Highlights

48 Hours After CEO Discloses $2.25B Shortage, Silicon Valley Bank Collapses

Posted by Karen Heumann, March 14, 2023

Silicon Valley Bank (SVB) surprised investors with news that the bank needed $2.25B to correct its financial deficiencies. The news caused panic and a massive sell-off resulting in the second-biggest bank collapse in US history. Customers withdrew $42B by the end of the next day, according to a California regulatory filing, decimating the remains of the 40-year-old investment bank, and leaving SVB with a negative cash balance of $958M. The California Department of Financial Protection and Innovation closed SVB, seized remaining cash deposits, and named the FDIC as receiver. The FDIC created the Deposit Insurance National Bank of Santa Clara to hold insured deposits from SVB.

Ryan Falvey, Managing Partner at Restive Ventures and a former SVB employee, pointed to the highly interconnected nature of the tech investing community as a key reason for the bank's sudden demise. "This was a hysteria-induced bank run caused by VCs," He explained. "This is going to go down as one of the ultimate cases of an industry cutting its nose off to spite its face."

Prominent funds including Union Square Ventures, Founders Fund, and Coatue advised their entire rosters of startups to pull funds out of SVB on concerns of a bank run. Anna Nitschke, Chief Financial Officer at Pear VC, an early-stage VC firm based in San Francisco, urged its portfolio network to withdraw money. "In light of the situation with Silicon Valley Bank that we are sure all of you are watching unfold, we wanted to reach out and recommend that you move any cash deposits you may have with SVB to another banking platform."

London-based VC firm, Hoxton Ventures, advised founders to withdraw two months' worth of "burn," or venture capital they'd use to finance overhead. Hussein Kanji, Hoxton's founding partner said, "We have seen some funds passing on a view that they remain confident in SVB. We are seeing other funds encouraging companies to withdraw their funds from SVB. It remains to be seen how this will all play out. "If the self-fulfilling prophecy occurs, the risks to you are asymmetric." Kanji said: "The big danger for startups is that their accounts will be frozen while the mess is being sorted."

Falvey noted. "When you say, `Hey, get your deposits out, this thing is gonna fail,' that's like yelling fire in a crowded theater," Falvey said. "It's a self-fulfilling prophecy." Falvey said that his analysis of SVB's mid-quarter update gave him confidence. The bank was well capitalized and could make all depositors whole, he said. He counseled his portfolio companies to keep their funds at SVB.

Ryan Falvey
@Ryan_Falvey

What we sent our founders this morning. We're standing by to help.

Restive Ventures @restivevc · 22m

How we're advising our founders today.

Dear Founders,

As you know, SVB failed this morning. A lightly redacted version of the note that we sent our LPs is below. Please keep this in your confidence but we wanted you all to know how we have communicated to our investors, if that's helpful to you.

As you can see we were affected by this and no, we're not worried. We'll get the deposits back and, if you're a client of SVB, you will too.

This is our advice:

1. If you're frantically trying to get deposits out, stop. The FDIC is in charge now. You will have, at a minimum, access to $250k on Monday. We suspect it will actually be all of it, or the vast majority, then or shortly thereafter. There is, unfortunately, no way to get money out of SVB until then.
2. Try to calm your teams and customers. Banks fail. The goal of our regulatory framework is to prevent that failure from harming the broader economy. We are the broader economy.
3. If you have immediate ST liquidity needs, reach out to us. We have significant cash on hand outside of SVB and are prepared to help.
4. If you are an SVB borrower or use SVB for any service besides holding funds, reach out to us. The FDIC has one goal: to repay depositors. If you used SVB for payments, FBO services, lines of credit, warehouse lines, venture debt, card services, treasury services or anything else, please let us know so we can help get you set up on a more stable solution.
5. If you have major counterparties who may rely on SVB for liquidity, please let us know. Ironically, that's probably other venture funds.

We all got into fintech because we saw opportunities to improve the financial services industry. The failure of this institution shows, in many ways, how much opportunity exists. However, as we told our LPs and industry partners, the single great strength of the U.S. banking industry is its safety and soundness. The highest principle of that is maintaining deposits.

We're sorry this is happening but we'll all come out stronger.

Best,

Ryan, Ty, Cameron & Kate

Another venture investor, TSVC partner Spencer Greene, also criticized investors who "were wrong on the facts" about SVB's position. "It appears to me that there was no liquidity issue until a couple of VCs called it," Greene said. "They were irresponsible, and then it became self-fulfilling."

Some SVB customers received emails assuring them that it was "business as usual" at the bank. "I'm sure you've been hearing some buzz about SVB in the markets today so wanted to reach out to provide some context," one SVB banker wrote to a client, according to a copy of the message obtained by CNBC. "It is business as usual at SVB," the banker wrote. "Understandably there may be questions and I want to make myself available if you have any concerns." SVB customers said CEO Greg Becker didn't instill confidence when he urged them to "stay calm" during a call. The stock's collapse reached 60% by the end of regular trading. Importantly, Becker couldn't assure listeners that the capital raise would be the bank's last, said a person on the call.

Customers still remaining with SVB have no idea if, when, and how they will get their money back. While insured deposits are expected to be available as early as Monday, the majority of deposits held by SVB were uninsured. A recent 10-K filing showed more than 90 percent of its deposits were uninsured, and the FDIC says today that "At the time of closing, the amount of deposits in excess of the insurance limits was undetermined." "The precipitous deposit withdrawal has caused the Bank to be incapable of paying its obligations as they come due," the California financial regulator stated. "The bank is now insolvent." The FDIC standard insurance covers up to $250k per depositor, per bank, for each account ownership category. The FDIC said uninsured depositors will get receivership certificates for their balances. The regulator said it will pay uninsured depositors an advanced dividend within the next week, with potential additional dividend payments as the regulator sells SVB's assets.

As of the end of December, SVB had $209B in total assets and $175.4B in total deposits. The FDIC is unclear what portion of those deposits were above the insurance limit. The last U.S. bank failure of this size was Washington Mutual in 2008, which had $307B in assets. Whether depositors with more than $250k ultimately get all their money back will be determined by the amount of money the regulator gets as it sells Silicon Valley Bank assets or if another bank takes ownership of the remaining assets. There are concerns in the tech community that until that process unfolds, some companies may have issues making payroll. Many CEOs are trying to transfer their funds out of SVB and set up corporate accounts at other banks. Avni Patel Thompson, Founder and CEO of Milo, said she was "shaking with relief" after she was able to move her company's funds. Nearly half of the U.S. technology and health care companies that went public last year after getting their early funding from venture capital firms were Silicon Valley Bank customers.

Modernizing the Universe with KIOXIA SSDs

KIOXIA

Posted by Karen Heumann, March 14, 2023

Innovation is boundless and KIOXIA never disappoints. At the cutting edge of technology is exploration of outer space and the potential the limitless universe provides. KIOXIA provides the tools that level up that exploration, innovation, and the accompanying scientific research.

HPE Spaceborne Computer-2 Program features KIOXIA SSDs in the HPE Edgeline Converged Edge System and HPE ProLiant Server to enable research science experiments at the International Space Station (ISS). The HPE Spaceborne Computer-2, the first in-space commercial edge computing and AI-enabled system to run on the ISS, is part of a greater mission to significantly advance computing and reduce dependency on communications as space exploration continues to expand. For example, astronauts can achieve increased autonomy by processing data directly on the ISS, eliminating the need to send raw data to Earth to be processed, analyzed and sent back to space.

As a sponsor of the HPE SBC-2, KIOXIA has provided flash-based SSDs, including KIOXIA RM Series Value SAS and KIOXIA XG Series NVMe™ SSDs, to enable these solutions. These flash-based SSDs are better-suited than traditional hard disk drive storage to withstand the power, performance and reliability requirements of outer space, as they have no moving parts, are less susceptible to electromagnetic waves and provide faster performance.

“Proving that data center-level compute processing can successfully operate in the harsh conditions of space will truly take something special,” noted Scott Nelson, Exec VP and Chief Marketing Officer for

KIOXIA America, Inc. “The synergies that exist when KIOXIA and HPE collaborate to leverage our respective technologies, allows us to explore and study at the very edge of scientific discovery. We can’t wait to see where the HPE Spaceborne Computer journey takes us.

KIOXIA has been collaborating with HPE to create best-in-class storage solutions for years, and the company’s products enable a broad range of HPE solutions, from mobile to data center to enterprise. Value SAS SSDs are part of the KIOXIA Life After SATA campaign, enabling customers to easily transition away from aging SATA SSDs, while delivering higher performance and reliability.
Designed to perform various high-performance computing tasks in space, including real-time image processing, deep learning, and scientific simulations, the HPE SBC-2 utilizes a combination of HPE’s edge computing solutions, including the HPE Edgeline Converged Edge System, a rugged and compact system, and the HPE ProLiant server for high-performing capabilities. The HPE SBC-2 targets a range of workloads and has already helped advance progress in healthcare, image processing, natural disaster recovery, 3D printing, 5G, AI, and more.

“One of the challenges with missions to the ISS and to places like Mars and other bodies is the latency between Earth and the spacecraft. Bringing more compute onboard means that more processing, for AI, for example, can be done locally at dramatically lower latency. The HPE Spaceborne Computer project’s impact, if successful, will help bring more modern technology to space instead of traditional lower-speed and lower-power hardened silicon,” Cliff Robinson, STH.

“It is an exciting time for Hewlett Packard Enterprise as we continue to play an important role in the expanding space economy. We are pleased to continue our longstanding collaboration with KIOXIA and partner together on our space computing initiatives to bring its storage solutions to the International Space Station with us,” said Jim Jackson, Chief Marketing Officer, at HPE. “By bringing KIOXIA’s expertise and its SSDs, one of the industry’s leading NAND flash capabilities, with HPE Spaceborne Computer-2, together we are pushing the boundaries of scientific discovery and innovation at the most extreme edge.”

The Race to Increase Flash Density While Avoiding Impacts to Flash Life & Performance

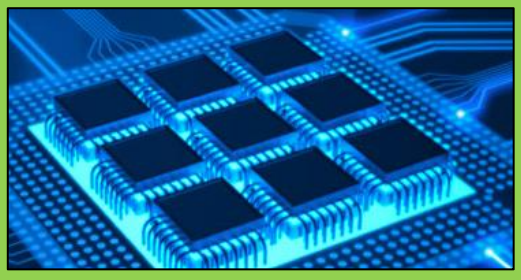

Posted by Mike Heumann, March 14, 2023

Flash is the de facto storage medium for online storage today, whether in enterprise storage systems and servers, consumer laptops and tablets, or smartphones and digital cameras. The speed and predictable wear characteristics of flash are far superior to that of hard disk drives (HDDs), and the density is roughly the same as hard disk drives, if not better when new form factors such as Enterprise and Data Center Standard Form Factor (EDSFF) are considered. The one place that flash has trailed HDDs is cost per byte. Because of this, the leading manufacturers of flash are continuing to innovate on storage density, which because of the nature of semiconductor electronics, drives down flash cost.

"Moore's Law for processors has arguably been lagging the last few years, but is alive and well for NAND flash," said Ben Whitehead, Technical Product Manager at Siemens EDA. "It's a good thing, because modern compute and networking have an insatiable appetite for fast storage." "For DRAM, it took something like 10 or 15 years of R&D to come to fruition, but for 3D NAND, the development was extremely fast. It's astonishing when you think of the usual pace of development," said Xi-Wei Lin, Semiconductor Technology at Synopsys. "Besides the technology itself, it's a killer app. Apple was first to put in a flash memory, to store data. Today, we buy iPhones still based on how much memory, and it's all flash. From there, big data, AI, and also analytics require high performance computation. Flash memory is filling this critical gap in latency between the hard disc drive and RAM memory. You can see the applications, especially in the data center, analytics, and gaming, because of the power, form factor, and the density cost."

Semiconductor manufacturers competing to increase memory chip density and reduce cost per are taking varying approaches to solve the bit density issue. For SK Hynix, the preferred approach is stacking layers. Intel showcased their three-dimensional NAND flash-memory chip at this year's International Solid State Circuits Conference (ISSCC). The chip stores five bits of data in each NAND flash cell. The 192-layer chip has 23 gigabit per square millimeter density, with the capacity to store 1.67 terabits of data. SK Hynix offers a 1-Tb NAND flash-memory chip with 300 layers, stores 3 bits per cell (called triple layer cell, or TLC), and has the highest write speed, 194 megabytes per second. The next few articles will focus on the efforts of leading flash vendors to increase density and reduce costs, while maintaining the performance and wear characteristics that make flash superior.

Approach 1:
Adding (Lots More) Layers to Flash ASICs

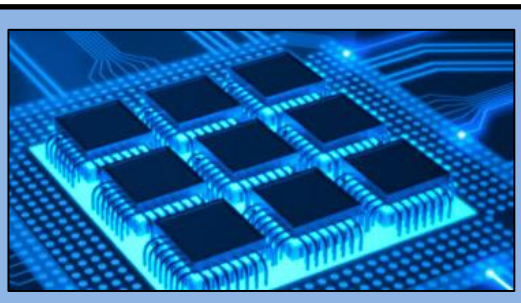

Increase the number of layers, and you increase planar density, right? Adding layers isn't free, but it can help to drive down costs, which is why nearly every flash manufacturer is looking into increasing the number of layers in a flash ASIC. Micron Technology was the first chip to pass the 200-layer mark, with its 232-layer NAND flash-memory chip last year. Micron also claimed the industry's fastest NAND I/O speed – 2.4 Gbps – and up to 100% higher write bandwidth and more than 75% higher read bandwidth per die than the prior generation. In addition, 232-layer NAND contains six-plane TLC production NAND, that Micron said is the most planes per die of any TLC flash and is capable of independent read capability in each plane. According to industry analysts, this may be the most impressive part of the innovation. Because of the six planes, this chip can behave as if it's six different chips.

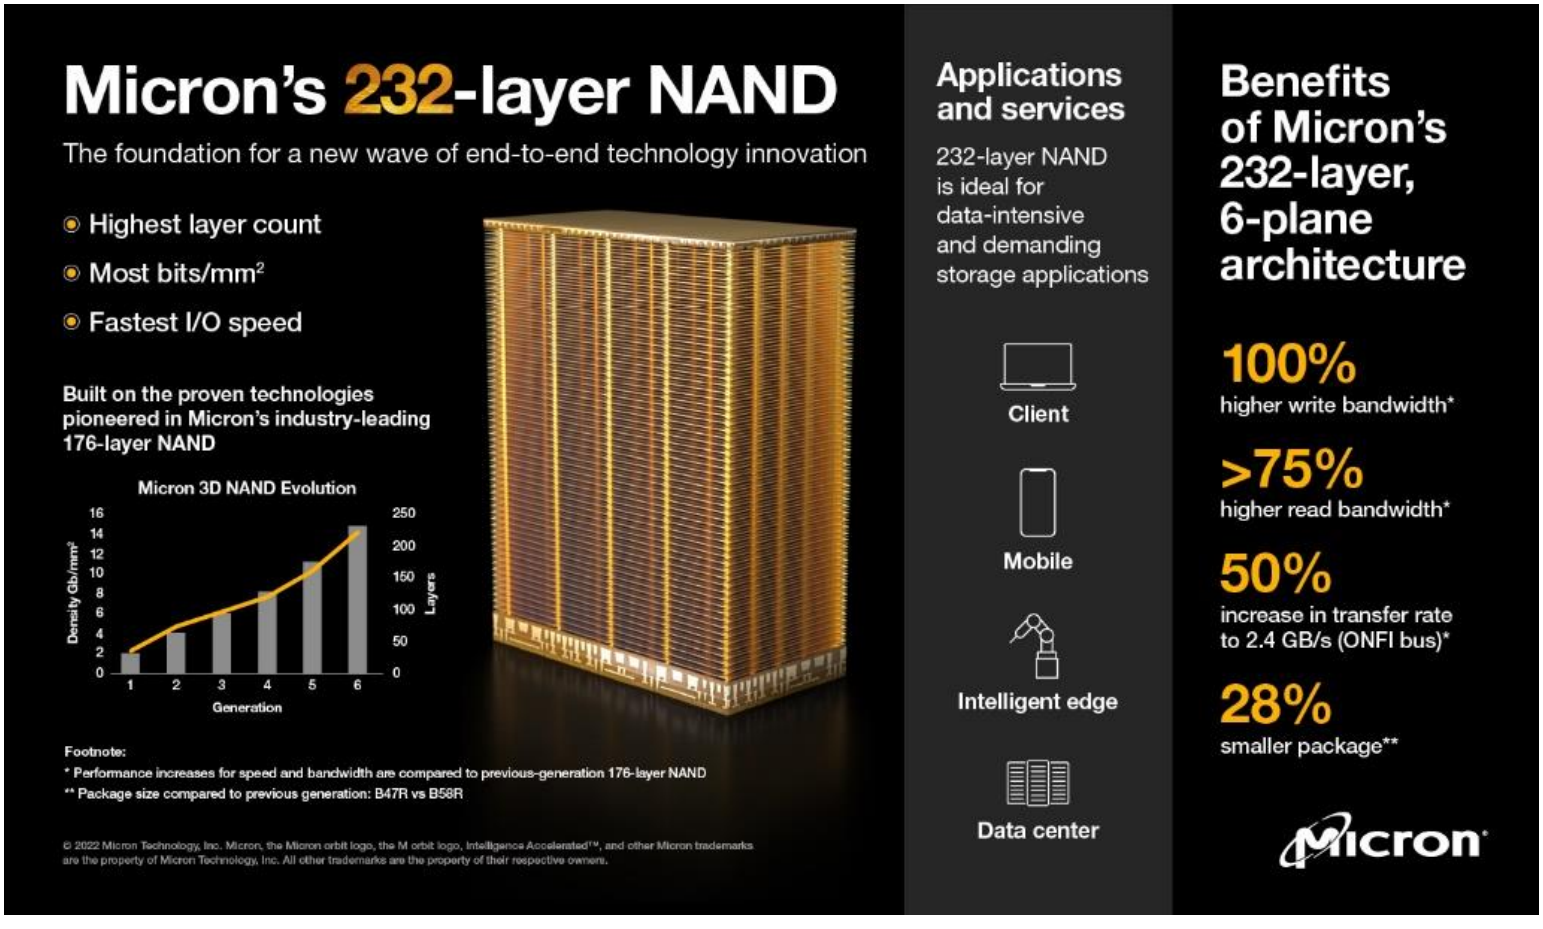

Fig. 3: Micron's 232-Layer NAND. Source: Micron

Samsung was first to market with "V-NAND" in 2013 which introduced alternating layers of polysilicon and silicon dioxide and swapped the floating gate for charge trap flash (CTF). FGs store memories in a conducting layer, while CTFs "trap" charges within a dielectric layer. The CTF design quickly became preferred because of manufacturing cost reductions, but certainly not the only one. Samsung's 6th generation V-NAND added 40 percent more cells to the 9x-layer (5th) generation and their V-NAND production process also now reduces the number of channel holes needed to connect the cells of a 256Gb die from over 930 million to 670 million. First generation 3D NAND design kept the peripheral circuitry to the side but eventually, 3D NAND vendors moved the peripheral circuitry under the CTF. In SK Hynix's terminology, it was now the Periphery Under Cell (PUC) layer, another variation of 3D NAND, with a smaller cell area per unit. At Samsung Tech Day 2022, Samsung unveiled its roadmap for the SSD ecosystem, and it promises significant progress. The company plans to "reach over 1,000 layers" in its NAND chip by 2030, though it did not say whether that would apply to consumer products.

Another approach is what is called "4-D NAND Flash", which is being pushed heavily by SK Hynix. Originally developed for smartphones, SK Hynix has achieved 238 layers on their latest 4-D TLC NAND flash device. The devices apply charge trap flash and peri-under-cell technologies to create what SK Hynix calls 4-dimensional flash structures, with a smaller call area per unit when compared to 3D flash ASICs. The 4D chips also achieve data transfer speeds of 2.4Gb/sec (50% over their previous flash chips). These devices have a capacity of 512Gb, and are expected to start mass production in the first half of 2023. The first use cases will be for client SSDs in PCs, followed by deployment for smartphones and SSDs for enterprise servers. SK Hynix expects to reach capacities of 1Tb/device in late 2023.

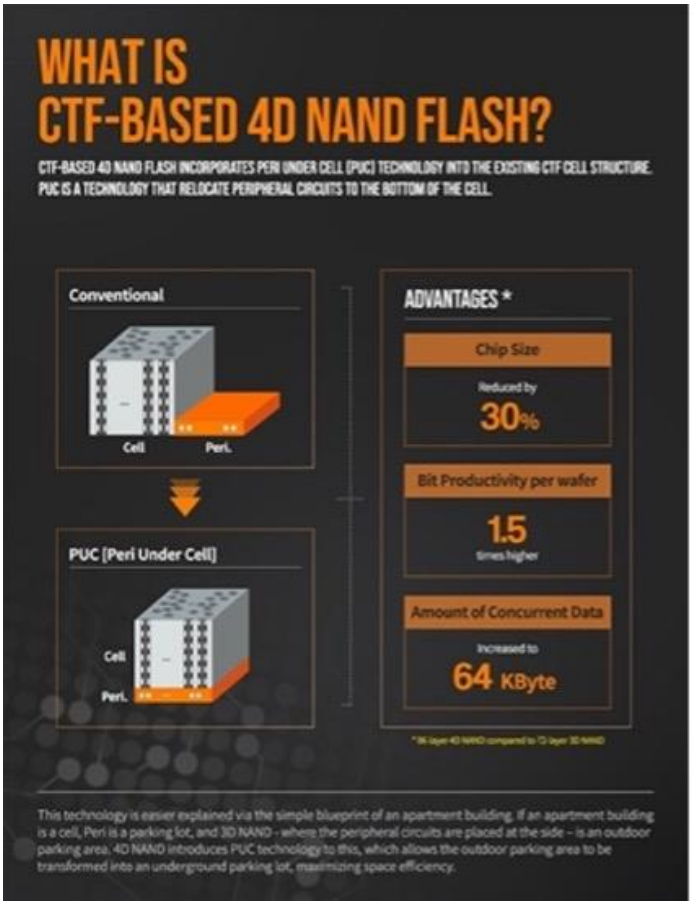

SK Hynix's explanation for 4D NAND. Source: SK Hynix global newsroom.

The Difficulties of Adding (Lots More) Layers to Flash ASICs

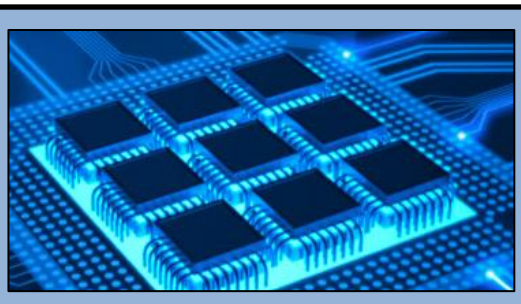

Further stacking seems reasonable, except for the unavoidable problem at the heart of the whole process. “The main challenges are in etch, because you have to etch very deep holes with a very high aspect ratio,” said Maarten Rosmeulen, imec. “If you look at the previous generation with 128 layers, this was about a 6-, 7-, or 8 micrometers deep hole of only about 120 nanometers diameter, extremely high-aspect ratio — or maybe a little bit higher, but not that much. There are advances in the etch technology to etch deeper holes in one go, but it won’t go faster. You can’t increase the speed of the etch. So if the process flow get dominated by the deposition and etch, and those process steps don’t increase in cost efficiency, then adding more layers is not as efficient anymore to reduce the cost.”

“Besides the etch, you also need to fill this hole with a very thin dielectric layer uniformly up and down,” said Xi-Wei Lin, Synopsys. “Typically, to deposit a layer of a few nanometers is not easy because of the chemistry of the wafer. Here, they have to go all the way down to be able to fill. There are sub-atomic layer deposition methods, but it’s still challenging. Another big challenge is stress. If you build up so many layers that go through a etch/deposition/cleaning/thermal cycle, that can cause stress locally and globally. Locally, because after you drill a hole, you need to cut a very deep trench through the full

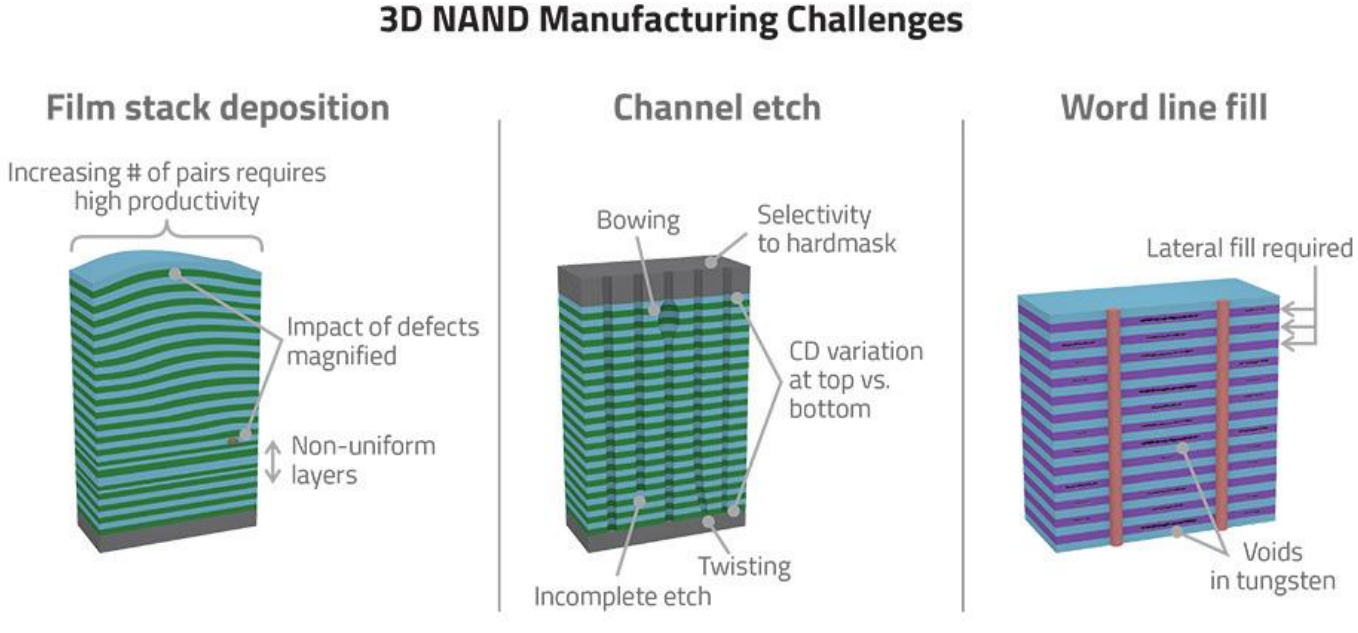

stack. It becomes a really high-aspect skyscraper, which is wobbly. And if you start going through some washing or other processes, a lot of things can happen to cause two skyscrapers to collapse against

each other. So then you've lost the yield. By putting so many materials on top of each other, and cutting different patterns, this can create global stress and can cause a wafer to warp, which will make it impossible to handle in the fab because a wafer has to be flat. And that's just for starters. Remember that the etch is going through layers of different materials."

At the 2021 ISSCC, Samsung presented the leading write throughput at 184 MS.s for a 3 bits per cell NAND flash memory and their chairman, Kinam Kim, predicted a 1000-layer flash could be possible by 2030. Achieving that is challenging because flash cells are made by etching deep, narrow holes through alternating layers of conductor and insulator, and then filling the holes with dielectric and other materials. Etching and filling deep-enough holes reliably and quickly through an increasing number of layers is a key limit to the technology. When the number of stack layers passes 300, it becomes increasingly challenging to improve NAND-memory performance because each layer in the stack has to be made thinner, increasing resistance, which introduces errors and reduces read and write speeds. Hynix used five different techniques to achieve high write throughput with 300 layers.

"Connecting many pipes with an increasingly fast (but never fast enough) host interface creates bottlenecks in very unexpected places," said Benjamin Whitehead, Technical Product Manager at Siemens. "Another design verification challenge is power. It has long been a lower priority by most storage controllers, but this has shifted now to a critical feature. Moving to smaller geometry nodes helps some but is expensive. Business models are intolerant of re-spins, not to mention the supply chain difficulty in getting in the queue. Time-to-market delays get a lot of visibility to upper management. There are even more growth drivers for storage, which require us to rethink how we verify designs. AI accelerators require much larger storage controllers, which may quickly consume your emulation and prototyping capabilities. Edge intelligence requires orders of magnitude more complex design verification. In-memory computing, like CSD, requires testing new processor combinations that mix RTOS and HTOS with previously unseen workloads."

Imec is exploring potential new structures for 3D NAND. It uses "trench architecture," a design variant in which the memory cells are part of the sidewall of a trench, with two transistors at opposite ends of the trench. Jan Van Houdt, imec fellow, explained its value: "The 3D trench architecture has the potential of double density as compared to the currently used gate-all-around (or cylindrical) architecture." However, he went on to point out a few drawbacks. "There are two high-aspect ratio (=challenging) etch steps instead of one, as well as a lower electric field in the tunnel oxide in the case of flash. The second drawback is not there when using ferroelectric FETs, which makes the trench version more appealing for ferro than for flash." The design is still in the prototype phase.

Approach 2:
Increasing the Number of Bits Per Cell Beyond Three

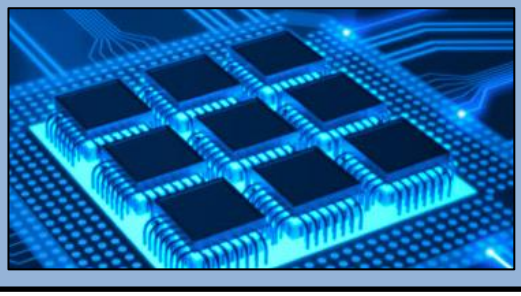

Today, triple level cell (TLC) is the most widely used NAND flash technology. TLD essentially stores three bits in each cell. Increasing this to 4, 5, or more bits per cell obviously would increase flash density and (hopefully) reduce flash cost per bit. Intel developed its new high-density 5-bit-per-cell chip using floating gate NAND cell technology which stores bits in a conducting layer. Most other manufacturers use charge-trap flash which stores charges in a dielectric layer, reducing manufacturing cost. Intel implemented special fast-read algorithms to overcome any concerns about lower endurance and speed. The new chip can be operated in either a 3-bit-per-cell or 4-bit-per-cell mode.

Layers or density? Note that the discussion was lost regarding Intel increasing density to contemplation of the degree to which layers can successfully be added. “When it comes to 3D technology, what sets one supplier’s offerings apart from another? It’s a question I’ve discussed often with customers and partners looking to accurately assess performance requirements for end applications. The answer is a bit complicated – but one thing that doesn’t factor into the equation is number of layers. It’s true. While it’s certainly understandable to look to number of layers as a technology gauge – after all, suppliers have promoted this metric since 3D flash memory devices were first introduced – getting to the most layers the fastest adds cost and is far less important than harnessing the greatest ROI and delivering a cost-effective, competitive storage solution.” Scott Nelson, Executive Vice President & CMO, KIOXIA.

Increasing the number of bits per cell can give bigger and cheaper memory, but it affects performance by compromising read and write speeds. This is why adding layers is still seen as one of the most

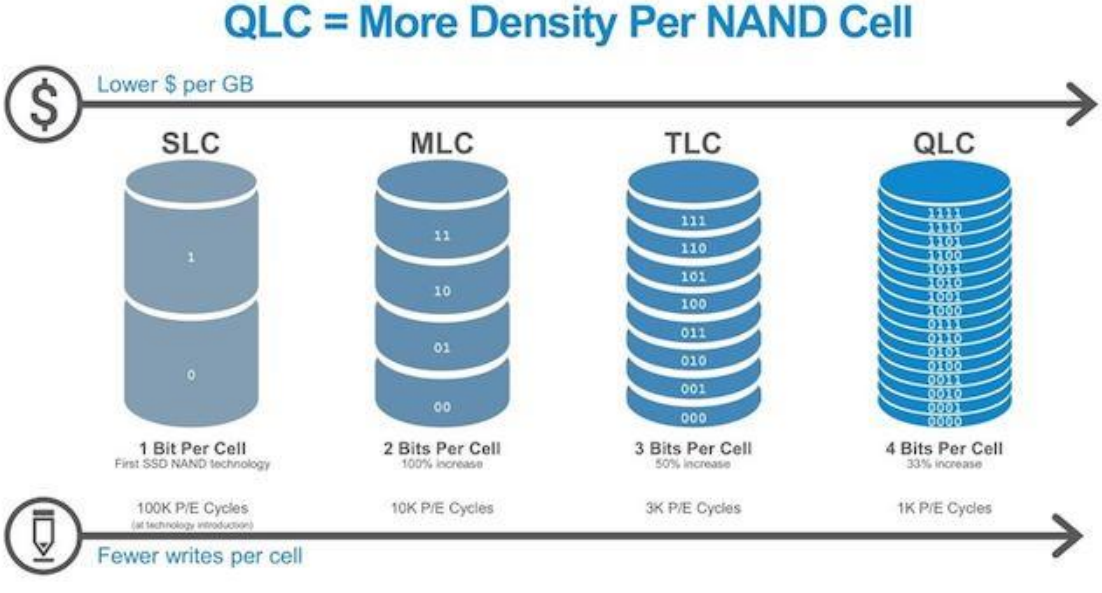

promising approaches to increasing NAND Flash density. However, there are not-insignificant challenges to doing this as well once you get to about 200 layers, especially around etching the channels that have to go through most or all of the layers.

In 2016, experts predicted 3D NAND would max at 300 layers. However, “[After SK Hynix’s 238 layers] I expect to see an increase in the number of layers over the next years at roughly the same speed,” said Roman Pletka, Senior Research Scientist, IBM Research. “However, increasing the number of layers is challenged from the technology point of view due to the high aspect ratio etching process, but also by CapEx because the time to manufacture a chip increases with the layer count. This is why we will see new scaling directions by making thinner layers, lateral scaling such as denser placement of the vertical holes, and the use of more efficient layouts such as shared bitlines and logical scaling (e.g., using split-gate architectures or storing more bits per cell). With these technologies, it is expected that the storage density of NAND flash keeps growing at the similar rate at least for the next 5 to 10 years.”

Others agree. “There’s no physical limit when people say we can’t go past this number of layers,” said Jim Handy, General Director at Objective Analysis. “In the world of semiconductors, there are always people saying we can’t do this. We can’t do lithography down below 20 nanometers. Now, they’re looking at 1 nanometer. Samsung talked about 1,000 layers. It could be that in 20 years we’ll be laughing that we once thought that was a lot.”

G2M Research Multi-Vendor Webinar Series

Our webinar schedule is below. We are offering a Cybersecurity series and an Enterprise Storage & Technology multivendor series.

"The Need for Speed: NVMe™, NVMe-oF™, and Data Processing Accelerators" webinar featured Tony Afshary, Vice President, Products and Marketing at Pliops; Rob Davis, Vice President of Storage Technology at NVIDIA; and Peter Onufryk, Intel fellow for NVMexpress. Companies are focused on storage/networking/processing acceleration. Higher-level networking protocols and custom protocols for specific workloads require "offloads" to lower CPU utilization and increase application performance. Advanced storage capabilities such as those offered by NVMe and NVMe-oF can also tax CPUs, reducing cycles available for workloads. And then there is security, data resilience, and other very real needs that take CPU cycles away from workloads. This webinar explored where non-hyperscalers go to accelerate their workload in the same way hyperscalers do. The webinar video is available to view and a copy of the slidedeck is available here.

Interested in Sponsoring a webinar? Contact **G2M** for a prospectus. We can create custom webinar, custom webinar series, and add or modify topics to specifically appeal to your target audience. View our webinars and access slide deck presentations on our website.

Enterprise Storage & Technology

Supercharging Oracle MySQL Performance, Scalability, & Efficiency with Pliops	March 21
Storage Architectures for Artificial Intelligence & Machine Learning	April 4
Software-Defined Flash Memory Architectures	May 9
Storage & Compute Architectures for Healthcare & Imaging Applications	June 27

NVMe & NVMe-oF – Past, Present, & Future	July 11
GPUs, SSDs, & Shared Memory: Accelerating Computing?	August 22
Securing Data – How Storage & Cybersecurity Technologies Can Work Together	Sept 26
The Open Compute Platform (OCP) Movement – Providing Compute-At-Scale Value to On-Premises Deployments	October 24
Storage Architectures for HPC Clusters	November 21
2024 Trends – Cloud, On-Premises, & Hybrid Compute/Storage	December 12

Cybersecurity

The Increasing Complexity of Cybersecurity Regulatory & Compliance for the Financial Services Industry	May 25
xDR- The Promise versus the Reality	August 3
10 Features of an Effective Attack Surface Management Tool	September 7
How Secure is the Cloud for Your Workloads?	October 12
Do You Need a SIEM? Use Cases Where a SIEM Makes Sense.	November 9

Upcoming Conferences

March 14-16	Gulf Information Security Expo, Dubai, UAE
March 20-22	Gartner Data & Analytics Summit, Grapevine, TX
March 20-23	GTC CPU Technology Conference, San Jose, CA
March 28-29	Gartner Security & Risk Management, Sydney, Australia
March 28-31	ISC West, Las Vegas
April 5-7	IST Information Security Expo, Tokyo, Japan
April 15-19	NABShow, Las Vegas
April 17-21	HIMMS Global Health Conference, Chicago, IL
April 17-21	Privacy Symposium, Venice, Italy
April 19-20	CyberSec Europe, Brussels, Belgium
April 24-27	RSA Conference, San Francisco
May 1-3	IAHSS AC&E, Nashville, TN
May 2-4	ACT Expo, Anaheim, CA
May 9-12	Black Hat Asia 2023, Singapore
May 15-17	Forth Roadmap Conference, Portland, OR
May 16-17	SIA GovSummit, Washington DC
May 17-18	Expo Summit Global, Santa Clara, CA
May 21-25	ISC, Frankfurt, Germany
May 22-25	Dell World, Las Vegas
May 22-25	Government Fleet Expo, Dallas, TX

June 2-6	School Transportation Network Expo East, Indianapolis, IN
June 4-8	Cisco Live, Las Vegas
June 5-7	Gartner Security & Risk Managemnt, National Harbor, MD
June 7-9	Synnex Red, White and You, Greenville, SC
June 11-14	36th Electric Vehicle Symposium & Expo, Sacramento, CA
June 11-16	2023 VLSI Symposium, Kyoto, Japan
June 14-16	Interop Tokyo, Chiba, Japan
June 20-22	HPE Discover, Las Vegas
June 20-22	Info Security Europe, London
July 14-19	School Transportation Network Expo, Reno, NV
August 5-10	Black Hat USA, Las Vegas
August 8-10	Flash Memory Summit, Santa Clara, CA
August 28-31	VMWare Explore, San Francisco, CA
August 30-Sept 1	Security Expo, Sydney, Australia
September 11-13	Gartner Security & Risk Management, London
September 11-13	Global Security Exchange, Dallas, TX
September 18-20	Crowdstrike fal.con, Las Vegas
September 18-21	SDC 2023, Fremont, CA
October 2-4	DattoCon, Miami, FL
October 3-4	CyberTech Europe, Rome
October 16-19	Gartner IT Symposium/Xpo, Orlando, FL
November 15-16	Microsoft Ignite, TBD
Nov 27- Dec 1	AWS re:Invent, Las Vegas