

SEMICONDUCTOR ENGINEERING

Home > Manufacturing & Process Technology > 3D NAND Flash Wars Begin

MANUFACTURING & PROCESS TECHNOLOGY

3D NAND Flash Wars Begin



Market overcrowding, more efficient manufacturing, and growing list of scaling issues create a challenging competitive landscape.

AUGUST 20TH, 2018 - BY: MARK LAPEDUS



3D NAND suppliers are gearing up for a new battle amid a period of price and competitive pressures, racing each other to the next technology generations.

Competition is intensifying as a new player enters the 3D NAND market—China's Yangtze Memory Technologies Co. (YMTC). Backed by billions of dollars in funding from the Chinese government, YMTC recently introduced its first 3D NAND technology. The move is fueling concerns that a new entrant could impact a deteriorating market. The 3D NAND business is heading toward a prolonged period of oversupply and price erosion.

3D NAND is the successor to today's planar NAND flash memory, and is used for storage applications such as smartphones and solid-state storage drives (SSDs). Unlike planar NAND, which is a 2D structure, 3D NAND resembles a vertical skyscraper in which horizontal layers of memory cells 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

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.

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

- 2.5D
- 3D-IC
- 7NM
- 10NM
- ANSYS
- APPLE
- APPLIED MATERIALS
- ARM
- ARTERIS
- ATRENTA
- AUTOMOTIVE
- BUSINESS
- CADENCE
- EDA
- EMULATION
- ESILICON
- EUV
- FINFETS
- GLOBALFOUNDRIES
- IBM
- IMEC
- INTEL
- IOT
- IP
- LAM RESEARCH
- LITHOGRAPHY
- MEMORY
- MENTOR
- MENTOR GRAPHICS
- MIT
- MOORE'S LAW
- NVIDIA
- NOX
- QUALCOMM
- RAMBUS
- SAMSUNG
- SECURITY
- SEMI
- SIEMENS
- SOFTWARE
- SONICS
- STACKED DIE
- SYNOPSYS
- TSMC
- VERIFICATION

RECENT COMMENTS

- Rafael Della Giustina on Inspecting Unpatterned Wafers
- Minho Park on Chip Aging Becomes

YMTC
Exhibit 2026
Micron v. YMTC
IPR2025-00191
Ex. 2026-001




Fig. 2: 3D NAND architecture. Source: [Western Digital](#)

3D NAND is quantified by the number of layers stacked in a device. As more layers are added, the bit density increases. Today, 3D NAND suppliers are shipping 64-layer devices, although they are now ramping up the next technology generation, which has 96 layers. And behind the scenes vendors are racing to develop and ship the next iteration, 128-layer products, by mid-2019, analysts said.

And in R&D, vendors also are working on the next technology generations, which are 256 and 512 layers. "It is kind of a race," said Jeongdong Choe, an analyst at TechInsights. "It's a race for the highest number of stacks."


Some are deviating from the roadmap. In one scenario, vendors will eventually move to half nodes to stay ahead of the game. Then, YMTC, which is behind the competition, plans to ship a 64-layer device by mid-2019, but it will skip the 96-layer generation and move directly to 128 layers. "Their mission is to catch Samsung and others. Maybe in 2020 or 2021, they will do 128," Choe said.

The incumbent 3D NAND suppliers—Intel, Micron, Samsung, SK Hynix and Toshiba—aren't standing still, and they will battle to stay ahead in the scaling race. But each vendor is taking a different approach to scale 3D NAND.

Regardless, 3D NAND scaling is difficult. And migrating from 96 layers and beyond is even more daunting due to a slew of technical and cost challenges.

For 96 layers and beyond, 3D NAND suppliers may need to move toward both old and new techniques in the fab. In fact, there is a re-emergence of cryogenic etching, which first appeared in the 1980s. New bonding and other technologies are in the works.

3D NAND SCALING



Year	2016-2017	2018-2019	2020-2021	2022-2023
Generation 3D	L48	L64	L72	L128
Die size (30/cell)	256-512 Gb	512Gb - 1Tb	512Gb-2Tb	1-3Tb
Hole CD	65-100	65-100	65-100	65-100
Sit pitch (# holes)	4	4	4-8	8
Vertical pitch	50-70nm	40-60	40-50	40-50
BL CD	20	20	20-40	~40
Multiple stacks	No	No	No	Yes (2-4)

Fig. 3: 3D NAND flash roadmap. Source: [Imec](#)

Falling flash

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

FPGAs Becoming More SoC-Like

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

BY ANN STEFFORA MUTSCHLER

Quantum Computing Becoming Real

Technology has the potential to reshape processing everywhere, starting with limited scientific and commercial applications.

BY ED SPERLING

200mm Fab Crunch

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

BY MARK LAPEDUS

RF SOI Wars Begin

5G is driving up demand for both 300mm and 200mm capacity. Both are in short supply.

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

Machine Learning's Limits

Experts at the Table, part 1: Why machine learning works in some cases and not in others.

BY ED SPERLING

- Rafael Della Giustina on Inspecting Unpatterned Wafers
- Minh Park on Chip Aging Becomes Design Problem
- Mark LaPedus on Next-Gen Memory Ramping Up
- Semiconductor Engineering ... Next-Gen Memory Ramping Up ON Get Ready For Nanotube RAM
- Semiconductor Engineering ... 3D NAND Flash Wars Begin ON Next-Gen Memory Ramping Up
- Semiconductor Engineering ... 3D NAND Flash Wars Begin ON NAND Market Hits Speed Bump
- Mark LaPedus on Next-Gen Memory Ramping Up
- memory consultant on Next-Gen Memory Ramping Up
- Michael Mingliang Liu on Syntiant: Analog Deep Learning Chips
- Michael Mingliang Liu on Chip Aging Becomes Design Problem

Fig. 3: 3D NAND flash roadmap. Source: [Imec](#)

Falling flash

The business environment poses another challenge. Last year, the NAND market was beset by product shortages, supply chain issues and a difficult technology transition.

It's a different story today, as the [3D NAND](#) market is expected to “collapse at the end of this year,” said Jim Handy, an analyst at Objective Analysis. “Already, we are seeing some price declines. Spot market prices have been going down all year.”

The situation is different than many down cycles, which are characterized by weak demand and oversupply. “We are on the verge of oversupply,” Handy said. “The issue is that people are getting more efficient in making 3D NAND. It's supply-driven. There is no shortage of demand.”

For NAND in general, average selling prices (ASPs) are expected to fall by 24% in 2018 and 23% in 2019, according to Gartner. In total, NAND revenue is projected to reach \$58.7 billion in 2018, up from \$53.7 billion in 2017, according to Gartner.




Fig. 4: Q2 NAND Revenue Forecast Source: [Gartner](#)

In the long term, though, some forecasts are slightly more upbeat. “If you look at this from a top level, it's a healthy market,” said Simon Yang, chief executive at YMTC. “If you look at China's consumption of memory chips, it's a pretty big number.”

Meanwhile, semiconductor equipment makers are keeping a close eye on the market. Some vendors have experienced a slowdown in memory orders, but the overall market is expected to grow. In total, the wafer fab equipment market is projected to grow from \$51 billion in 2017 to between \$56 billion to \$58 billion in 2018, according to TEL. “(The) equipment market is moving to the next phase amid expanding applications for semiconductors,” said Toshiaki Kawai, president and chief executive of [TEL](#), in a recent presentation.

Besides the uncertain business climate, there are also challenges on the technology front. For years, the industry sold planar NAND devices for storage applications. NAND flash consists of a memory cell, which stores bits of data. The latest NAND devices store multiple bits of data (3 or 4 bits per cell). In NAND, the data remains stored even after the power is turned off in a system.

A planar NAND cell is based on a floating gate transistor structure. Over the years, vendors have scaled the cell size from 120nm to the 1xnm node

off in a system.

A planar NAND cell is based on a floating gate transistor structure. Over the years, vendors have scaled the cell size from 120nm to the 1xnm node regime today, enabling 100 times more capacity. At 15nm/14nm, though, planar NAND is running out of steam.

That's why the industry is moving to 3D NAND. In planar NAND, memory cells are connected via a horizontal string. In 3D NAND, the string is folded over and stood up vertically. In effect, the cells are stacked in a vertical fashion as a means to scale the density.

The vertical stack has several levels or layers. The bit density increases more layers are added. For example, Toshiba's 64-layer device (3-bit-per-cell) is a 512Gb device, which has a 65% larger capacity per unit chip size over its 48-layer chip.

Toshiba's latest 96-layer product (4-bit-per-cell) boasts a capacity of 1.33 terabits² with a 40% smaller die size over a 64-layer product. "QLC will have a game-changing impact across many different markets," said Scott Nelson, senior vice president of the Memory Business Unit at Toshiba.

Generally, suppliers are scaling 3D NAND roughly one technology generation every year. In 2018, suppliers are migrating from 64- to 96-layer products. Then, vendors are expected to move from 96 to 128 layers in 2019, followed by 256 layers in 2020/2021, and 512 in 2022/2023, according to Imec.

Others are following a different cadence. YMTC will move from 64 to 128 layers, thereby skipping 96. YMTC is skipping 96 layers for several reasons. First, a 64-layer device is price-competitive and will remain the sweet spot for some time. Then, from a density perspective, YMTC says that its 64-layer device is close to its competitors' 96-layer products.

"If you look at our current pace, we are moving pretty fast," YMTC's Yang said. "For the generation after 64, we still plan to have a time lag within 12 to 18 months. We plan on our next generation to go directly to 128. Based on that pace, we'll get very close on par (to others)."

The jump from 128 to 256 layers is not straightforward, though. Some will move to half nodes before taking the leap to 256 layers. For example, Samsung will move from 128 layers to somewhere around 180 or 190 layers, according to TechInsights.

Scaling 3D NAND

Regardless, to scale 3D NAND, suppliers are taking one of two approaches—single deck or string stacking. Both approaches are viable, but they are different, with various tradeoffs.

"The first way to scale these devices is to go to more and more layers. 96 is happening today. We see a path to do a single deck of pairs up to 256," said Rick Gottscho, CTO of [Lam Research](#), during a recent presentation. "The second way of scaling these devices is to take one deck and stack another deck on top. That creates a whole other set of challenges."

Samsung is embracing the single-deck approach. In its latest device, which

"The second way of scaling these devices is to take one deck and stack another deck on top. That creates a whole other set of challenges."

Samsung is embracing the single-deck approach. In its latest device, which is actually 92 layers, Samsung stacks all 92 layers in the same monolithic die, analysts said.

Others are taking the string-stacking approach. In a 64-layer device, for example, some developed two separate 32-layer parts. Then, they stacked one on top of the other, enabling a 64-layer chip.

Then, for 96 layers, some combine two separate 48-layer chips. In both cases, the two chips are separated by an insulating layer.

Both approaches, single deck and string stacking, are viable. "Double stacks are probably becoming more of the norm at 96. There might be some doing single stacks," said Mahendra Pakala, managing director of process development at [Applied Materials](#).

Each approach has some technical and cost issues. In string stacking, for example, a vendor is making two devices. In effect, the vendor is doubling the number of steps to make a single device, which translates into cost and cycle time.

In the single-deck approach, a vendor is making a single device in one shot. This, in theory, reduces cost and cycle time. But in the fab, the single-deck approach is difficult. Some believe this approach may run out of steam over time.

Both approaches follow the same process steps. In the fab, 3D NAND is different from planar NAND. In 2D NAND, the process is dependent on shrinking the dimensions 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.

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), the process involves depositing and stacking alternating thin films on the substrate.

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 different materials. For example, Samsung deposits alternating layers of silicon nitride and silicon dioxide on the substrate. "You deposit oxide-nitride or oxide-poly, depending on the kind of device you are fabricating," Lam's Gottscho said during the presentation.

It's possible to stack hundreds of layers on the substrate. But as more layers are added, the challenge is to stack the layers with the exact thickness and good uniformities at high throughputs. The big challenges are stress and defect control. In addition, the stack tends to bow under stress.

Film stack deposition

are stress and defect control. In addition, the stack tends to bow under stress.

Film stack deposition




Fig. 5: Film stack deposition challenges. Source: Lam Research.

That becomes more apparent in the single-deck approach. For this, the supplier would stack 96 layers of films on the substrate. “That’s a lot of deposition. If you look at any other device, such as traditional DRAM devices, logic devices or previous 2D NAND flash, they didn’t have 96 layers of deposited films,” Gottscho said.

There are solutions. For example, Lam has released a product that performs backside deposition, which compensates for front-side stress.

Another way to avoid stress is to use string stacking. For example, you deposit the layers on one 48-layer device, and then repeat the process on the other device, forming a 96-layer product.

Generally, a 48-layer alternating stack deposition process is mature and produces relatively less stress, but there are challenges. “You need to get one deck lined up with the other. If they both are highly deformed, you are going to have big alignment errors,” Gottscho said.

High-aspect ratio etch

Following this step, a 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 this, the etch tool must drill tiny circular holes or channels from the top of the device stack to the bottom substrate. The channels enable the cells to connect with one another in the vertical stack. A device may have 2.5 million tiny channels in the same chip. Each channel must be parallel and uniform.

This step is performed using today’s reactive ion etch (RIE) systems. In simple terms, the etcher creates tiny channels by bombarding the surface with ions. “That etch is very difficult and very time-consuming,” Lam’s Gottscho said. “There’s a fundamental law of aspect ratio scaling in etching that says the higher the aspect ratio, which is the deck of the deposited layers and the smaller the hole, the slower the etch.”

Then, as the etch process penetrates deeper into the channels, the number of ions may decrease. This slows down the etch rate. Even worse, unwanted CD variations may occur.

Then, as the etch process penetrates deeper into the channels, the number of ions may decrease. This slows down the etch rate. Even worse, unwanted CD variations may occur.




Fig. 6: Channel etch challenges. Source: Lam Research.

A 64-layer device has an aspect ratio of 60:1, compared to 40:1 for a 32-/48-layer device. Still, today's etchers can do the job, at least up to a point. "The 32-, 48- and 64-layer devices use conventional etching tools for the HAR channel hole," TechInsights' Choe said.

Based on this premise, it's conceivable that suppliers can migrate from 96 to 128 layers and beyond using string stacking. In theory, using traditional etch tools, a vendor could process two 64-layer devices, enabling 128 layers.

The single-deck approach is another story, as aspect ratios climb beyond 70:1. "For 96 layers, we can etch with a one-step etch. But you might have etch damage or the profiles are not good. If we use one-step etching, that's quite difficult," Choe said.

For a single-deck 96-layer device and beyond, the industry requires traditional etch tools for the HAR step. "However, another plasma tool and methods are needed. Cryogenic etching is an example," Choe said.

Traditional etchers involve a process of alternating etch and passivation steps at room temperature. In contrast, cryogenic etching is conducted at cryogenic temperatures. They use fluorine-based high-density plasmas.

"Cryogenic etch is not new. People have used it for other applications," Applied's Pakala said. "Atoms move around at high temperatures. If you don't want the atoms while etching, you reduce the temperatures."

cryogenic temperatures. They use fluorine-based high-density plasmas.


“Cryogenic etch is not new. People have used it for other applications,” Applied’s Pakala said. “Atoms move around at high temperatures. If you don’t want the atoms while etching, you reduce the temperatures.”

Cryogenic etching is difficult and expensive, however. “We are back to the future. What we’re doing is introducing cryogenic etching. It’s been in the literature since the mid-1980s, but it was very much ahead of its time,” Lam’s Gottscho said. “It’s a difficult technology, but we’ve made great progress. The advantage of cryogenic etching is that you get more reactants down at the etch front at the bottom of this high aspect ratio feature. That enhances the etch rate. That’s an expensive technology to implement, but the benefits outweigh those added costs.”

More steps

Following the process, each vendor follows a different flow. In some flows, the channel is lined with polysilicon and filled with silicon dioxide.

Then, the original nitride layers in the stack are removed. A gate dielectric is deposited, followed by a conductive metal gate fill using tungsten for the wordlines. That’s a simplified version of a complex process.






Fig. 7: 3D NAND process flow Source: [Objective Analysis](#)

Generally, this entire process is conducted in one continuous flow in the fab. A vendor will first take a substrate and build logic circuitry on top of it, followed by the NAND structure.

YMTC, however, has another approach. The company processes the circuitry on one wafer and the NAND structure on another wafer. Then, the two wafers are bonded and connected electrically using millions of metal vertical interconnect access structures. YMTC’s approach, dubbed Xtacking, reduces the manufacturing cycle time by 20% and allows for higher bit density.

It will take time before YMTC ramps up and moves into production, so the incumbent players will continue to dominant the competitive landscape for the foreseeable future.

To be sure, though, it’s a good time for OEMs. 3D NAND products will be plentiful at competitive prices.

Related Stories

- [NAND Market Hits Speed Bump](#)
- [Next-Gen Memory Ramping Up](#)



TAGS:

3D NAND

APPLIED MATERIALS

CHEMICAL VAPOR DEPOSITION

CHINA

CRYOGENIC ETCH

DEFECTIVITY

DEPOSITION

FLASH MEMORY

GARTNER

HAR ETCH

IMEC

INTEL

LAM RESEARCH

MICRON

OBJECTIVE-ANALYSIS

REACTIVE ION ETCH

SAMSUNG

SK HYNIX

SOLID STATE DRIVES

STRESS

STRING STACKING

TECHINSIGHTS

TEL

TOSHIBA

VARIATION

XTACKING

YANGTZE MEMORY TECHNOLOGIES CO.

YIELD

YMTC



Mark LaPedus (all posts)

Mark LaPedus is Executive Editor for manufacturing at Semiconductor Engineering.

Leave a Reply

Comment

INTERNET ARCHIVE

Wayback Machine

https://semiengineering.com/3d-nand-flash-wars-begin/

Go

JUL

AUG 22 2018

DEC

2020

About this capture

22 Aug 2018 - 24 Jun 2025

T

Twitter

LinkedIn

Google+

Email

Print

TAGS:

3D NAND

APPLIED MATERIALS

CHEMICAL VAPOR DEPOSITION

CHINA

CRYOGENIC ETCH

DEFECTIVITY

DEPOSITION

FLASH MEMORY

GARTNER

HAR ETCH

IMEC

INTEL

LAM RESEARCH

MICRON

OBJECTIVE-ANALYSIS

REACTIVE ION ETCH

SAMSUNG

SK HYNIX

SOLID STATE DRIVES

STRESS

STRING STACKING

TECHINSIGHTS

TEL

TOSHIBA

VARIATION

XTACKING

YANGTZE MEMORY TECHNOLOGIES CO.

YIELD

YMTC



Mark LaPedus (all posts)

Mark LaPedus is Executive Editor for manufacturing at Semiconductor Engineering.

Leave a Reply

Comment

Name *

Email *

Website

Post Comment

Connecting Cars Reliably

Avinash Ghirnikar

< Previous Article

Next Article >

Planning For A Digital Automotive Supply Chain

Robert Vamosi

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

CONNECT WITH US

Facebook

Twitter @semiEngineering

LinkedIn

YouTube

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

Document title:	Semiconductor Engineering ... 3D NAND Flash Wars Begin
Capture URL:	https://web.archive.org/web/20180822030129/https://semiengineering.com/3d-nand-flash-wars-begin/
Page loaded at (UTC):	Fri, 15 Aug 2025 09:34:31 GMT
Capture timestamp (UTC):	Fri, 15 Aug 2025 09:35:40 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:	11
Capture ID:	r4oymNqRtLPhvakJTKnevF
Display Name:	Alan.Man