

SEMICONDUCTOR ENGINEERING

Home > Manufacturing, Packaging & Materials > 3D NAND Race Faces Huge Tech And Cost Challenges

MANUFACTURING, PACKAGING & MATERIALS

3D NAND Race Faces Huge Tech And Cost Challenges



Shakeout looms as vendors struggle to find ways to add more layers and increase density.

MAY 23RD, 2019 - BY: MARK LAPEDUS



Amid the ongoing memory downturn, 3D NAND suppliers continue to race each other to the next technology generations with several challenges and a possible shakeout ahead.

Micron, Samsung, SK Hynix and the Toshiba-Western Digital duo are developing 3D NAND products at the next nodes on the roadmap, but the status of two others, Intel and China's Yangtze Memory Technologies Co. (YMTC), is less certain. Currently, Intel is evaluating its 3D NAND business amid losses in this market, and is mulling over the idea of finding a new NAND partner or exiting the market, analysts said. No decision has been reached. Meanwhile, it's unclear if YMTC will ship its initial 3D NAND products this year, as previously planned.

Nonetheless, the 3D NAND market could become a war of attrition amid technical and cost challenges. Some will keep up with the roadmap, which extends to at least 2024 and perhaps beyond, while others may fall behind or drop out of the race.

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: What is 3D NAND? Source: Lam Research

3D NAND is quantified by the number of layers stacked in a device. As

TECHNICAL PAPERS

Machine Learning Based Prediction: Health Behavior On BP

OCTOBER 12, 2018 BY TECHNICAL PAPER LINK

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

Knowledge Centers

Entities, people and technologies explored

LEARN MORE

RELATED ARTICLES

Engineering Talent Shortage Now Top Risk Factor

New market opportunities and global competitiveness are limited by qualified people.

BY MARK LAPEDUS

Making Chip Packaging Simpler

The promise of advanced packaging is being able to

Overall site

Search

SPONSORS



NEWSLETTER SIGNUP


Email:

Interests:

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

Subscribe

POPULAR TAGS



RECENT COMMENTS

- David Botma on Circuit Aging Becoming A Critical Consideration
- Charles R. on Circuit Aging YMTC Exhibit 2032 Micron v. YMTC IPR2025-00191 Ex. 2032-001

Fig. 1: What is 3D NAND? Source: Lam Research

3D NAND is quantified by the number of layers stacked in a device. As more layers are added, the bit density increases, enabling SSDs with more storage capacity. In 2013, [Samsung](#) shipped the world's first 3D NAND part, a 24-layer 128Gbit device. Today, suppliers are ramping up 96-layer devices (256Gbit) with the first 128-layer products (512Gbit) due out by late 2019.

Then, in 2021, vendors are expected to ship 192-layer devices with 256 layers in the works. "We're in a race," said Jeongdong Choe, an analyst at TechInsights. "It's a race for the highest number of stacks."

In R&D, vendors also are working on 500-layer 3D NAND, which is slated in the 2024 timeframe. The industry is also working on devices beyond 500 layers using new die stacking and bonding techniques. But to extend 3D NAND at 128 layers and beyond, vendors require new equipment and materials, more fabs, and billions of dollars in funding.

2D to 3D NAND

NAND is one of several technologies in the memory/storage hierarchy in systems. In the first tier, SRAM is integrated into the processor to enable fast data access. DRAM, the next tier, is used for main memory. Disk drives and NAND-based SSDs are used for storage.

NAND is a non-volatile memory that stores and retrieves data in a memory cell. Each cell can store multiple bits of data (3 or 4 bits). In NAND, the data remains stored even after the power is turned off in systems.

For years, the mainstream NAND technology was planar, based on a floating gate transistor structure. Over time, vendors scaled the cell size in planar NAND from 120nm to the 1xnm node, enabling 100 times more capacity.

Today, planar NAND is reaching its limit at the current 15nm/14nm node. "The floating gate is seeing an undesirable reduction in the capacitive coupling to the control gate," said Jim Handy, an analyst with Objective Analysis.

That's why the industry is migrating to 3D NAND. In planar NAND, a series of memory cells are connected together in 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 to scale the density.

"3D NAND flash memory has enabled a new generation of non-volatile solid-state storage useful in nearly every electronic device imaginable," said Timothy Yang, a software applications engineer at [Coventor, a Lam Research Company](#), in a blog. "3D NAND can achieve data densities exceeding those of 2D NAND structures, even when fabricated on later generation technology nodes. The methods used to increase storage capacity come with potentially significant tradeoffs in memory storage, structural stability, and electrical characteristics."

3D NAND has several benefits. "The first benefit to consider is that 3D

Making Chip Packaging Simpler

The promise of advanced packaging is being able to integrate heterogeneous chips, but a lot of work is needed to make that happen.

BY ED SPERLING AND MARK LAPEDUS

Finding Defects In Chips With Machine Learning

Better algorithms and more data could bolster adoption, particularly at advanced nodes.

BY MARK LAPEDUS

The Case For Embedded FPGAs Strengthens And Widens

Combining the flexibility of a FPGA with the performance and cost benefits of an SoC is pushing this technology well into the mainstream.

BY BRIAN BAILEY

Single Vs. Multi-Patterning EUV

Why this choice isn't as obvious as it might look.

BY MARK LAPEDUS

EUV Arrives, But More Issues Ahead

Improvement still needed for uptime, defectivity, line edge roughness and process flows.

BY KATHERINE DERBYSHIRE

3D NAND Metrology Challenges Growing

Rising costs and gaps in equipment emerge as technology scales; new tools under development.

BY MARK LAPEDUS

MicroLEDs: The Next Revolution In Displays?

Technology offers improved brightness, colors, and lower power, but they're expensive and difficult to manufacture.

BY MARK LAPEDUS

RECENT COMMENTS

- David Botma on Circuit Aging Becoming A Critical Consideration
- Charles R. on Circuit Aging Becoming A Critical Consideration
- Jon Peddie on Circuit Aging Becoming A Critical Consideration
- Tanj Bennett on 3D NAND Race Faces Huge Tech And Cost Challenges
- Brian Bailey on DAC 2019: Day 1
- steven Kim on 3D NAND Race Faces Huge Tech And Cost Challenges
- Danning on 3D NAND Race Faces Huge Tech And Cost Challenges
- Wayne Chen on Manufacturing Bits: June 10
- Mark LaPeducus on 3D NAND Race Faces Huge Tech And Cost Challenges
- Katherine Derbyshire on Partitioning In 3D
- Geoff Koch on In Automotive, A Move From Microcontrollers To Massively Complex SoCs
- david on Training Tomorrow's Chip Designers
- TimingsEverything on DAC 2019: Day 1
- Sagar Patel on The Growing Uncertainty Of Sign-Off At 7/5nm
- Sylvain Muckenhirn on MicroLEDs: The Next Revolution In Displays?
- Lydell Aaron on How To Integrate An Embedded FPGA
- Gopal P on Controlling Variability And Cost At 3nm And Beyond
- Gil Russell on How Far Can AI Go?
- Kevin Cameron on In-Memory Computing Challenges Come Into Focus
- Semi_Eng_Fan_4321 on Partitioning In 3D
- Brian Bailey on The Case For Embedded FPGAs Strengthens And Widens
- Richard F. Wahl on 3D NAND Race Faces Huge Tech And Cost Challenges
- Endric Schubert / Missing Link Electronics on The Case For Embedded FPGAs Strengthens And Widens
- Mritunjay Kumar on SiC Demand Growing Faster Than Supply
- Mark LaPeducus on 3D NAND Race Faces Huge Tech And Cost Challenges
- Eric Klien on 3D NAND Race Faces Huge Tech And Cost Challenges
- David Leary on What's After FinFETs?
- Stephen Fairbanks on The

generation technology nodes. The methods used to increase storage capacity come with potentially significant tradeoffs in memory storage, structural stability, and electrical characteristics.”

3D NAND has several benefits. “The first benefit to consider is that 3D NAND allowed a relatively seamless transition from MLC to TLC technology with very little impact to performance and endurance. The larger NAND cell created by the addition of the third dimension enables a path for higher densities needed by SSDs to meet the constantly growing demand for storage capacity. Another benefit to consider is the increased cell margin of 3D NAND. This gives NAND designers the flexibility to improve write and read times through enhanced array architectures as Vt placement become less critical,” said Daniel Doyle, senior manager for NAND component marketing for Micron. “In general, higher stacks increase capacity, or have lowered costs for the same capacity. However, NAND designers are working to enable higher speeds using innovative concepts that are leading to higher speeds as we scale capacity.”

Initially, vendors struggled in making these devices in the fab. Still, they managed to scale the technology from 24/32-layer 3D NAND chips in 2014, followed by 48- and 64-layer devices. And over time, they became more proficient in making it, which is why it’s become the mainstream NAND technology.

Today, suppliers are ramping up 96-layer 3D NAND. For example, a 96-layer device from the Toshiba-Western Digital duo is a 512Gbit device with bit densities of 5.95Gbit/mm2. In comparison, a 64-layer is a 256Gbit device with a die size of 75.2mm2 and a bit density of 3.40Gbit/mm2.

The next technology on the roadmap is 128 layers, which is due out by year’s end. Recently, Toshiba-WD described the world’s first 128-layer device, a triple-level-cell 512Gbit product with a bit density of 7.80Gbit/mm2. “128 might be possible this year, at the end of the third or early fourth quarter of this year, although this is a custom sample but not mass production. Mass production should be early next year. Then, you have 192. That might be three stacks,” TechInsights’ Choe said.

In 3D NAND scaling, though, the cost-per-bit benefit is less dramatic. “When you go to 96 layers, the cost reduction is maybe 10% to 15%. When you go to 128 layers, it’s may be another 5%,” said Handel Jones, chief executive of International Business Strategies (IBS).

The migration toward 96-layer devices and beyond will present some challenges. Compounding the challenges are the current business conditions. For more than a year, the NAND industry has been engulfed in an oversupply mode with falling prices.

NAND remains a tough market amid lackluster demand. “Our current forecast is for a 29% decline in the NAND flash market this year from \$59.4 billion in 2018 to \$42.2 billion in 2019,” said Bill McClean, president of IC Insights. “We forecast that bit volume will increase 35% this year after jumping by 28% last year. There is definitely some elasticity of demand at work here as lower per bit prices spur increased usage of NAND memory.”

Responding to the downturn, NAND vendors have scaled back their production of 3D NAND, hoping for a recovery later this year. “Without

Huge Tech And Cost Challenges

- David Leary on What's After FinFETs?
- Stephen Fairbanks on The Growing Uncertainty Of Sign-Off At 7/5nm
- John Wuethrich on Issues In Designing 5G Beamforming Antennas
- John Hallman on Can The Hardware Supply Chain Remain Secure?
- Clarisse GINET - CEO @Texplained on Can The Hardware Supply Chain Remain Secure?
- Capt. Obvious on Road Not Fully Constructed
- Mark LaPedus on Finding Defects In Chips With Machine Learning
- IQBAL SINGH Sabharwal on Cars, Security, And HW-SW Co-Design
- Applied_maths_guy on Finding Defects In Chips With Machine Learning
- Henry on EUV Mask Gaps And Issues
- Kevin Cameron on The Growing Uncertainty Of Sign-Off At 7/5nm
- masa on Return Of The Organic Interposer
- Carl on The Limits Of Energy Harvesting
- realjjj on The Limits Of Energy Harvesting
- SMS on Engineering Talent Shortage Now Top Risk Factor
- Alain Bensoussan on Gaps Emerge In Automotive Test
- Mdf on Quantum Effects At 7/5nm And Beyond
- Nitin on Finding Defects In Chips With Machine Learning
- Sam on Dude, Where's My Autonomous Car?
- Brian engineer on Make Your Own Energy
- Infrequent Flyer on Road Not Fully Constructed
- Jack McIntosh on Road Not Fully Constructed
- realjjj on Make Your Own Energy
- Mark LaPedus on Electric Cars Gain Traction, But Challenges Remain
- Suresh on China's Foundry Biz Takes Big Leap Forward
- Hank Walker on Electric Cars Gain Traction, But Challenges Remain
- Katherine Derbyshire on Finding The Source Of EUV Stochastic Effects
- Katherine Derbyshire on Finding The Source Of EUV Stochastic Effects
- C.S. Pawl on Understanding Memory

work here as lower per bit prices spur increased usage of NAND memory.”

Responding to the downturn, NAND vendors have scaled back their production of 3D NAND, hoping for a recovery later this year. “Without China’s entry, the market would have been likely to recover in 2021,” Objective Analysis’ Handy said. “China will probably be an important factor in the NAND flash market in 2021, and that will extend today’s downturn through that year, with no recovery until 2022.”

China’s YMTC is the wild card in 3D NAND. By year’s end, YMTC hopes to ship its initial product, a 64-layer device. YMTC could intensify the competition in the market, if it can execute. “My projection for YMTC is that they will struggle,” Handy said. “Although YMTC plans to create its own 3D NAND technology, I suspect that the company will eventually need the help of a partner who already understands 3D NAND volume production.”

Others are struggling for different reasons. Intel, for example, is losing money in 3D NAND, prompting it to rethink its position in the market.

For years, Intel and Micron were co-developing two types of memory technologies—3D NAND and 3D XPoint. 3D XPoint is a next-generation memory based on [phase-change](#) technology.

Recently, Intel and Micron have ended the memory alliance and are going their separate ways. While Intel will continue to develop 3D XPoint, it’s unclear if the company will move beyond its current 96-layer 3D NAND devices in the market.

“This year we expect to not be profitable in NAND,” said Robert Swan, [Intel’s](#) new chief executive during the company’s recent analyst meeting. “So we’re really evaluating the continued progress in NAND, whether the technology can bring down the cost curve. And we’ll evaluate that during the course of this year. We’re not going to put any more NAND capacity in place for the foreseeable future until we bring down the cost curves on 64- and 96-layer and beyond.”

Intel has not made a final decision. Nonetheless, the 3D NAND market is ripe for a shakeout. “There are too many vendors,” IBS’ Jones said. “There is no strategic advantage right now for Intel to be in 3D NAND. And if it’s losing money or even not contributing to their 60% gross profit margin or 20% operating income, why keep it? 3D XPoint is a different case.”

3D NAND scaling approaches

The other players, meanwhile, will move forward with 3D NAND at 128 layers and beyond, but it won’t be so simple. “Beyond 96 layers, we expect continued scaling with both an increase in layer count and a reduction in cell dimensions,” said Ceredig Roberts, senior technology director at Micron. “The major challenges in continuing to scale NAND will be maintaining cell performance and reliability as we scale the cell size. This includes mitigating the reduction in cell current and increased die and wafer stress levels of future nodes.”

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

The Source Of EUV Stochastic Effects

- Katherine Derbyshire on Finding The Source Of EUV Stochastic Effects
- C.S. Pawl on Understanding Memory
- Marc-David Levenson on Finding The Source Of EUV Stochastic Effects
- Priya Mukundhan on 3D NAND Metrology Challenges Growing
- DNR on 3D NAND Metrology Challenges Growing
- Barry Weavers on Finding The Source Of EUV Stochastic Effects
- Priya Mukundhan on 3D NAND Metrology Challenges Growing
- Twinkle Pandagre on VCSEL Technology Takes Off
- Brian Bailey on Target: 50% Reduction In Memory Power
- 10957 reader on Single Vs. Multi-Patterning EUV
- C. Hodson on Target: 50% Reduction In Memory Power
- Gil Russell on 3D NAND Metrology Challenges Growing
- neil on Target: 50% Reduction In Memory Power
- Computational Storage Dude ON In-Memory Computing Challenges Come Into Focus
- Mark LaPedus on 3D NAND Metrology Challenges Growing
- Peter manoj on 3D NAND Metrology Challenges Growing
- David Stoddard on Silicon CMOS Architecture For A Spin-based Quantum Computer
- Jacob Monberg on Making high-capacity data caches more efficient
- John Evans on The 7nm Pileup
- Nick T on Mathematic Model Helps Predict Markets That Will Drive Semiconductor Growth
- Enchuan on Single Vs. Multi-Patterning EUV
- Suresh Kumar K Regonda on Finding Defects In Chips With Machine Learning
- Chandra on Machine Learning Meets IC Design
- Sally, displaced and re-engineered self on Engineering Talent Shortage Now Top Risk Factor
- Subodh Taigor on The Case For Antifuse OTP NVM For Secure & Reliable SoCs
- Wolfgang Roethig on IC Test: Doing It At The Right Place At The Right Time
- Katherine Derbyshire on EUV Arrives, But More Issues Ahead

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.

To make 3D NAND, suppliers have several options. One of the first manufacturing decisions is to determine which scaling approach is the best path. For this, there are two approaches—single deck or string stacking.

In a 96-layer device, some are stacking all 96 layers on the same chip. This is referred to as the single-deck approach. Others are using string stacking. For example, in a 96-layer device, some are stacking two 48-layer devices on top of each other, which are separated with an insulating layer.

In the fab, string stacking is a relatively easier approach. In string stacking, though, a vendor is making two devices. In effect, the vendor is doubling the number of steps, which translates into cost and cycle time.

“Companies have different strategies. Some would rather go with existing equipment and then do multi-tier integration. Multi-tier integration requires more process steps, but they can come to the market quickly. Single-tier can save the number of process steps, but developing such equipment and processes will take a little bit of time,” said Gill Lee, managing director of memory technology at [Applied Materials](#).

At 128 layers, vendors will use both approaches. Most will stack two 64-layer devices on each other. In contrast, Samsung plans to use the single deck approach for 128 layers.

For now, 128 layers represent the limit for the single deck approach unless the industry comes up with a new breakthrough. So string stacking will become the norm beyond 128.

Beyond 128 layers, some vendors may stack two or more devices. For 192-layer devices, which are due out in 2021, vendors could string stack three 64-layer devices, according to TechInsights' Choe.

String stacking won't last forever and could run into issues at 500 layers. At this point, vendors are exploring another approach—die stacking. “It's kind of a die-on-die approach,” Choe said.

This involves stacking 3D NAND dies, which are connected using [through-silicon vias](#) (TSVs), he said. Wafer bonding is another approach. In theory, using these approaches, the industry could stack a 500-layer die on top of another one, and so on.

Deposition, etch challenge

It's not that simple, however. String or die stacking are only a part of the 3D NAND equation. Building a device involves an assortment of process steps and challenges.

Key Process Steps in 3D NAND Memory Cell Formation



- Doing It At The Right Place At The Right Time
- Katherine Derbyshire on EUV Arrives, But More Issues Ahead
- K. Goldberg on EUV Arrives, But More Issues Ahead
- jeffl on Making Chip Packaging Simpler
- Tejinder Singh on EUV Arrives, But More Issues Ahead
- Guest on EUV Arrives, But More Issues Ahead
- Riko R on A Different Kind Of Material World
- semi probes on The Mighty Sensor In The Fab
- Tony on Engineering Talent Shortage Now Top Risk Factor
- Pete Clarke on The Growing Challenge Of Thermal Guard-Banding
- Henry on More Lithography/Mask Challenges (part 1)
- Scott Jewler on Finding Defects In Chips With Machine Learning
- Mark LaPedus on Finding Defects In Chips With Machine Learning
- Carl Zweben on Shedding Pounds In Automotive Electronics
- John Eastwood on Making Autonomous Vehicles Safer
- Scott Jewler on Finding Defects In Chips With Machine Learning
- Navathej on Power Aware Intent And Structural Verification Of Low-Power Designs
- Guest on Defect Detection Strategies and Process Partitioning for SE EUV Patterning
- David J Leary on Chip Aging Accelerates
- nikilla on CMOS Image Sensors (CIS): Past, Present & Future
- Vojislav on In-Memory Computing Challenges Come Into Focus
- Brendan on Engineering Talent Shortage Now Top Risk Factor
- WILSON PARDI JUNIOR on Engineering Talent Shortage Now Top Risk Factor
- Ludo Vandenberg on Lithography Challenges For Fan-out
- Jeff on The Paradox Of Automotive Electronics
- Allen Rasafar on Transistor Options Beyond 3nm
- Allen Rasafar on Week In Review: Manufacturing, Test
- Ed on Engineering Talent Shortage Now Top Risk Factor
- Mark LaPedus on Embedded Phase Change Memory Emerges

Key Process Steps in 3D NAND Memory Cell Formation




Figure 2: 3D NAND memory and key process challenges. Source: Lam Research

The actual 3D NAND flow starts with a substrate. Then, using [chemical vapor deposition](#), vendors deposit 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. For its 9x-layer devices, Samsung uses the single deck approach, where it stacks all layers on the same substrate.

“When we talk about 96 layers, we are depositing actually twice that many, because there are pairs of oxide and nitride layers,” said Bart van Schravendijk, CTO for dielectrics at [Lam Research](#). “We are already depositing something like 192 layers. The key to those layers is they need to be very uniform, and more specifically, uniformity of the nitride layer becomes the key. That needs to be tightly controlled to enable the narrow threshold voltage distributions required for triple-level cell and quadruple-level cell. And from layer to layer, we then need to have extreme repeatability.”

As you add more layers to the stack, stress and defect control become more challenging. At 128 layers, these challenges escalate.

String stacking is another approach. In a 128-layer device, for example, you deposit 64 layers on two separate substrates and then connect them. A 192-layer chip might consist of three 64-layer devices.

This isn’t as easy as it seems. “The move beyond 128 layers will bring additional wafer shape requirements to handle high wafer bow and increased deck-to-deck overlay requirements,” said Scott Hoover, principal yield consultant at [KLA](#).

Following this step is 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.

The aspect ratios are 70:1 for a 96-layer device. Amazingly, 1 trillion tiny holes are etched on every wafer, according to Lam. Each channel must be parallel and uniform.

To accomplish this feat, a thin carbon-based material is first deposited on

- Manufacturing, Test
- Ed on Engineering Talent Shortage Now Top Risk Factor
 - Mark LaPedus on Embedded Phase-Change Memory Emerges
 - SStones on Engineering Talent Shortage Now Top Risk Factor
 - Américo Dias on Engineering Talent Shortage Now Top Risk Factor
 - Kati on Dodging The Next Generation Of Car Thieves
 - Santosh K Kurinec on Engineering Talent Shortage Now Top Risk Factor
 - Link on Engineering Talent Shortage Now Top Risk Factor
 - Andrew D Reynolds on Engineering Talent Shortage Now Top Risk Factor
 - P. Miguel Santos on Engineering Talent Shortage Now Top Risk Factor
 - CH on Using Sensor Data To Improve Yield And Uptime
 - David Silva on A Brief History of Test
 - Priya Mukundhan on Lithography Challenges For Fan-out
 - Mohammed Abdelwahid on The Paradox Of Automotive Electronics
 - Katherine Derbyshire on What's For Dinner?
 - dick freebird on Why Analog Designs Fail

The aspect ratios are 10:1 for a 50-layer device. Amazingly, 1 trillion tiny holes are etched on every wafer, according to Lam. Each channel must be parallel and uniform.

To accomplish this feat, a thin carbon-based material is first deposited on the stack. This material, called a hard mask, stabilizes the stack during the etch process.


Today's hard masks work. But as you increase the layer count, you need a thicker hard mask to reduce the stress, which could slow down the etch rate. Then, you may need a stronger hard mask like a pure diamond material, but this isn't feasible yet. So vendors must find ways to bolster today's carbon-based hard masks.

The next step is to pattern holes on the top of the hard mask. This seems simple, but pattern placement errors can crop up. "The placement issues may create etch slanting. This is also known as tilting, which makes it an even harder challenge to control the etch profiles and align the high-aspect ratio features amongst themselves, and to where they need to land," said Ofer Adan, director of metrology and process control at Applied Materials. "So it's increasingly important to maintain uniformity across the device CDs and their placement, as any slight deviation from a grid pattern could cause shorting or crosstalk between the memory devices."

After that comes the HAR etch process itself, which is conducted using today's reactive ion etch systems. In this two-step process, the etcher drills part of a tiny channel hole in the device. Then, the sidewall of the hole is passivated to prevent it from caving in. The process is repeated until a channel hole is drilled from the top of the stack to the substrate.

"The memory hole etch is probably the most difficult step in 3D NAND manufacturing. You need to etch many microns deep and you need to be able to tightly maintain the profiles to very specific dimensions," Lam's Schravendijk said. "When you are in that hole, you need to keep on digging. That's really the challenge. As you get deeper, you need neutrals that provide sidewall passivation, and you need ions on the bottom to dig deeper and deeper. As the aspect ratio increases, the number of ions and neutrals reaching the bottom tend to go down further and further."

As the etch process penetrates deeper into the channels, the etch rates tend to decrease. Even worse, unwanted CD variations may occur.






Fig. 3: Channel etch challenges in 3D NAND. Source: Lam Research

For a single deck process, today's HAR etchers will extend to 128-layers before the technology runs out of steam. To go beyond that, the industry is exploring cryogenic etch. [Cryogenic etch](#) is a one-step process, which simultaneously removes materials and passivates the sidewalls at cold temperatures. But it's unclear if this will work for 3D NAND. It's difficult to control and it requires specialized cryogenic gases in the fab.

The other option is string stacking. This appears to be easier, but the challenge is to align two or more stacks with one another. "With increasing stack height and the move to multi-deck structures, coupled with extreme wafer-level bow and in-die stress induced distortion, deck-to-deck channel hole alignment will be challenging," KLA's Hoover said.

From there, vendors have different flows. In some cases, the next step is called the staircase etch process, where you have a pattern that resembles a staircase on the sides of the device.

The staircase pattern is critical. This is how vendors eventually connect the peripheral logic on the bottom of the device to the control gates inside the chip. In this process, you pattern a small step, etch the structure and then trim it, then you repeat the process until you have the desired number of steps.

This is complex. A 96-layer device requires 12 lithography steps and 96 etch steps. A 128-layer device requires 128 etch steps and so on. "This series of process steps requires precise etch step profiling, trim etch uniformity and pull-back CD control for the WL (wordline) contact," said Steve Shih-Wei Wang, a process specialist at Lam, in a blog. "As you add more 3D NAND layers at a given cell density, the WL staircase also needs to lengthen and takes more space. For example, in the case of a 32-layer NAND device, the WL staircase stretches out 20um from the edge of the cell array. For a 128-layer architecture, the WL staircase would extend out 80um. Current WL staircase designs may be a key obstacle to cell efficiency and scaling of this type of 3D NAND architecture, due to this linear scaling effect."

More steps

The next step is to create columns next to the channel holes using an etch process. Slits are formed in the columns. Then, the original alternating layers of silicon nitride are removed. A silicon nitride charge trap material is deposited in the structure, which forms the gates.

The next step is to create columns next to the channel holes using an etch process. Slits are formed in the columns. Then, the original alternating layers of silicon nitride are removed. A silicon nitride charge trap material is deposited in the structure, which forms the gates.

Finally, the device is filled with a tungsten conductive metal gate material. “You get into these stack challenges, for example, like misalignment,” Lam’s Schravendijk said. “Misalignment then becomes an issue for the subsequent steps, where we want to fill the inside of the memory hole with solid material. If you have a void, it’s like having a hollow tree. Hollow trunks are how trees start to die. We prefer them to be filled, so preventing or minimizing any misalignment is key.”

Clearly, 3D NAND is a difficult technology. Still, vendors hope to move from one technology generation to the next almost every year. Each vendor wants to be first at each node. But not all will be able to keep up. In fact, it already looks like some have stumbled in the competitive landscape.



TAGS:

3D NAND

3D XPOINT

APPLIED MATERIALS

COVENTOR

CRYOGENIC ETCH

HIGH ASPECT RATIO ETCH

IC INSIGHTS

INTEL

INTERNATIONAL BUSINESS STRATEGIES

KLA

LAM RESEARCH

MICRON

OBJECTIVE-ANALYSIS

PHASE-CHANGE MEMORY

REACTIVE ION ETCH

SAMSUNG

SK HYNIX

STAIRCASE ETCH

TECHINSIGHTS

THROUGH-SILICON VIAS

TOSHIBA

TSVS

WESTERN DIGITAL

YANGTZE MEMORY TECHNOLOGIES CO.

YMTC

Mark LaPedus

(all posts)

Mark LaPedus is Executive Editor for manufacturing at Semiconductor Engineering.

7 comments

Eric Klien Says:

May 23, 2019 at 9:04 am

“In 2014, Samsung shipped the world’s first 3D NAND part, a 24-layer 128Gbit device.”

Samsung shipped 24 layers in 2013 and 32 layers in 2014.

Reply

Mark LaPedus Says:

May 23, 2019 at 1:39 pm

Hi Eric, Thanks. I changed the date. FYI. If you check Samsung’s releases, they went into “mass production” with a 24-layer part in 2014.

Reply

Richard F. Wahl Says:

May 23, 2019 at 4:02 pm

Reply

 *Richard F. Wahl Says:*
May 26, 2019 at 4:02 pm

Hello Mark,
Thanks for a very informative overview of the 3D NAND field. I'm curious what your thoughts are regarding YMTC and the trade war with America. How do you think their San Jose research location will be affected?

Thanks

Reply

 *Mark LaPedus Says:*
June 7, 2019 at 6:53 pm

Hi Richard. It's unclear right now. It's too early to say.

Reply

 *Danning Says:*
June 11, 2019 at 8:06 am

Very good article, thanks

Reply

 *Steven Kim Says:*
June 13, 2019 at 5:26 pm

Impressive and very informative to me. Tx.

Reply

 *Tanj Bennett Says:*
June 14, 2019 at 1:00 pm

It seems like this process should be asymptotic for cost per bit as the layers increase. Why should a 192 layer chip have cheaper bits than an 128 layer chip further down the experience curve and with an easier process? Those extra layers have got to be far more expensive than the silicon at the bottom. You can't even point to faster production pace per bit because the majority of the cycle time is just repeated layering and repeated step etch, which probably dwarfs the time to make the CMOS part on the base. The markets which demand more than 64GB per chip for some reason of packing density seem almost zero by now. So, why is the race to add layers, instead of race to make it cheaper by maturity of process?

Reply

Leave a Reply

[Reply](#)

Leave a Reply

Comment

[illegible]

Name*

(Note: This name will be displayed publicly)

Email*

(This will not be displayed publicly)

Post Comment

Post Comment

Manufacturing Bits: June 18

Mark LaPeduc

Power/Performance Bits: June 18

Jesse Allen

SEMICONDUCTORENGINEERING

ABOUT

About us

Contact us

Advertising on SemiEng

Newsletter SignUp

NAVIGATION

Homepage

Special Reports

System-Level Design

Low Power-High Perf

Manufacturing, Packaging & Materials

Test, Measurement & Analytics

IoT, Security & Automotive

CONNECT WITH US

Facebook

Twitter @semiEngineering

LinkedIn

YouTube

Copyright ©2013-2019 SMG | Terms of Service | Privacy Policy

Document title:	3D NAND Race Faces Huge Tech And Cost Challenges
Capture URL:	https://web.archive.org/web/20190619034609/https://semiengineering.com/3d-nand-race-faces-huge-tech-and-cost-challenges/
Page loaded at (UTC):	Fri, 15 Aug 2025 17:43:05 GMT
Capture timestamp (UTC):	Fri, 15 Aug 2025 17:45:09 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:	13
Capture ID:	u6Yb9XtLb5jds8ib7knrLx
Display Name:	dani.green