

involves a process of stacking individual 3D NAND devices on top of each other. For example, if one stacks three 64-layer 3D NAND devices on top of each other, the resulting chip would represent a 192-layer product. The trick is to link the individual 64-layer devices with some type of interconnect scheme.

String stacking is already in the works. For example, Micron Technology, according to multiple sources, recently demonstrated a 64-layer 3D NAND device by stringing two 32-layer chips together.

This is not a simple technology to develop, however. And even with string stacking, 3D NAND would top out at or around 300 layers, according to experts.

All told, 3D NAND is projected to remain viable at least until 2020 and perhaps beyond. “This is a 10-plus year roadmap and we are just at the beginning of it,” said Yang Pan, chief technology officer for the Global Products Group at [Lam Research](#).

In any case, OEMs will need to get a handle on the 3D NAND manufacturing issues in order to have more realistic expectations about their design schedules. To help OEMs, Semiconductor Engineering has taken a look at some of the more challenging process steps for 3D NAND. This includes alternating step deposition, high aspect ratio etch, metal deposition and string stacking.

Why 3D NAND?

In today systems, the memory hierarchy is fairly straightforward. SRAM is integrated into the processor for cache. DRAM is used for main memory. And disk drives and NAND-based SSDs are used for storage.

NAND, a nonvolatile memory technology, is based on the traditional floating gate transistor structure. Thanks to 193nm immersion lithography and multiple patterning, vendors have extended planar NAND down to the 1xnm node regime.

But at 1xnm, there are issues cropping up. “In fact, the floating gate is seeing an undesirable reduction in the capacitive coupling to the control gate,” said Jim Handy, an analyst with Objective Analysis.

So, today’s planar NAND will soon stop scaling, prompting the need for 3D NAND. Basically, 3D NAND resembles a vertical skyscraper or layer cake. The layers, which are horizontal, are the active wordlines. “The bitlines also run horizontally in the metal layers on the top of the chip,” Handy said. “The vertical channels are the NAND ‘strings’ that attach to the bitlines.”

Meanwhile, vendors are at various stages of ramping up the technology. Samsung, the leader in 3D NAND, last year shipped its third-generation 3D NAND device—a 48-layer chip. In addition, Micron and its 3D NAND partner, Intel, have recently begun shipping their first 3D NAND chip—a 32-layer device. Both SK Hynix and the SanDisk/Toshiba duo are separately sampling 48-layer chips.

2016 is expected to be a big year for 3D NAND. At the end of 2015, there was a total of 160,000 wafer starts per month (wspm) in terms of worldwide installed capacity for 3D NAND, according to Lam Research.

10nm Versus 7nm

The economics and benefits of moving to the next process node are not so obvious anymore.

BY MARK LAPEDUS

7nm Lithography Choices

Four possible scenarios for patterning the next generation of chips.

BY MARK LAPEDUS

FinFET Scaling Reaches Thermal Limit

Advancing to the next process nodes will not produce the same performance improvements as in the past.

BY BRIAN BAILEY

Pain Points At 7nm

More rules, variation and data expected as IP and tools begin early qualification process.

BY ED SPERLING

What’s Next For DRAM?

As DRAM scaling runs out of steam, vendors begin looking at alternative packaging and new memory types and architectures.

BY MARK LAPEDUS

The Week In Review: Manufacturing

Intel buys multi-beam firm; Intel supplier awards; RF SOI; image sensors.

BY MARK LAPEDUS

7nm Fab Challenges

FinFET formation, mask challenges and back-end-of-line issues will make this node difficult and expensive.

BY MARK LAPEDUS

Alternative To X86, ARM Architectures?

Support grows for RISC-V open-source instruction set architecture.

BY JEFF DORSCH

- AdeleHars on Will Little Guys Win IoT Race?
- J.Chonan on Save Power And Area By Eliminating Redundant Resets
- Mike Bruzzone on Executive Insight: Simon Segars
- Jesse Enjaian on Lawyers, Insurance And Self-Driving Cars
- Kev on DAC Day One: EDA Through Different Glasses
- Dev Gupta on Fab Investment Increases In China
- Philippe Faes on Preparations for DAC
- Dev Gupta on Electromigration: Not Just Copper Anymore
- TheDoktor42 on The Numbers Game
- Ed Sperling on Will Little Guys Win IoT Race?

2016 is expected to be a big year for 3D NAND. At the end of 2015, there was a total of 160,000 wafer starts per month (wspm) in terms of worldwide installed capacity for 3D NAND, according to Lam Research. "We estimate that the industry will ship approximately 350,000 to 400,000 wspm of 3D NAND capable capacity by the end of 2016," Lam's Pan said.

Still, 3D NAND represents a fraction of the total installed capacity for NAND (2D and 3D), which is roughly 1.3 million to 1.4 million wspm. "Eventually, we expect a significant majority of the NAND installed base to become 3D capable," Pan said.

Vendors are ramping up these devices in new or converted 3D NAND fabs. In total, the equipment cost for a 2D NAND fab ranges from \$30 million to \$45 million for 1,000 wspm, according to Christian Dieseldorff, an analyst with the Industry Research & Statistics group at [SEMI](#).

In comparison, the equipment cost for a 3D NAND fab ranges from \$50 million to \$65 million for 1,000 wspm, Dieseldorff said. "3D NAND equipment costs are higher because more equipment like CVD and etch tools are needed," he said.

Some vendors are retrofitting their current fabs into 3D NAND facilities. "We expect 2X to 4X more space is needed when converting from 2D to 3D. In this case, there is a high degree of re-use because most equipment is already there. Again, additional CVD and etch tools are needed," he said.

Still, a 3D NAND fab is not as expensive as a leading-edge logic fab. A 7nm logic processes, for example, will require a \$160 million fab equipment investment for every 1,000 wspm, according to Gartner.

The new litho: alternating stack deposition

In any case, 3D NAND represents a major departure from today's planar NAND. In 2D NAND, the fabrication process is dependent on advanced lithography. In 3D NAND, though, vendors are using trailing-edge 40nm to 20nm design rules. Lithography is still used, but it isn't the most critical step. So for 3D NAND, the challenges shift from lithography to deposition and etch.

In fact, 3D NAND introduces a number of new and difficult process steps to the semiconductor industry. "By moving the bit-string into the third-dimension, this technology eases many of the patterning-scaling challenges," said David Fried, chief technology officer at [Coventor](#). "But it has introduced several fairly complex and new processes. Uniformity of these processes is critical. So, from my perspective, the challenges here are focused on variability control of several key processes."

The 3D NAND flow starts with a substrate. Then, vendors undergo the first major 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. In simple terms, a layer of material is deposited on the substrate. Then, another layer of material is deposited on top of that. The process is repeated several times until a given device has the desired number of layers.

This process is much like making a layer cake. In simple terms, a layer of material is deposited on the substrate. Then, another layer of material is deposited on top of that. 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, to make its 3D NAND devices, Samsung deposits alternating layers of silicon nitride and silicon dioxide on the substrate, according to Objective Analysis. In contrast, Toshiba's 3D NAND technology consists of alternating layers of conductive polysilicon and insulating silicon dioxide, according to the firm.

Alternating stack deposition must have precise control with good uniformities and low defectivity. "Initially, the uniformities must be good," Applied's Ping said. "It's all going back to stress control because the alternating films are different. For each film there could be a mismatch. Stress could show up."

The challenges escalate as vendors increase the number of layers in a device. "Repeatability at every single step is also critical and it has to be done at high productivity in order to keep costs down," Lam's Pan said.

High aspect ratio etch

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

Tiny trenches or channels are etched from the top of the device to the substrate. To illustrate the complexity of this step, Samsung's 3D NAND device has 2.5 million tiny channels in the same chip. Each of them must be parallel and uniform.

Today's high-aspect ratio etch tools can handle the requirements for 32- and 48-layer devices. For these chips, the aspect ratios range from 30:1 to 40:1. "This etch is really complex. Uniformity is absolutely critical to the performance of the memory device," Coventor's Fried said. "The statistics are also staggering. Once the etch is complete, the amount of processing that takes place inside that hole is also pretty impressive."

The problem? Current high aspect ratio etch tools are either not ready or struggling to meet the demands for 64-layer devices and beyond. At 64 layers, the aspect ratios are 60:1 to 70:1. "This is too high for current etching capability," Applied's Ping said. "The etching and hard mask technologies are not necessarily available for 60:1 or 70:1."

So going forward, NAND vendors are simultaneously following two paths. First, they will wait for the next-generation high-aspect ratio etch tools and other technologies to arrive. And then, provided the etchers are ready on time, they may scale today's 3D NAND in the following progressions—32 and 48 layers, to 64 layers, to 96, and then to 128.

In the second path, NAND vendors also will develop next-generation string stacking technology (see below for details).

Charge trap versus floating gate

Before moving to string stacking, vendors will continue to scale today's 3D

stacking technology (see below for details).

Charge trap versus floating gate

Before moving to string stacking, vendors will continue to scale today's 3D NAND. Besides deposition and etch, today's 3D NAND undergoes other complex steps, including the formation of the gate.

For this, the industry is moving in two directions. Samsung, SK Hynix and the SanDisk/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. "In floating gate, the gate is actually a conductor," Objective Analysis' Handy said. "A charge trap layer, which actually looks like a floating gate, is an insulator."

Floating gate involves some difficult patterning steps. "It's hard to pattern things on the sides of a vertical hole that you've made. You have to go through a lot of process steps," Handy said.

Charge trap also has some drawbacks. "The advantage with charge trap is that you don't have to pattern it. Charge trap is easier to make that way," he said. "Excluding Spansion, which ships over 80% of all bytes in charge trap, nobody else has been able to make charge trap cost effectively."

Metal deposition

Once the gate is developed, the next step is difficult. The device requires contacts. The device is backfilled with a conductor using a metal deposition step.

"There is a challenge in the metal deposition area," said Dave Hemker, senior vice president and chief technology officer at Lam Research. "We're seeing a lot of customers' backfilling it with tungsten. And that's a tricky deposition, because you are doing a non-line-of-sight deposition. So you basically have these caves and tunnels in there. You have to go back in there after the fact and put in tungsten metal. If you don't engineer the process right, you may put in this pre-cursor that wants to plate out metallic tungsten. Given its own way, it could plate out right when it gets into the hole. So you have a lot of ways to create voids."

String stacking

There are other difficult steps in the flow, but the biggest challenge is clear. Until the industry solves the high aspect ratio etch issues, today's single-string 3D NAND technology is arguably stuck at 48 and/or 64 layers.

And even when the etchers are ready, today's single-string 3D NAND hits the wall at 128 layers. "This is because the aspect ratio is limited by the process," Applied's Ping said. "So you must figure out a way to bypass the limitations."

So what's the answer? String stacking. In this approach, vendors will stack individual 3D NAND devices. Each device might be separated by an insulating layer. "When you do string stacking, you finish one string," Ping said. "Then, you are repeating the steps. It's difficult, but you can do it."

So what's the answer? String stacking. In this approach, vendors will stack individual 3D NAND devices. Each device might be separated by an insulating layer. "When you do string stacking, you finish one string," Ping said. "Then, you are repeating the steps. It's difficult, but you can do it."

For example, a vendor will develop a 48-layer device. To devise that chip, it will go through the same process flow, such as alternating layer deposition, etch and others.

Then, the vendor will develop a separate 48-layer chip using the same flow. The process is not limited to 48-layer chips. A vendor could also stack multiple 32-layer chips. And if the technology is available, a vendor could stack 64-, 96- and perhaps 128-layer devices.

In theory, though, vendors may opt for string stacking with 32- and 48-layer chips. There is less stress for an individual 32- or 48-layer device, as compared to a 96- or 128-layer chip.

Ultimately, though, 3D NAND with string stacking may run out of steam at or near 300 layers. "It will hit a problem in terms of yield," Ping said. "When you stack, the yield loss from defects continues building up. That will be the limitation. Plus, everything will be limited by stress. If you put too much film, then the stress presents a limitation."

Still to be seen, however, is how vendors will connect the individual 3D NAND devices together in string stack. For this, the industry is looking at various interconnect schemes. "There will be four or five different options," he said. "You can build a shared bit line in the middle. Then, another option is that you build a string, which contacts each string directly."

To be sure, though, there are still many unknowns and challenges with string stacking. The industry also faces several challenges even without string stacking. In either case, the industry must continue to master and perfect the various process steps with 3D NAND. Otherwise, the technology will remain costly, at least for the vast majority of OEMs.

Related stories

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

TAGS:

3D NAND

APPLIED MATERIALS

CHARGE TRAP FLASH

CHEMICAL VAPOR DEPOSITION

COVENTOR

CVD

DEPOSITION

ETCH

FLOATING GATE

GARTNER

INTEL

LAM RESEARCH

MICRON

OBJECTIVE-ANALYSIS

SAMSUNG

SANDISK

SEMI

SK HYNIX

SPANSION

STRING STACKING

THIN FILMS

TOSHIBA



Mark LaPedus (all posts)

Mark LaPedus is Executive Editor for manufacturing at Semiconductor Engineering.

you stack, the yield loss from defects continues building up. That will be the limitation. Plus, everything will be limited by stress. If you put too much film, then the stress presents a limitation.”

Still to be seen, however, is how vendors will connect the individual 3D NAND devices together in string stack. For this, the industry is looking at various interconnect schemes. “There will be four or five different options,” he said. “You can build a shared bit line in the middle. Then, another option is that you build a string, which contacts each string directly.”

To be sure, though, there are still many unknowns and challenges with string stacking. The industry also faces several challenges even without string stacking. In either case, the industry must continue to master and perfect the various process steps with 3D NAND. Otherwise, the technology will remain costly, at least for the vast majority of OEMs.

Related stories

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

TAGS:

3D NAND

APPLIED MATERIALS

CHARGE TRAP FLASH

CHEMICAL VAPOR DEPOSITION

COVENTOR

CVD

DEPOSITION

ETCH

FLOATING GATE

GARTNER

INTEL

LAM RESEARCH

MICRON

OBJECTIVE-ANALYSIS

SAMSUNG

SANDISK

SEMI

SK HYNIX

SPANSION

STRING STACKING

THIN FILMS

TOSHIBA



Mark LaPedus (all posts)

Mark LaPedus is Executive Editor for manufacturing at Semiconductor Engineering.

SEMICONDUCTORENGINEERING

ABOUT

About us

Contact us

Advertising on SemiEng

NAVIGATION

Homepage

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 ... How To Make 3D NAND
Capture URL:	https://web.archive.org/web/20160613061552/https://semiengineering.com/how-to-make-3d-nand/
Page loaded at (UTC):	Fri, 15 Aug 2025 17:52:17 GMT
Capture timestamp (UTC):	Fri, 15 Aug 2025 17:53:12 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:	dVVTzK18ZVSCE8eB7m6g4d
Display Name:	dani.green