 2031
Micron v. YMTC
IPR2025-00191

Ex. 2031-001




Fig. 2: 3D NAND architecture. Source: Western Digital




Fig. 3: Samsung's V-NAND. Source: Samsung

Still, NAND [flash](#) vendors continue to convert a large percentage of their fab capacity from planar to 3D NAND. But the migration is taking longer and some vendors can't bring up 3D NAND fast enough. A few are struggling to ship 3D NAND.

During this period, demand has exceeded supply for both planar and 3D NAND, causing a shortage of parts that started in July of 2016 and has extended until now. In fact, suppliers planned for 40% bit growth in the overall NAND market in 2017, but the current demand picture calls for about 45% bit growth this year, according to Objective Analysis.

"The shortages have become worse," said Jim Handy, an analyst with Objective Analysis. "Everybody was planning for about 40% bit growth. They planned to reach that by building new capacity and by converting over to 3D NAND. The 3D NAND conversion has not gone as smoothly as anticipated. Because of that, production bit growth has not been able to keep pace with demand bit growth."

Prices have been relatively flat for NAND, but the supply issues are far from over. "We are expecting to see shortages until the middle of 2018," Handy said. "What we are looking at is how long it's going to take for the last problem to be solved in order to make 3D NAND a cost effective part to build. Until that happens, we expect the shortages to grow more acute, because demand will continue to grow while production does not."

In addition, there are other issues taking place in the NAND market today. Here are the main ones:

- Suppliers are shipping their previous-generation 3D NAND parts, which involve 32- and 48-layer devices. But the current transition to the next iterations - 64- or 72-layer 3D NAND devices - has been

Lines blur as processors are added into traditional FPGAs, and programmability is added into ASICs.

BY ANN STEFFORA MUTSCHLER

Big Trouble At 3nm

Costs of developing a complex chip could run as high as \$1.5B, while power/performance benefits are likely to decrease.

BY MARK LAPEDUS

Non-Traditional Chips Gaining Steam

Flexible hybrid electronics are showing up in a variety of markets where electronics never existed before.

BY JEFF DORSCH

Investors Back IoT Startups

Total private funding was more than \$1.35B in 2H 2017, but investors also are showing more caution about where they put their money.

BY JEFF DORSCH

200mm Fab Crunch

Shortages of used equipment and lower margins mean this problem isn't getting solved anytime soon.

BY MARK LAPEDUS

EUV's New Problem Areas

Random variations will require new methodologies, tools and cooperation among different companies.

BY MARK LAPEDUS

Designing Hardware For Security

Most attacks in the past focused on gaining access to software, but Meltdown and Spectre have changed that forever.

BY ED SPERLING

- John Schneider on The Battle To Embed The FPGA
- Tom on CEO Outlook On Chip Industry
- Gary Huang on Chip Dis-Integration
- Trelliant Fang on Advanced Packaging Confusion
- Dave Noble on CEO Outlook On Chip Industry
- Tony In Chicago on Chip Dis-Integration
- Mr Dadjeans on Delivering On The Promise Of Self-Driving Cars
- Del Hert on The Week In Review: Manufacturing
- Gary P on IBM Takes AI In Different Directions

Here are the main ones:

- Suppliers are shipping their previous-generation 3D NAND parts, which involve 32- and 48-layer devices. But the current transition to the next iterations—64- or 72-layer 3D NAND devices—has been relatively difficult.
- In R&D, suppliers are now working on the next products—96- and 128-layer 3D NAND. The complexity of these devices may force vendors to make some hard decisions and go down one of two paths. One is to scale 3D NAND, which is challenging. The other is a stacking technique that adds cost. Either way, the 3D NAND node cadence could extend from one to two years.
- Toshiba, the world's second largest NAND vendor, is looking to sell its memory unit amid financial troubles. Toshiba's woes haven't caused jitters in the supply chain—yet.
- China is spending billions of dollars to get a larger foothold in the IC market, including the development of 3D NAND. China isn't expected to be a factor for some time, however.

Why 3D NAND?

It's not all doom and gloom in the market, as 2017 is expected to be a big year for 3D NAND. "By the end of this year, the bit shipments for 3D NAND are expected to exceed 2D NAND," said Er-Xuan Ping, managing director of memory and materials within the Silicon Systems Group at [Applied Materials](#).

A year ago, there was a total of 350,000 to 400,000 wafer starts per month (wspm) in terms of worldwide installed capacity for 3D NAND, according to Applied Materials. That figure is expected to range from 500,000 to 700,000 wspm by year's end, according to the company.

Moreover, the conversion from planar to 3D NAND is causing a surge in capital spending. In total, capital spending for the overall flash memory segment is expected to reach \$23.6 billion in 2017, up 24% over 2016, according to Pacific Crest Securities. In total, semiconductor capital spending is projected to reach \$76.6 billion, up 18% over 2016, according to the firm.

"We continue to see strong momentum in 3D NAND," said Yang Pan, chief technology officer for the Global Products Group at [Lam Research](#). "It is the biggest driver for semi fab equipment."

For years, meanwhile, planar NAND has experienced enormous growth for use in data storage applications, such as flash drives, smartphones and SSDs. But suppliers are converting from planar to 3D NAND for several reasons.

Over the years, vendors have scaled the NAND cell size from 120nm to the 1xnm node regime today, enabling 100 times more capacity. At 1xnm, though, problems are beginning to crop up. "The floating gate is seeing an undesirable reduction in the capacitive coupling to the control gate," Objective Analysis' Handy said.

Basically, planar NAND is running out of steam. Plus, the overall wafer cost is lower for 3D NAND. In planar, a 300mm wafer costs \$1,200. There are 5.6-terabytes per wafer, translating to a cost-per-gigabyte of \$0.21, according to Objective Analysis. In 3D NAND, the wafer cost is \$2,000. But there are 17.2-terabytes per wafer, translating to a cost-per-gigabyte of

is lower for 3D NAND. In planar, a 300mm wafer costs \$1,200. There are 5.6-terabytes per wafer, translating to a cost-per-gigabyte of \$0.21, according to Objective Analysis. In 3D NAND, the wafer cost is \$2,000. But there are 17.2-terabytes per wafer, translating to a cost-per-gigabyte of \$0.12, according to the firm.

On top of that, 3D NAND has some advantages in systems. "The need for data is increasing," Applied's Ping said. "3D NAND is far superior over 2D NAND in terms of speed and reliability."

A big driver is in the data center, where SSDs based on 3D NAND are displacing traditional hard disk drives (HDDs). Generally, SSDs are more expensive than HDDs, but SSDs consume less power and reduce space. "NAND flash memory's transition from 2D to 3D technology has enabled enterprise SSD capacities already exceeding that of enterprise HDDs," according to Forward Insights. "The introduction of QLC (four-bit-per-cell) technology is expected to drive SSD capacities even higher."

3D NAND is seeing some new competition, however. One ReRAM-like technology, dubbed 3D XPoint, is supposedly faster than NAND.

Still, 3D NAND is expected to become a huge market. The shift towards 3D NAND started in 2013, when Samsung shipped the world's first 3D NAND device. Today, suppliers are shipping 32- and 48-layer devices with 64- and 72-layer chips just beginning to ramp up. 96 and 128 layers are in R&D.

"We will continuously push the limits of generations of industry-first V-NAND production, in moving the industry closer to the advent of the terabit V-NAND era," said Kye Hyun Kyung, executive vice president of the Flash Product and Technology team at [Samsung Electronics](#). "We will keep developing next-generation V-NAND products in sync with the global IT industry so that we can contribute to the timeliest launches of new systems and services, in bringing a higher level of satisfaction to consumers."

3D NAND SCALING




Fig. 4: 3D NAND flash roadmap. Source: Imec

In planar NAND, memory cells are connected via a horizontal string. In 3D NAND, though, the string is folded over and stood up vertically. In effect, the cells are stacked in a vertical fashion. Based on polysilicon, the string or strips are used for the wordlines in the structure. The bitlines run perpendicular to the wordlines.

The vertical stack has several levels or layers. The bit density increases as you add more layers, but it also brings more complexity into the fab.

or strips are used for the wordlines in the structure. The bitlines run perpendicular to the wordlines.

The vertical stack has several levels or layers. The bit density increases as you add more layers, but it also brings more complexity into the fab. "Generally, the move to 64-/72-layers has been challenging due to the increased capex intensity and slower yield ramp," said Greg Wong, an analyst with Forward Insights. "The yields are improving. Could they be better? Yes, but I don't see it as being a showstopper."

The challenges escalate as the industry moves beyond 64-/72-layer devices. And the bit density doesn't necessarily scale on the same curve. "Increasing the number of layers is getting more challenging, but there's at least four to five more generations to go for 3D NAND," Wong said. "You also get a huge jump in bits per wafer going from 2D to 3D NAND, but a lesser one going from one 3D NAND generation to the next. If you just relied on 3D-to-3D NAND conversions, it may not be enough to meet the demand. So there may be a need for some capacity increases beyond 2018."

How to make 3D NAND

In the fab, meanwhile, 3D NAND represents a departure from planar NAND. In 2D NAND, the fabrication process is dependent on scaling the dimensions of the memory cell using lithography.

Lithography is still used for 3D NAND, but it isn't the most critical step. So for 3D NAND, the challenges shift from lithography to deposition and etch.




Fig. 5: Key process steps in 3D NAND. Source: Lam Research

The 3D NAND flow starts with a substrate. Then, vendors undergo the first challenge in the flow—alternating stack deposition. Using chemical vapor deposition (CVD), alternating stack deposition involves a process of depositing and stacking thin films layer by layer on the substrate.

This process is much like making a layer cake. First, a layer of material is deposited on the substrate, followed by another layer on top. The process is repeated several times until a given device has the desired number of layers.

Each vendor uses a different set of materials to create a stack of layers. For example, Samsung deposits alternating layers of silicon nitride and silicon dioxide on the substrate.

In theory, a supplier can stack an unlimited number of layers. But as more

For example, Samsung deposits alternating layers of silicon nitride and silicon dioxide on the substrate.

In theory, a supplier can stack an unlimited number of layers. But as more layers are added, the challenge is to stack the layers with the exact thickness and good uniformities. It also must be done at high throughputs.

The big challenges are stress and defect control. “Throughout the 3D NAND manufacturing process, stress induced during film deposition needs to be carefully controlled,” Lam’s Pan said. “This is particularly important as the number of layers goes up.”

Film stack deposition




Fig. 6: Film stack deposition challenges. Source: Lam Research.

High-aspect ratio etch


Following the alternating stack deposition step, a carbon-based hard mask is applied on the film stack and holes are patterned on the top. Then, here comes the hardest part of the flow—high-aspect ratio (HAR) etch.

For 3D NAND, the etch tool must drill tiny circular holes or channels from the top of the device stack to the bottom substrate. The aspect ratio is around 40:1 for 32- and 48-layer devices with 64 layers moving to 60:1.

To illustrate the complexity of this step, Samsung’s 3D NAND device has 2.5 million tiny channels in the same chip. Each channel must be parallel and uniform. And each channel is three or more microns deep.

In this process, the tool etches holes using ions. But as the etch process penetrates deeper into the channels, the number of ions may decrease. This, in turn, slows down the etch rate. Even worse, unwanted CD variations may occur.

“High-aspect ratio channel hole formation continues to be the most critical and challenging module,” Lam’s Pan said. “Managing the fundamental trade-offs between profile, selectivity, and CD during HAR etch requires continuous technology and product innovation.”






Fig. 7: Channel etch challenges. Source: Lam Research.

Today, HAR etch and other tools are capable of developing 64-layer 3D NAND devices. It might be a different story for the next iterations—96- and 128-layer devices, as well as beyond.

“96 layers might be a tipping point,” Applied’s Ping said. “The current dielectric etcher may face challenges at 96 layers. The challenge is coming from the hard mask.”

At 96 and 128 layers, the HAR etcher must etch tiny holes in a deeper structure, roughly six or so microns in depth. Today’s etchers can perform this complex HAR etch task at 60:1, if there is no unwanted interactions with the current hard mask.

The problem? “Right now, the difficulty is the hard mask,” Ping said. “The hard mask is interacting with the dielectric etcher. You lose some of the energy and power from the etcher if you have a hard mask on top.”

Simply put, today’s HAR etch tools and hard mask materials may run out of steam at 96 or 128 layers. At this point, the industry faces some tough choices. Vendors could go down one of two paths to scale 3D NAND—single string or string stacking.

Single string versus string stacking

Today, many vendors are following the single string approach, where you stack all of the layers on a single string or die.

In contrast, string stacking involves a process of stacking individual 3D NAND devices on top of each other, which are separated by an insulating layer. For example, if one stacks two 64-layer 3D NAND devices on top of each other, the resulting chip would represent a 128-layer product.

“Some will choose to push the limit (with the single string approach). Others will choose string stacking. As the equipment and materials change, there is no single recipe,” Ping said.

Beyond 128 layers, though, the single string approach may hit the wall. By then, the industry may need to use string stacking or develop a new technique. “3D NAND is nowhere at the point that it will stop,” he said. “There are several more generations. I definitely see five more.”

Both the single string and string stacking have some advantages and disadvantages. With string stacking, 3D NAND vendors can reduce their risks. For example, a vendor has developed a mature 64-layer process. Then, the vendor can use the same process and equipment to develop and stack two separate 64-layer devices.

disadvantages. With string stacking, 3D NAND vendors can reduce their risks. For example, a vendor has developed a mature 64-layer process. Then, the vendor can use the same process and equipment to develop and stack two separate 64-layer devices.

But this also adds more cost to the equation. In that example, a supplier is doubling the number of steps to make a single device. And it's doubtful that suppliers can pass on these costs to customers, meaning a vendor could take a hit on their margins.

So others may continue with the single string approach at 64, 96 and 128 layers. String stacking is the back-up plan if this approach falters.

3D NAND based on a single string enables a less expensive part, as the process involves only one pass. But it also brings more challenges into the fab flow. For example, a 96-layer device is expected to be 1.2X to 1.3X taller than a 64-layer device, according to Applied Materials. To obtain the desired height, a vendor may reduce the thickness of each layer in the alternating deposition step.

But reducing the film thickness could also increase the resistance in the structure, especially for a 128-layer device. "One way is to reduce the film," Ping said. "But it's probably very difficult. Pair thickness reduction is hard from a device point of view."

Then, the stack must undergo a HAR etch step to create the channels. But as stated above, the industry faces a roadblock at 96 and/or 128 layers. It's difficult to perform HAR etch with today's hard masks. Therefore, the industry will require a next-generation hard mask at 96 layers and beyond. Even with a new hard mask, the industry may also require a new class of HAR etchers capable of 80:1, which represents the physical limit.

Another issue is channel mobility. In 3D NAND, the goal is to move the current through a polysilicon-based vertical channel. But polysilicon suffers from lackluster mobility and the degradation increases as the channel height becomes taller.

Consequently, the industry is exploring futuristic channel materials. "We are looking at III-V materials for the channel. If the roadmap continues, we may consider those materials (at 100 layers and beyond)," said Arnaud Furnemont, memory director at [IMEC](#).

Gates and metal fill

Following the HAR etch process, the next step is the formation of the gate. Samsung, SK Hynix and the Western Digital/Toshiba duo are making use of charge trap flash technology. This technology uses a non-conductive layer of silicon nitride. The layer wraps around the control gate of a cell, which, in turn, traps electrical charges to maintain cell integrity. In contrast, the Intel/Micron duo are not using charge trap. Instead, they have extended the floating gate structure to 3D NAND.

Once the gate is developed, the device requires a conductive tungsten wordline fill using a metal deposition step. "With the increasing number of layers, wordline fill of narrower, higher aspect ratio and more complex structures is another significant challenge," Lam's Pan said. "Atomic layer deposition, along with other process optimization, is required for next-

wordline fill using a metal deposition step. “With the increasing number of layers, wordline fill of narrower, higher aspect ratio and more complex structures is another significant challenge,” Lam’s Pan said. “Atomic layer deposition, along with other process optimization, is required for next-generation devices.”

For this step, the industry is migrating from [CVD](#) to [ALD](#). With CVD-based tungsten fill, the stress increases as the stack becomes taller. In addition, the fluorine content in the tungsten causes resistance. Meanwhile, ALD basically reduces the stress and fluorine content in the process.

All told, 3D NAND faces several challenges, although the industry continues to work on the solutions. “There are different ways for 3D NAND scaling, whether it’s more layers, more bits per cell, optimizing array efficiency, and/or optimizing peripherals,” Pan said. “This is going to be a multi-year roadmap.”

Related Stories

- [What’s Next For NAND?](#)
- [How To Make 3D NAND](#)



TAGS:	3D NAND	3D POINT	APPLIED MATERIALS	CHEMICAL VAPOR DEPOSITION	FLASH MEMORY	
FORWARD INSIGHTS	HIGH ASPECT RATIO ETCH	IMEC	INTEL	LAM RESEARCH	MICRON	NAND
OBJECTIVE-ANALYSIS	PACIFIC CREST SECURITIES	SAMSUNG	SK HYNIX	TOSHIBA		
WESTERN DIGITAL						



Mark LaPedus [\(all posts\)](#)

Mark LaPedus is Executive Editor for manufacturing at Semiconductor Engineering.

4 comments



Realjji Says:

June 21, 2017 at 1:12 am

Weird that everybody assumes that string stacking would be with 0 scaling, what would be the point? They need, let’s say, a 30% cost reduction from each gen , how they get there can be a combination of many factors. One would assume that string stacking is paired with horizontal scaling (including array efficiency) ,vertical (2 options, more layers same height and/or taller stack),improved throughput , yield and w/e else can be tuned for lower costs. So you either add lots of layers and that’s that or you go string stacking+ little bits here and there that get you to reasonable cost reductions.
If 4 bits per cell can end up serving a large enough market, that helps too at some point.

Reply



Gale Morrison Says:

June 21, 2017 at 8:15 am



Gale Morrison Says:

June 21, 2017 at 8:15 am

Excellent work Mark. Really enjoyed this.

Reply



Memister Says:

June 23, 2017 at 3:54 am

The addition of layers increases channel height, and so the read current is reduced inversely. Going to new channel materials to compensate seems to offset the cost reduction.

Reply



Kiers Says:

October 24, 2017 at 5:39 pm

Defect control is vital! Engineers must beware: "The curse of dimensionality" (adding another dimension). Caveat Fectum!

Reply

Leave a Reply

Comment

Name * (Note: This name will be displayed publicly)

Email * (Note: This will not be displayed publicly)

☐

Save my name and email in this browser for the next time I comment.

Post Comment

current is reduced inversely. Going to new channel materials to compensate seems to offset the cost reduction.

Reply

 Kiers Says:
October 24, 2017 at 5:39 pm

Defect control is vital! Engineers must beware: "The curse of dimensionality" (adding another dimension). Caveat Fectum!

Reply

Leave a Reply

Comment

Name * (Note: This name will be displayed publicly)

Email * (Note: This will not be displayed publicly)

Save my name and email in this browser for the next time I comment.

Post Comment

SEMICONDUCTORENGINEERING

ABOUT

About us

Contact us

Advertising on SemiEng

Newsletter SignUp

NAVIGATION

Homepage

Special Reports

System-Level Design

Low Power-High Perf

Manufacturing & Process Tech

Packaging, Test & Materials

IoT, Security & Automotive

KNOWLEDGE CENTERS

Knowledge Centers

Startup Corner

CONNECT WITH US

Facebook

Twitter @semiEngineering

LinkedIn

YouTube

Copyright ©2013-2018 SMG | Terms of Service | Privacy Policy

Document title: Semiconductor Engineering ... NAND Market Hits Speed Bump

Capture URL: <https://web.archive.org/web/20180625071314/https://semiengineering.com/nand-market-hits-speed-bump/>

Page loaded at (UTC): Fri, 15 Aug 2025 09:35:46 GMT

Capture timestamp (UTC): Fri, 15 Aug 2025 09:37:20 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: 12

Capture ID: vhfFhSFpty7FBKn7bFq3zx

Display Name: Alan.Man