DRAM and Advanced Packaging Master Class

PREPARED REMARKS | June 25, 2026

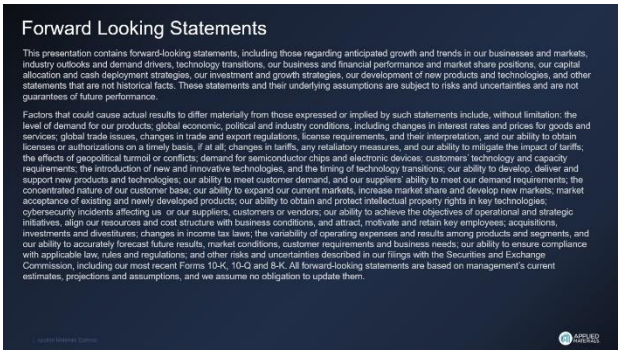

- Hello and welcome back to the Applied Materials Master Class Series.

- Several years ago, we anticipated that the AI wave would drive the semiconductor industry to $1 trillion dollars in annual sales by 2030.
- And through that time, we modeled the wafer fab equipment market composition to be around 1/3 leading-edge foundry-logic, 1/3 ICAPS, and 1/3 memory – with memory evenly divided between DRAM and NAND.

- Today, we see AI driving the semiconductor industry to around $1 trillion dollars this year, and new, incremental applications like agentic, edge and physical AI growing the industry to much higher levels over the next several years.
- These AI waves are fueling demand for faster, more energy-efficient chips and systems, and this is creating a new WFE spending mix.

- We now expect leading-edge foundry-logic to outgrow ICAPS and drive well over 50% of foundry-logic in the years ahead.
- And in memory, we expect DRAM WFE spending to be well over two times NAND spending.
- Applied is well positioned for this new mix with the highest process equipment market share in leading-edge foundry-logic – which we covered in our April Master Class – as well as both DRAM and advanced packaging, which we're covering today.

- In a moment, Kevin Moraes will summarize our strategy and explain why DRAM and advanced packaging are growing with AI.
- Next, Sony Varghese will share the roadmap for both standard DRAMs and high-bandwidth memories.
- Then, Jinho An will discuss how advanced packaging is enabling faster and more energy-efficient AI chips and systems.
- Next, Lior Engel will explain how we are bringing eBeam process control to advanced packaging.
- Finally, Kevin will give examples of how inflection-focused innovation is growing our business, and then the five of us will address your questions.
- Now Kevin, it's over to you!

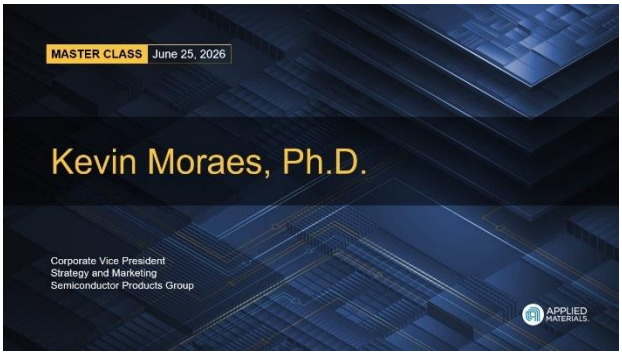

- Thank you, Mike.

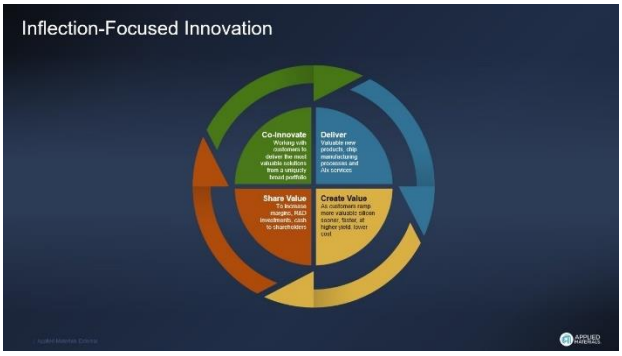

- Applied's strategy is inflection-focused innovation.
- We collaborate with our customers to predict roadmap inflections and develop solutions to the highest-value technology challenges.

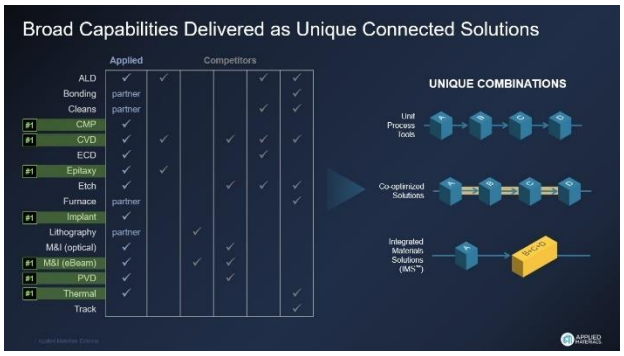

- We do this using our broad portfolio of unit process systems, co-optimized solutions and integrated solutions.
- In our previous Master Class on logic, we showed you how our strategy has given Applied the #1 materials engineering share position in Gate-All-Around – in both transistors and wiring – with combined share of around 50%.

- I'll next share three examples of how we used our strategy to become #1 in DRAM and advanced packaging as well.
- Back in 2013, Applied's DRAM process equipment share was below 15%, and we've since gained around 10 points of share and become #1.

- In our first Master Class 5 years ago, we explained two DRAM inflections.

- The first was patterning where we co-optimized our Draco™ CVD hard mask patterning film, Sym3™ etch and PROVision™ eBeam metrology to help customers scale DRAM capacitors.
- Today, these products are used by all our customers and generate over half a billion dollars in annual sales.

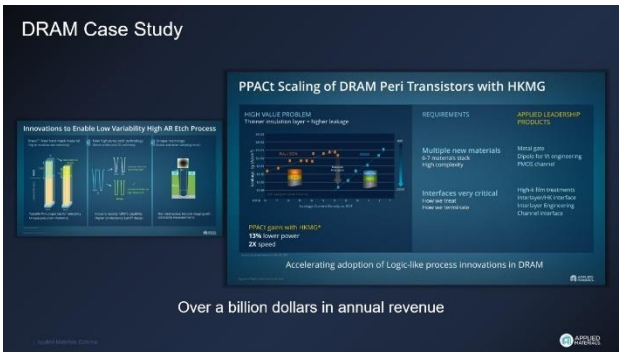

- Another inflection was transitioning the DRAM periphery logic transistors from polysilicon to high-k metal gate and incorporating more advanced wiring technologies to increase DRAM performance and energy efficiency.
- This change has increased our DRAM revenue by more than a billion dollars per year.

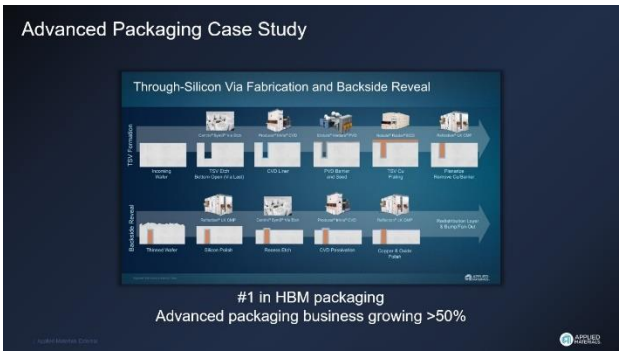

- A third inflection was high-bandwidth memory (HBM).
- We formed a packaging group that combined our materials deposition and removal technologies to enable through-silicon vias and microbumps.
- High-bandwidth memory took off, and we've become #1 in HBM packaging equipment sales.
- In fact, HBM is growing rapidly this year and helping expand our overall advanced packaging business by over 50%.

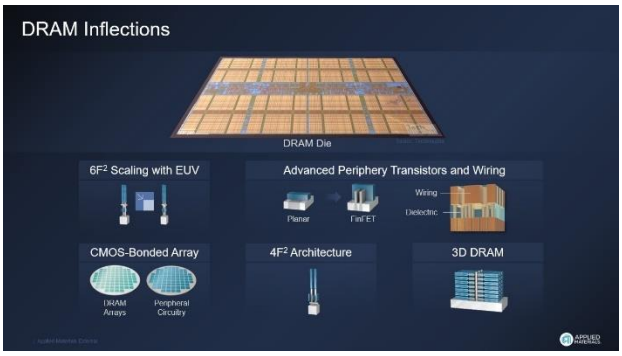

- In today's Master Class you'll learn about five emerging DRAM inflections and two advanced packaging inflections.
- In $6F^2$ DRAM, we see customers increasing the use of EUV patterning to further scale DRAM memory cells and periphery logic.
- Next, we see DRAM customers following the foundry-logic roadmap, adopting more advanced transistors and wiring in the periphery logic to improve performance and power.
- In addition, we see a new CMOS-bonded array architecture in which memory and periphery logic are built on separate wafers that are later bonded together.
- Next, we see a new $4F^2$ DRAM architecture based on vertical cell transistors.
- Finally, we see 3D DRAM, built using the same deposition and etch technologies Applied leads in today.

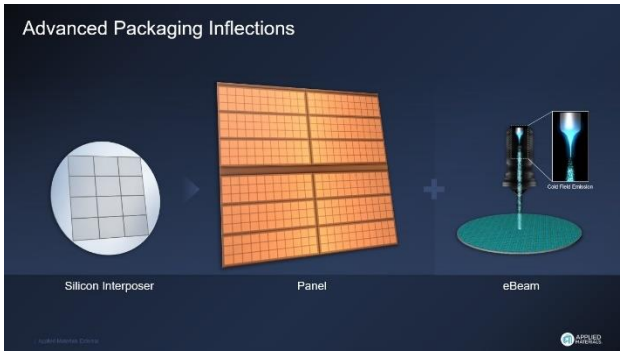

- In advanced packaging, we see the industry moving from silicon wafer interposer substrates to panels which enable larger-body AI accelerators.
- We also see the eBeam technologies used in leading-edge foundry-logic and DRAM being adopted in advanced packaging.

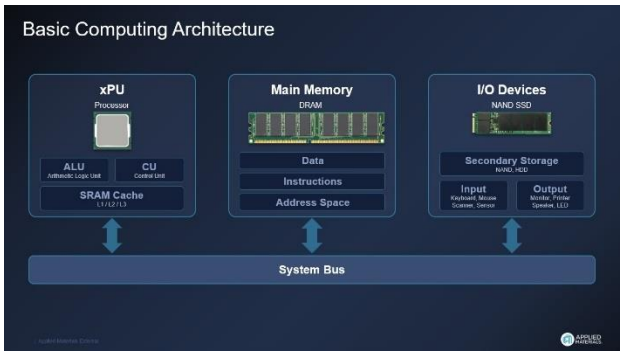

- Before Sony, Jinho and Lior discuss these inflections, I'll provide a non-technical overview of some of the key ingredients inside the computer to help explain the unique and critical role DRAM and advanced packaging play in the AI era.
- Regardless of workload, computers need processors to execute software instructions and manipulate data.
- CPUs, GPUs and TPUs can process extraordinary amounts of data each second.
- Keeping them fueled with data to avoid idle states is imperative to performance and performance per watt.
- Since processors have limited space for memory, we prefer to keep data that is not actively being processed nearby rather than on-chip.

- The perfect memory would have low latency, high bandwidth, random accessibility, unlimited endurance, non-volatility and low cost.
- But the perfect memory doesn't exist, and over the decades, a number of memories with different characteristics have been invented and refined.

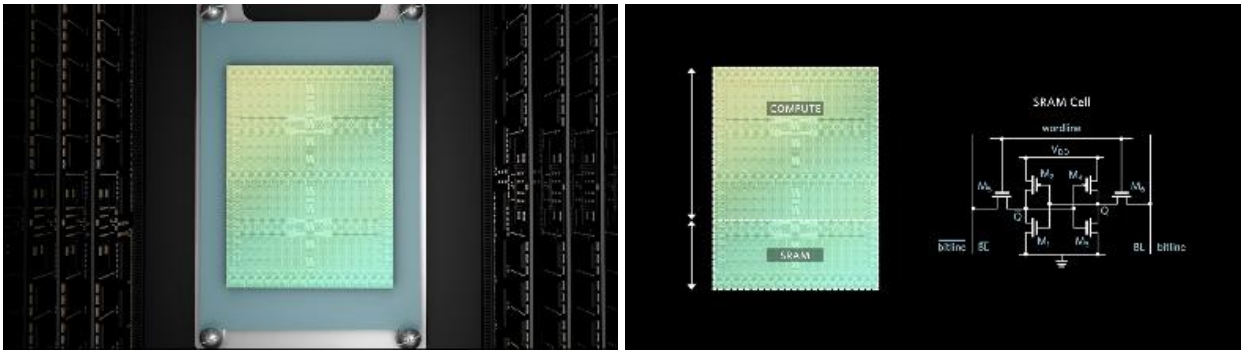

(animation)

- Let's take a closer look at three key memories.
- SRAM is built using logic process technology and can be integrated onto the processors, enabling cache memories with zero-wait-state performance.
- SRAM is extremely fast.
- The downside is SRAM typically requires 6 transistors for every bit of data stored which makes it particularly expensive.
- And after decades of optimization, designers are finding it increasingly difficult to scale SRAM more than marginally with each new node.
- SRAM caches can easily take up a third or more of the processor's die area, and using advanced packaging to relocate the caches without sacrificing performance is a major focus of the industry today.

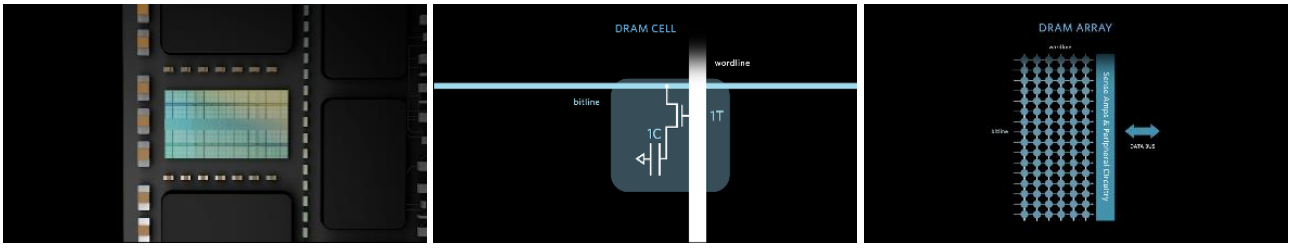

(animation)

- Next comes DRAM.
- DRAM process technology is mostly incompatible with logic technology and resides off-chip.
- It shares the random-access capability of SRAMs but needs only one transistor per bit of storage, making it much more space-efficient and cost effective.
- Each DRAM cell has a transistor that controls the binary state of the bit -- and a capacitor that holds a charge level that represents zeros and ones.

- Each cell is connected to a wordline and a bitline, which together enable random read and write access.
- Zooming out, each DRAM chip has gigabytes of memory cells along with periphery logic regions that route data to and from the memory cells – and to the rest of the computer system.

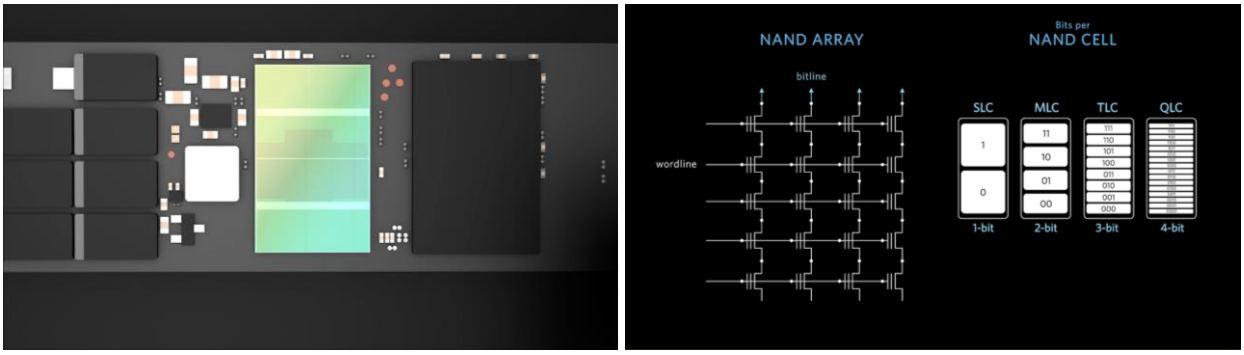

(animation)

- Finally, NAND.
- Each NAND cell includes a transistor that traps electrons, with the resulting threshold voltage representing zeros and ones.
- NAND memory cells do not have the space overhead of unique bitline and wordline connections.
- Instead, they are arranged in long, narrow strings that offer higher density and lower cost versus DRAM.
- A special attribute of NAND is that the transistors can be programmed as multi-level cells whereby 4, 8 or 16 distinct voltage levels are used to logically represent 2, 3 or 4 bits of data in each cell.
- The transition from single-level cell NAND to these multilevel cell variants has dramatically increased NAND storage density relative to DRAM.
- In fact, it can take nearly 200 HBM DRAM wafers to hold the data stored on a single QLC NAND wafer.

- But the density and cost benefits of NAND come at the expense of performance and endurance.
- Strings of NAND cells are organized into larger pages and blocks.

- Individual bits are not randomly accessible; instead, an entire page of data — typically several kilobytes — must be read to access any particular bit.
- Write operations occur at the page level.
- And erase operations can only be done at the block level, wherein millions of NAND cells are erased at the same time.
- NAND chips constantly move fresh data to empty cells so blocks of stale data can be erased.
- Erasing NAND requires more than 20 volts, and the operations degrade NAND reliability, leading to endurance limits.
- All of this helps explain why NAND is not a perfect substitute for DRAM — and why designers instead use NAND more like a fast hard disk drive for data that seldom changes.

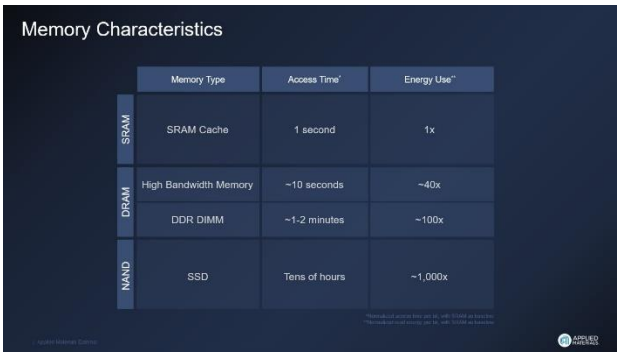

- This table summarizes memory performance as measured in access time and energy per bit transfer.
- If you normalize accessing a bit of data from SRAM to one second, then accessing the same bit from HBM takes around 10 seconds, from a DRAM module 1 to 2 minutes, and from a NAND SSD tens of hours.
- Energy consumption per bit transfer also multiplies substantially as data moves farther away from the processor, from on-chip SRAM to high-bandwidth memory, DRAM module and NAND SSD.

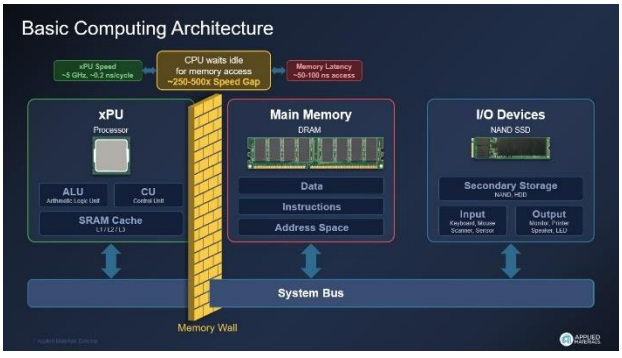

- Returning to the overall system architecture, a key challenge today is what we call the “memory wall.”

- It describes the performance gap between increasingly fast processors and memory which struggles to keep the processors busy.
- In AI inferencing, memory access and data movement can account for the vast majority of overall system activity and energy use.

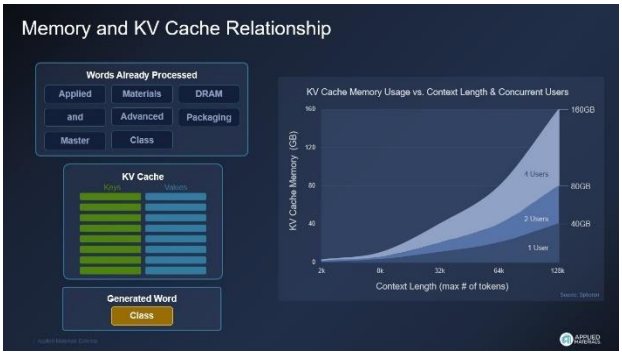

- Next I'll relate what we've discussed about memory to AI and what's called the KV cache.
- In AI workloads, items to be processed such as language are divided into small segments called tokens that can be manipulated by AI models.
- Key-value pairs are mathematical representations of the tokens that are stored as the model works to understand context.
- In agentic AI, context windows grow longer, and sessions persist across increasing numbers of steps, causing KV cache data to expand.
- DRAM is extremely fungible and preferred for performance-critical KV cache data that is frequently updated.
- NAND provides a cost-effective repository for overflow KV cache data that seldom changes.
- Bringing together everything we've discussed, a key insight is the closer we put memory of all kinds to the processors, the faster the system and the lower the power.
- This explains the opportunity to scale the entire system using the advanced packaging techniques we'll describe today.
- I hope my presentation helps you understand why we believe leading-edge foundry-logic, DRAM and advanced packaging will be attractive areas of the WFE market for as long as AI continues to proliferate and grow.
- It should also help you understand why we now expect the DRAM WFE market to be over twice the size of the NAND market for the foreseeable future.
- Now I'd like to turn the meeting over to Sony Varghese who will discuss the DRAM roadmap. Sony?

- Thank you Kevin.
- I've spent over 20 years in the DRAM industry, working at one of our customers for over a decade.
- There has never been a more exciting time for DRAM, both in the value DRAM brings to AI – and in the innovation that's happening, from the memory cells, to the periphery transistors and wiring, to the advanced packaging.
- I'll discuss five DRAM inflections that combine to improve DRAM density, performance and power consumption.

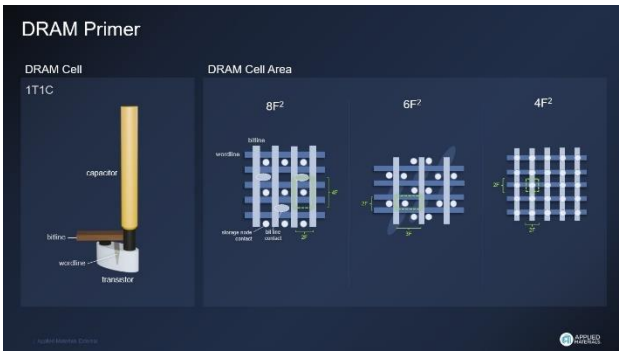

- Let's start by defining the key DRAM structures.
- As Kevin mentioned, every DRAM cell needs a transistor and a capacitor.
- To hold a one in storage, the capacitor's charge needs to be refreshed around 15 times every second.
- You've heard about $6F^2$, which was preceded by $8F^2$ and will be followed by $4F^2$.
- F^2 refers to the square of the smallest wire width in the DRAM cell, and the coefficient refers to the multiple of that area needed to construct a complete memory cell.

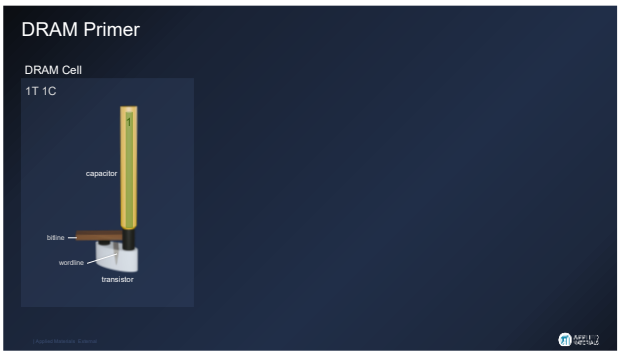

- Within the memory cell, the capacitor acts like a tiny storage container, and the transistor acts as a switch.
- The switch turns on to read or write data, and off to hold it in place.

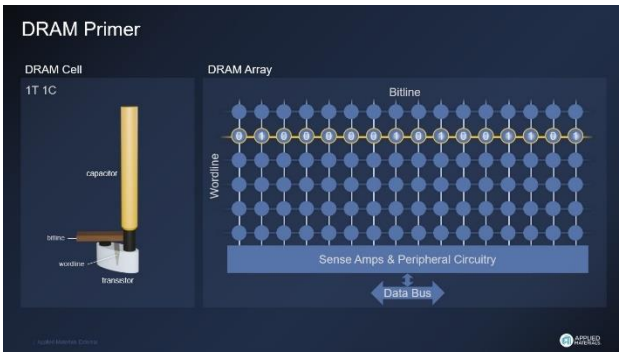

- Cells are organized in rows called wordlines and columns called bitlines.
- When a wordline is activated, all the cells in that row are turned on.
- Bitlines connect the cells to the surrounding control circuitry which manages all key operations: reading, writing, and refreshing data so it is not lost.
- It also measures the electrical charges to determine whether each cell holds a zero or a one.
- As the memory cells shrink, they store less charge, making it harder to reliably distinguish zeroes from ones.
- In NAND, sensing reliability challenges drove the industry from 2D shrinking to 3D stacking of larger, more reliable storage elements.
- In DRAM, we still have opportunities to drive 2D scaling even further before we need 3D stacking.

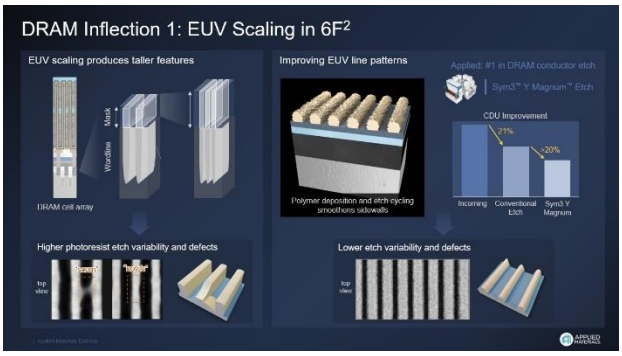

- We see several more inflections in $6F^2$ DRAM in the coming years.
- The first DRAM inflection is greater use of EUV to increase DRAM density and capacity by scaling DRAM capacitor arrays, wordlines, bitlines and shallow trench isolation.
- The EUV inflection is positive for Applied because our Sym3™ Y Magnum™ system is the most widely adopted etch technology for EUV patterning in DRAM.
- The system acts on EUV photoresists, cyclically etching and redepositing the material in the same chamber to straighten lines that are uneven due to EUV stochastic errors.
- By making EUV line patterns smoother before they are etched into the wafer, we help increase yields and decrease line resistance to improve chip performance and power.

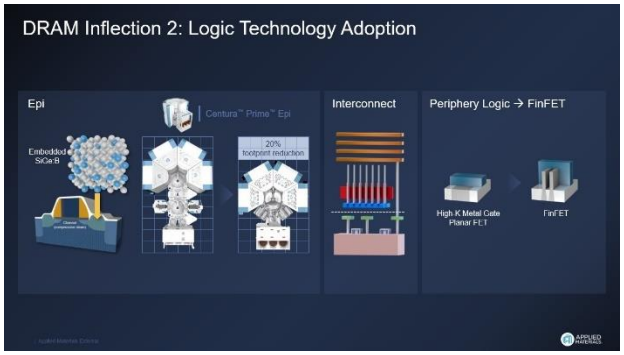

- The second DRAM inflection is that the transistor and wiring technologies pioneered in leading-edge foundry-logic are now being adopted in DRAM.
- I'll share three examples.
- First, our customers are adding a new epitaxy step.
- Applied today introduced a Centura™ Prime™ Epitaxy system that has been specially designed for DRAM fabs.
- We selectively grow boron-doped, embedded silicon germanium in the source/drain regions of DRAM periphery transistors.

- The resulting channel strain improves DRAM performance and power consumption.
- A unique feature of the new system is a 20% smaller footprint which helps DRAM makers generate more output per square meter of cleanroom.
- Second, as continued DRAM scaling with EUV packs more DRAM bits and periphery transistors on each chip, designers are adding more copper interconnect layers.
- Just as in foundry-logic, these additional layers increase the market for Applied's leadership products in wiring, including PVD and CMP.
- Third, DRAM makers are preparing to transition from high-k metal gate periphery logic transistors to FinFETs.
- Just as in foundry-logic, FinFET transistors offer better control of the silicon channel which increases performance and reduces power consumption.
- FinFETs also increase DRAM scaling in two ways.
- First, FinFET transistors are more space efficient than high-k metal gate transistors.
- Second, FinFETs have lower transistor to transistor variability.
- I described how the periphery logic senses the capacitor charges that represent zeros and ones, and how the sensing margins are becoming narrower with scaling.
- FinFETs allow us to continue to scale density AND maintain reliability.

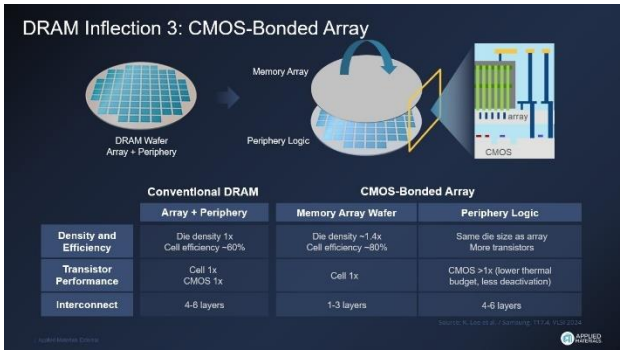

- The third DRAM inflection is called CMOS-bonded array.
- DRAM makers can build memory arrays on one wafer, periphery logic on a second wafer, and bond the two together.
- The approach lends itself to $6F^2$, $4F^2$ and 3D DRAM as well.
- There are several benefits to decoupling logic from memory.
- One is density: placing memory and logic on different wafers allows each to be scaled more aggressively.

- A second benefit is transistor performance.
- The high temperatures needed for DRAM capacitor fabrication deteriorate FinFET transistor performance which is maximized as the array steps move to a different wafer.
- A third benefit is interconnect performance.
- In wafer-to-wafer bonding, the connections between memory cells and periphery logic are shorter and more vertical.
- Designers can eliminate many of the long, horizontal connections that reduce speeds and increase capacitance, power and heat.
- From a business perspective, the wafer bonding also adds new wiring steps that Applied can serve.

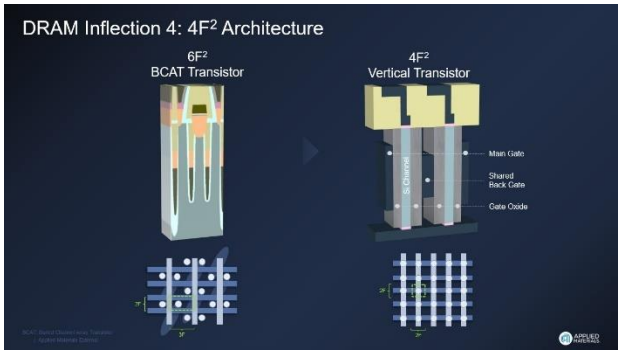

- A fourth DRAM inflection is the space-efficient $4F^2$ architecture in which vertical transistors are deployed in the memory cells to further scale both the transistors and the capacitors.
- I'll describe significant $4F^2$ materials engineering challenges and the products Applied is delivering to solve them.

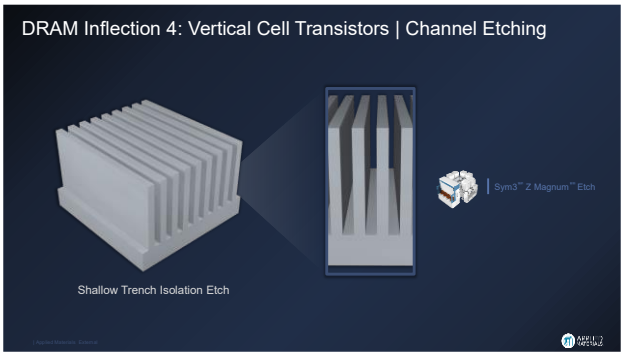

- The first is etching the vertical cell transistor's narrow, high-aspect-ratio silicon channels.
- The shapes needs to be exactly as intended, and the billions of channels across each chip on the wafer need to be uniform.

- Applied's Sym3™ Z conductor etch system has been optimized for this application.

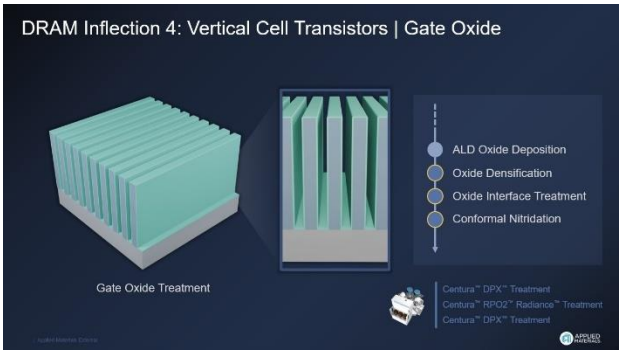

- A second challenge is managing the materials interface between the channel and gate oxide to reduce electron leakage, improve data retention, support capacitor refresh operations, and minimize standby power.
- Applied delivers three helpful steps.
- First, Applied's Centura™ DPX™ solution uses plasma to densify the gate oxide to improve its quality.
- Next, our Centura™ RPO2 Radiance system engineers the gate oxide interface to reduce defects that trap electrons and holes.
- Finally, Centura™ DPX™ builds a thin nitride film barrier that promotes adhesion of the subsequent transistor metal gate layers and prevents metal impurities from diffusing into the gate oxide.

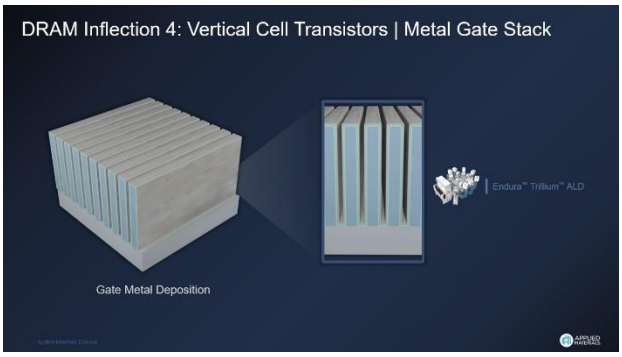

- A third challenge is creating the transistor's multi-layer metal gate stack.
- The Trillium™ ALD system launched during our Logic Master Class in April is also being used in vertical cell DRAM to enable uniform metal gate stack deposition.

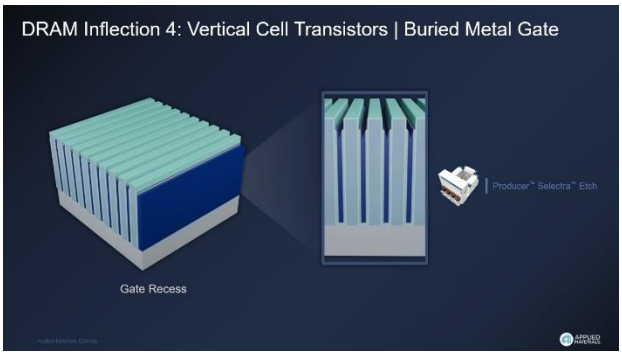

- A fourth challenge is reducing the area occupied by the metal gate stack to maintain the tight $4F^2$ design parameters.
- Etching is used to recess or bury the metal gate within a pre-formed trench in the silicon substrate.
- Applied's selective etch technology tightly controls the recess shapes, depth and uniformity.

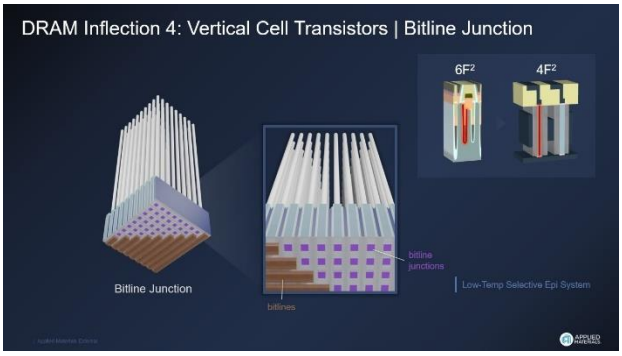

- A fifth and particularly important challenge is forming the bitline junction to the cell transistor.
- A region of the silicon channel is doped with phosphorous ions, creating an electrical connection between the two.
- This new step is the first to require low-temperature selective epitaxy in DRAM.
- In $6F^2$ DRAM, the cell transistors are U-shaped, the wordline is placed within the U, and the bitline connections are made at the top of the structure where both ends of the transistor channel are accessible.
- But the $4F^2$ transistors are I-shaped, making it difficult to dope the lower end of the silicon channel.
- The junctions are deep and narrow, and reliability depends on keeping the dopants contained to target regions of the transistor – and applied uniformly across all the transistors on the wafer.
- In $4F^2$ DRAM, traditional high-temperature anneal processes are not viable because they damage the delicate capacitor structures.

- To solve this, Applied has developed a low-temperature selective epitaxy that grows doped crystalline silicon only on the target region of the transistor channel.

- From a roadmap perspective, we see multiple nodes of $6F^2$ innovation before the arrival of $4F^2$ and 3D DRAM.
- However, R&D pathfinding in 3D DRAM is happening today.
- The work includes identifying the ideal materials and structures, and learning how to deposit, modify and remove them to create reliable, high-performance DRAMs that are scalable to much greater heights over time.
- A key insight from the earlier portion of today's class is that 3D DRAM will be very different from 3D NAND, demanding a unique set of materials and materials engineering capabilities that are already leadership areas for Applied today.
- Earlier we noted that DRAM is a random-access, read-write memory that is heavily used in computing to deliver high performance, low power and high endurance.
- In contrast, NAND is better suited to reading and writing blocks of data that rarely change.

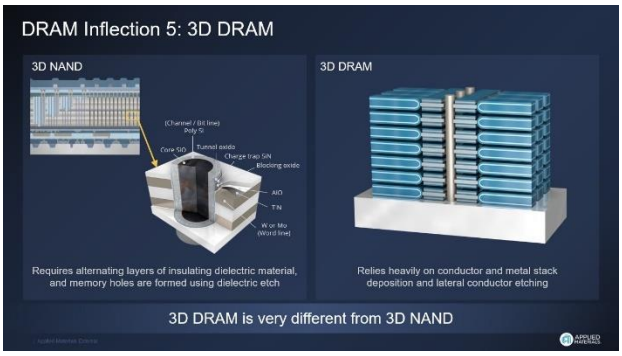

- 3D NAND is built by depositing alternating layers of insulating dielectric material, and memory holes are formed using dielectric etch.
- In contrast, 3D DRAM stacks will rely heavily on conductor and metal stack deposition and lateral conductor etching.

- Applied's leadership in epitaxy, conductor etch, eBeam metrology and inspection, and other technologies give us a head start in collaborating with customers to develop 3D DRAM prototypes.

- In fact, 3D DRAM is one of the technology inflections targeted for customer co-innovation at the EPIC Center in Silicon Valley.
- Because we are doing pathfinding work today, we already have line of sight to many of the key process steps and high-value problems we can address with our products.
- I'll discuss some of the major 3D DRAM materials engineering challenges and the solutions Applied is providing.

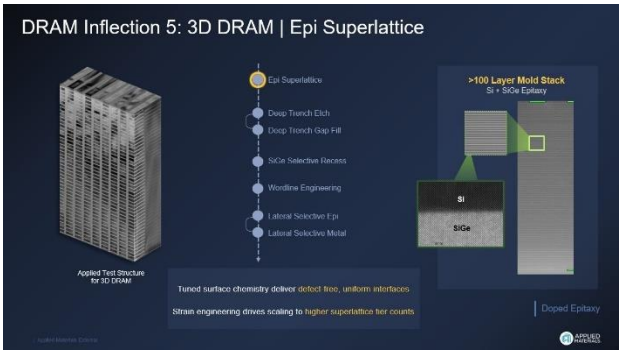

- 3D DRAM requires building an ultra-tall stack of more than 100 alternating layers of silicon and silicon germanium.
- This superlattice is formed using advanced epitaxy.
- Maintaining tier-to-tier uniformity across hundreds of horizontal layers is essential, since variations impact electrical performance.
- A major materials engineering challenge is managing interfacial stress within the silicon germanium lattice which has a higher lattice constant than silicon.
- Without control, lattice mismatches can generate defects that increase DRAM cell transistor leakage and variability.

- Applied is using its epitaxy leadership to tune surface chemistries and process conditions, to help customers achieve high film quality and uniformity.
- Another challenge is structural integrity and uniformity of tall and complex 3D DRAM materials stacks.
- The technology we use to treat silicon nanosheets in gate-all-around transistors can help strengthen the silicon channels in 3D DRAMs as well.
- The 3D DRAM layers will be highly sensitive to impurities and particles.
- Applied's integrated materials systems can be used to maintain an ultra-high-vacuum, clean deposition environment.

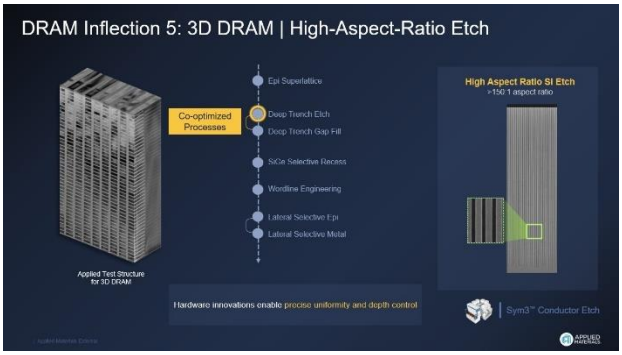

- Once the 3D DRAM stack is formed, it must be etched at an ultra-high-aspect-ratio of more than 200 to 1.
- The etch pattern is defined from above.
- A CVD hard mask is deposited on top of the stack, patterned and opened with a hard mask etch step to create a stencil for a second etch step that will transfer the pattern throughout the stack below.
- Next, Applied's Sym3™ conductor etch system is well suited to etching the tall materials stack, providing excellent uniformity and depth control through hardware innovations like high-voltage bias pulsing and ion directionality.

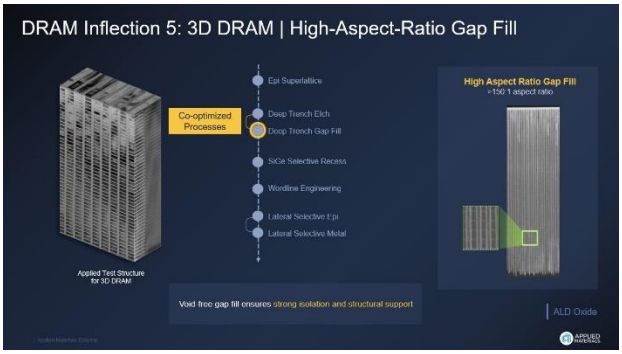

- After the deep etch, a high-quality dielectric gap fill is needed to provide electrical isolation and structural integrity.
- ALD oxide is a good candidate for void-free gap fill.

- Next, precise lateral recess etching is required to selectively remove sacrificial silicon germanium layers and define where subsequent structures like the metal wordline will be formed.
- Even nanometer-scale variances in lateral recess etch depth can degrade transistor performance and cause leakage.
- A highly selective recess etch is needed to remove all the silicon germanium while leaving the adjacent materials and structures undisturbed, and Applied has industry-leading selective etch technology.

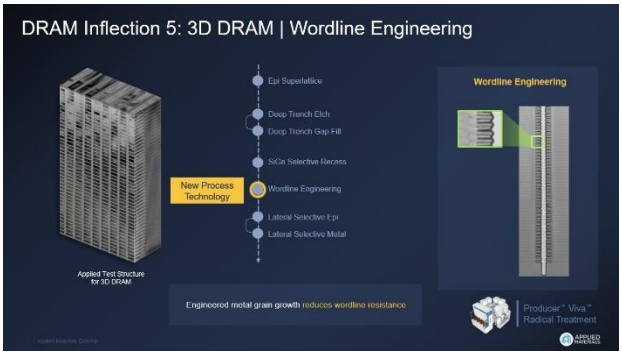

- The recess step sets the geometry of the wordline that will be deposited within the tall 3D DRAM stack.
- A metal such as moly will be used for the wordline, and as stacks grow to 200 layers or more, managing wordline sheet resistance will be critical for performance, power consumption and switching uniformity across layers.
- One way to reduce resistance is by growing large, well-formed grains of moly that reduce electron scattering.
- Applied's Viva™ High Power solution uses a radical and thermal treatment to precisely control grain growth.

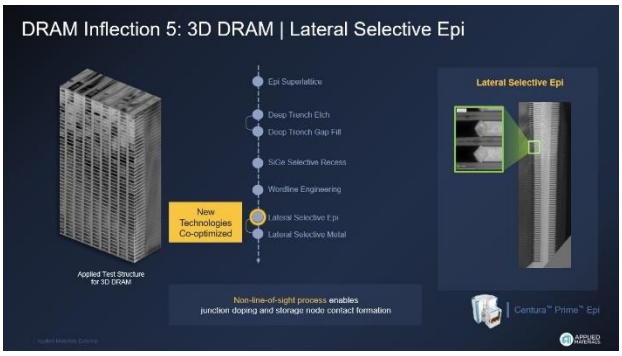

- Next, selective epitaxy is used to form the lateral transistor junctions and storage-node contacts.
- This epitaxy step is especially challenging because it is non-line-of-sight and must be uniform across hundreds of layers.
- Applied's Centura™ Prime™ Epi system is tuned for highly selective deposition and leaves no unwanted material that could create shorts in buried structures that cannot be accessed or repaired.

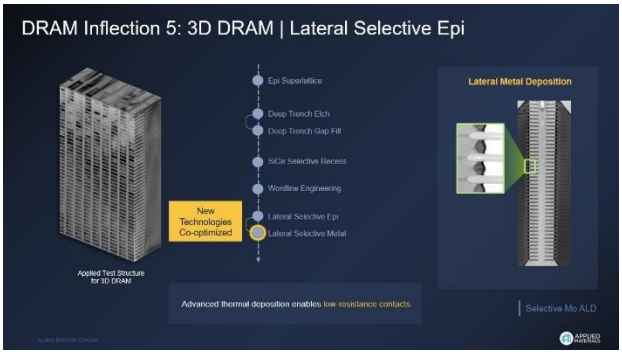

- Finally, selective ALD deposition of a metal such as molybdenum silicide forms low-resistance contacts across all the layers.
- Lateral selective deposition is challenging, and plasma energy cannot be used because ions travel directionally.
- So thermal processes must be used instead, and the selective ALD moly technology we developed for the foundry-logic market can be used in this 3D DRAM application.

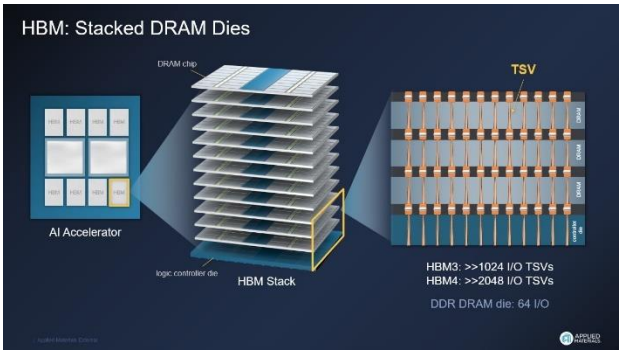

- I'll conclude my section with an overview of high-bandwidth memory.
- HBM DRAMs are like standard DRAMs in most areas of wafer processing technology.
- The most significant design difference is the TSV channel holes which are large, consuming as much die area as 2,000 memory bits.
- As a result, HBM dies can be twice the size of standard DRAM dies, and the industry needs to produce 3 to 4 HBM DRAM wafers to deliver the same number of bits.
- Comparing HBM4 to HBM3, major differences include a much wider, 2048-bit interface, a doubling in channel count, and a tighter TSV pitch.
- At the base of every HBM stack sits a logic die that helps transfer data between the HBM chips and the processor's memory controller.
- In earlier HBM generations, the base die performed relatively simple control functions and was built on non-leading-edge logic.
- In HBM4, the base die has more transistors to deliver higher data bandwidth, improve signal integrity, and enable power management.
- The HBM4 base die is moving to more advanced logic nodes with FinFET transistors and advanced wiring which are better suited to running these more complex functions at lower power.
- As a result, the growth in HBM demand increases Applied's foundry-logic business as well.
- My colleague Jinho will explain how TSV DRAM wafers are later processed to create and connect the through-silicon vias.
- Thank you for listening, and now, Jinho, it's over to you.

- Thanks Sony.
- I worked at the world's largest memory company on technologies including HBM before joining Applied 5 years ago to develop new advanced packaging platforms.

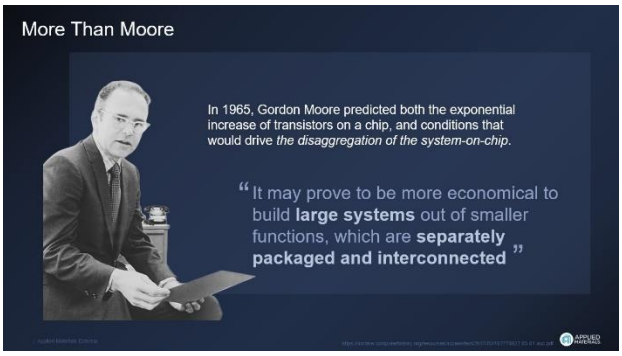

- At that time, we held our first advanced packaging master class and proposed that the future of chiplets that Gordon Moore predicted back in 1965 was going to be a key driver of the industry's progress in the AI Era.
- Today, no one doubts that packaging has become as strategic to the computing industry as on-chip scaling of transistors and wiring.

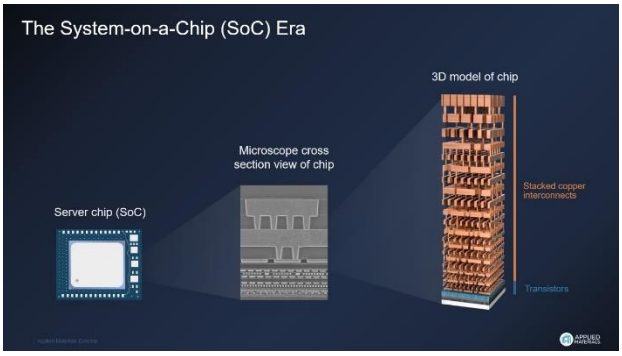

- Let's take a moment to revisit how chip design was done when Moore's Law was working well, in the "system on a chip" era.
- On the left is a server chip with a variety of functional blocks.
- With each new node, designers added more performance-enhancing transistors and on-chip SRAM memory.
- They also integrated more system functions that were previously handled off-chip on application-specific ICs.
- Looking at the chip in cross section, the tiny transistors were connected to successive layers of wiring, as we saw in our previous Master Class.
- The more transistors, the more layers of wiring.
- And the farther from the transistors, the thicker the wires and the fewer the connections.

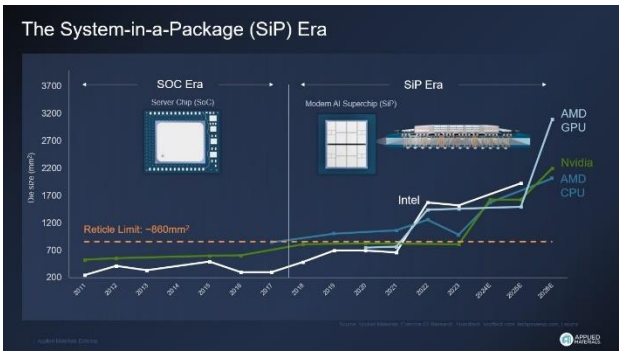

- Looking at a modern AI superchip, you can see that as the pace of Moore's Law slowed, designers adopted a "system-in-a-package" strategy.
- Each processor approaches the "reticle limit," or the roughly 860-square-millimeter size that an EUV litho scanner can print.
- And to keep up with growing AI performance and data needs, there are now multiple processors in the design.
- Some system functions have moved from on-chip to off-chip.
- Looking at the system-in-a-package in cross section, we see many layers of wiring, like in the system on a chip, but significantly more complex.
- I'll next aim to make all of these easy to recognize and understand.

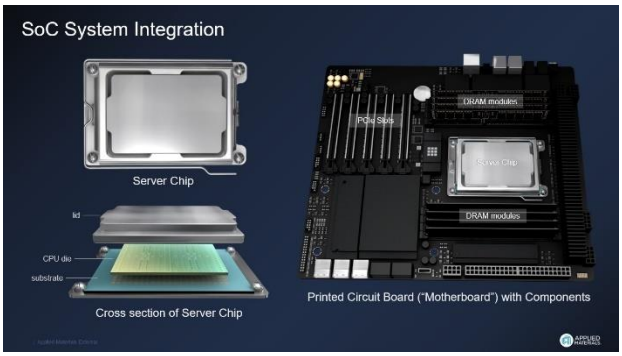

- To define the elements, let's begin with a look back at packaging and system integration during the system-on-a-chip era.
- This is a server chip and printed circuit board from 2016.
- Looking at a cross section, we see that the CPU die sat on a chip substrate.
- And the chip substrate was connected to an epoxy printed circuit board.
- This "motherboard" had wiring traces that connected the processor to DRAM modules which each had a number of similarly mounted DRAM chips.

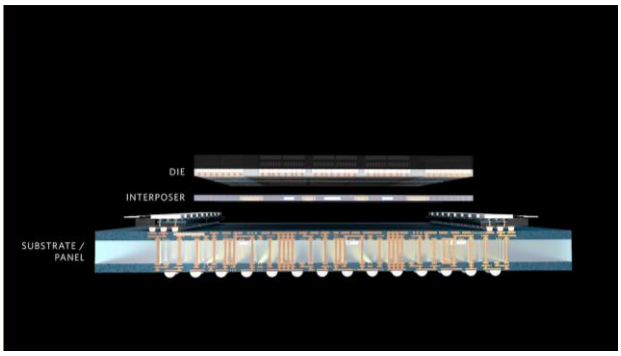

(animation)

- Even in today's most advanced packages, we still combine these three ingredients: dies, substrates and motherboards.
- But we add a fourth element: the interposer, which sits between the dies and the package substrate.
- The interposer is where much of today's advanced packaging innovation is focused.
- To understand the need for the interposer, recall that in on-chip wiring, the closer to the transistors, the thinner the wires and the more the connections.
- In the system in a package design, there are now far too many transistors and connections for any normal chip substrate to handle.
- The interposer has multiple layers of wiring called redistribution layers – or RDLs.
- It acts like a miniature circuit board within the advanced package, enabling thousands of short, high-density connections between the multiple dies above and the substrate below.
- Compared to die-to-substrate packaging, interposers dramatically reduce interconnect length which reduces resistance and power consumption while increasing bandwidth and performance.
- In short, the interposer marks the transition from packaging as a protector and connector to packaging as a critical enabler of energy efficient performance.

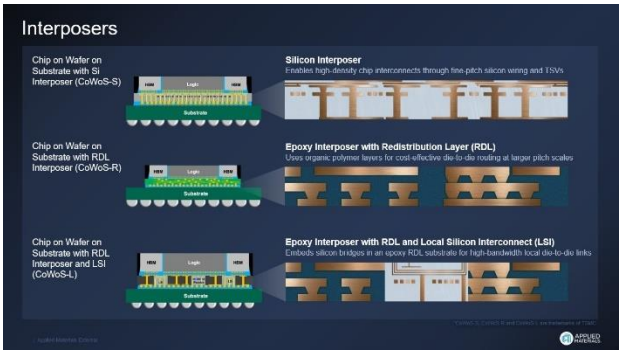

- The first interposers were made of silicon because the industry has decades of experience building layers of fine wiring on silicon wafers.

- Another approach is building layers of wiring into epoxy as an interposer material and embedding silicon bridges into the epoxy, right beneath the dies.
- The bridges connect the dies to one another and the substrate below.
- The key benefit is cost as the epoxy substrates and silicon bridges cost less than silicon wafers.

- This is a good time to define a few packaging terms.
- 2D refers to connecting packaged chips side by side on a circuit board, as in the older server example.
- 2.5D packaging is using interposers to connect multiple dies to one another on an intermediate substrate as we just discussed.
- 3D packaging is stacking dies on top of dies using through-silicon-vias and microbumps or hybrid bonds which I'll soon describe.
- High-bandwidth memory is a great example of 3D stacking.
- And 3.5D packaging is connecting 3D chip stacks on interposers or substrates.

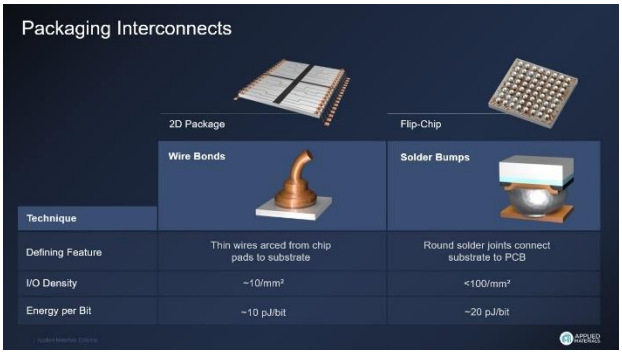

- Next, let's talk about packaging interconnects beginning with traditional 2D packaging.
- I/O and power circuits are often placed along the edges of chip designs.
- In simple chips, wires are bonded to these pads and the circuit board so power can be brought to the chip and data can be routed between the chip and the rest of the system.

- One of the original interconnect technologies is wire bonding which has around 10 connections per square millimeter.
- In the 1990s, as transistor counts increased, the industry brought redistribution layers to chip designs, thereby re-arranging the chip's I/O connections as a rectangular grid pattern that covers the entire bottom of the chip's surface with evenly spaced connection points.
- This change supports the use of flip-chip packaging which offers hundreds of I/O connections to the chip's substrate.
- Round solder bumps connect the substrate to the motherboard, and there are well under 100 bumps per square millimeter.

- 2.5D and 3D packaging use microbumps with I/O densities below 1,000 per square millimeter.
- 3D packaging typically relies on through-silicon vias – or TSVs – the outlines of which are defined during wafer fab production.
- After wafer processing and TSV fabrication, small copper pillars and microbumps can be used to connect the chips using a technique called thermos-compression bonding – or TCB – with I/O density approaching 10,000 per mm^2.
- Hybrid bonding is the ultimate interconnect technology.
- It completely eliminates bumps and pillars.
- Instead, we align the on-chip copper wiring of two chips and use heat to fuse the copper of both chips.
- I/O densities approaching 1 million per mm^2 are possible.
- An important observation is that the smaller the packaging interconnect, the lower the energy needed to transfer data, which is measured in picojoules per bit transfer.
- In advanced packaging, bringing chips closer together and decreasing the interconnect pitch increases performance and reduces power consumption.
- Finally, co-packaged optics puts tiny fiber-optic connections right next to the chip, so data moves as light instead of electricity, further increasing I/O bandwidth and energy efficiency

- I'll talk more about this later.

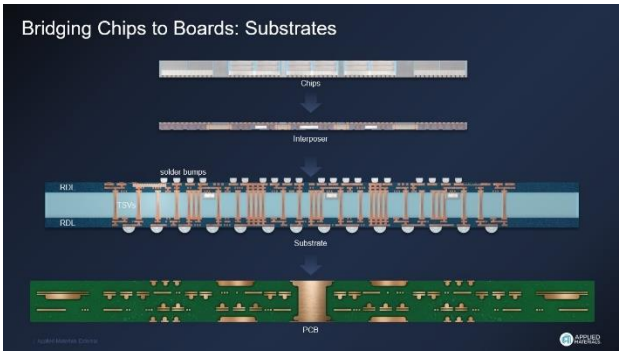

- Next I'll take a moment to describe substrates, which hold a single chip or an interposer full of chips.
- A key function of the substrate is to support chips, remaining flat and rigid even as temperatures fluctuate.
- The substrate includes large redistribution layers that route power and signals from chips above to the circuit board below.
- The RDLs are built one layer at a time, beginning with a first layer that is built on a sacrificial substrate that is later removed.
- Smaller bumps are added to the top of the substrate to connect chips and interposers, and larger bumps are connected to the back of the substrate to connect to the circuit board.
- There can be large differences in how the silicon chips, substrate materials and substrate wiring react physically to changes in heat.
- Soft organic substrates are often infused with glass that adds structural integrity.
- The larger the substrate, the higher the risk of warpage which is a key source of alignment and connection errors that can impact performance, yield and reliability.
- Lior will discuss the growing need for packaging process control.

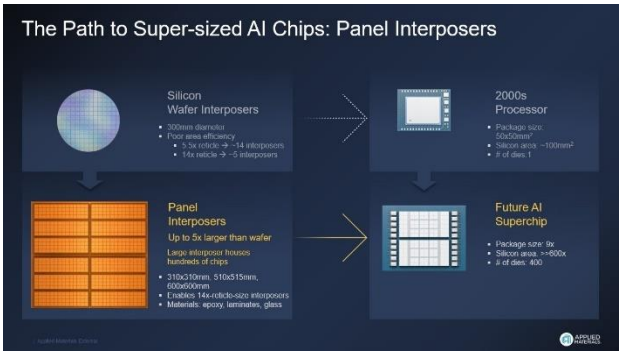

- Now I'll discuss a major emerging trend called panel which will give the industry a larger interposer form factor that enables bigger AI superchips and higher manufacturing output.

- AI accelerator designers commonly describe the size of their products using reticle-equivalent area which I defined earlier, and they hope to deliver 14x reticle-size products within a few years.
- But silicon interposers are limited to the size of a 300 millimeter wafer minus the edge loss that comes from packing rectangular interposers on a round host.
- Only around 4 to 5 interposers of this 14x size can be made from a silicon wafer.
- Fortunately, the packaging industry has long used much larger form factors for chip substrates.
- And today, we aim to enable panel interposers as large as 310 by 310 millimeters, 510 by 515 millimeters, or even 600 by 600 millimeters.
- The bigger the panel, the more interposers that can be built at a time, increasing output and reducing cost.
- Panels can be made of epoxy, advanced laminates or glass.
- Glass is an interesting material because silicon chips and certain types of glass expand at similar rates in the presence of heat.
- In addition, glass is exceptionally flat and rigid.
- It also has electrical insulation properties which enables high signal integrity at high RDL wiring speeds.
- Kevin described how Applied's pioneering work in TSV technology helped Applied become the #1 high-bandwidth memory equipment provider.
- Today we are working years ahead of the panel inflection to build a comprehensive portfolio and help our customers and the AI ecosystem accelerate the panel inflection.

- Our portfolio includes digital lithography used to pattern RDL circuitry that exceeds reticle dimensions.
- Our unique litho technology dynamically compensates for the topographical variations in large multilayer substrates.
- Around 5 years ago, we acquired panel PVD technology, and we've since expanded the platform to include CVD and etch.

- Our long heritage in the display industry helped us refine our large-area CVD capability for panel, and we've ported our display eBeam review technology to panel as well.
- Just recently, we acquired large-area copper electroplating with the acquisition of NEXX.
- Applied's broad panel portfolio enables us to pioneer co-optimized solutions for panel manufacturing challenges.
- One example is through-glass vias which are challenging because copper and glass expand and contract at different rates which can cause panel cracking and failure.
- Applied addresses this by tightly co-optimizing through-glass via etch, CVD liner deposition, PVD seed layer deposition, plating, anneal and CMP steps.
- No other equipment company can do this.

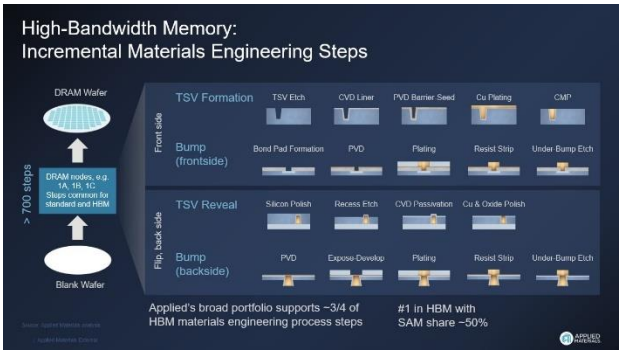

- Next I'll explain how HBM memories are packaged.
- Earlier I described that the spaces for the TSVs are patterned when the wafer is processed.
- Around 19 materials engineering steps are used to complete the TSVs and add metal pillars and microbumps that allow the HBM chips to be connected to one another in a thermo-compression bonding machine.
- Applied offers products for 15 of the steps, which occur on both sides of the wafer.
- Focusing on the front side, the first step is etch which opens the spaces used for the vias.
- Next, an Applied CVD system deposits a highly conformal dielectric barrier film that will electrically isolate the copper TSV from the silicon.
- Then, an Applied PVD system deposits a tantalum nitride barrier that promotes better adhesion of the subsequent liner to the CVD dielectric, a tantalum liner that prevents copper from diffusing into silicon, and a copper seed layer that enables good copper fill.
- Next, an electroplating step deposits copper into the vias.
- Finally, an Applied CMP system removes excess copper and creates a perfectly flat surface for the next process steps.

- Once the TSVs are completed, microbump pillars are formed on the top of the wafer, which is then bonded to a temporary carrier and flipped upside down.
- The back side of the wafer is ground to remove bulk silicon, and CMP is used to thin the wafer to just above the back side of the TSV.
- A highly selective recess etch exposes the TSVs without damaging adjacent structures.
- Next, an Applied CVD system deposits a passivation film that electrically isolates the exposed TSVs and protects the bottom side of the chip from copper migration.
- We use a special, low-temperature CVD film because high temperatures would damage the delicate TSV and DRAM features.
- Our CMP system is again used to make all the structures level and ensure a flat surface for the back-side metal pad formation steps.
- Applied's leadership in the most critical of these TSV formation and interconnect steps has given us the #1 market position in HBM packaging.

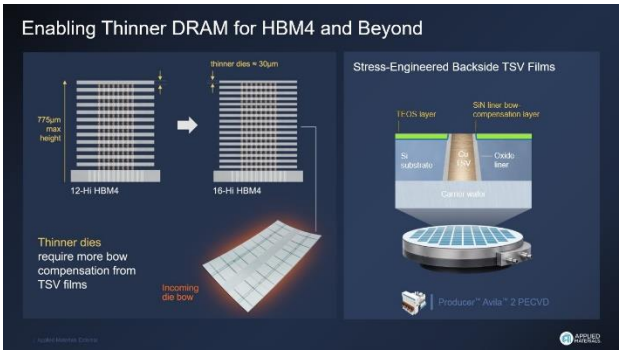

- Applied today introduced a new product that helps our customers work with thinner HBM DRAMs to squeeze more of them into a given stack height.
- HBM dies are thinned to around 1/25th the height of a standard DRAM wafer which makes them susceptible to bowing.
- The higher the stack, the higher the risk of alignment errors, bonding failures and yield issues.
- Our new Avila™ 2 CVD system uses a multi-step sequence to increase the rigidity and reliability of HBM dies.
- It first deposits a thin silicon nitride layer that electrically isolates exposed TSVs on the back side of the wafer.
- The step balances frontside film stress to help prevent bowing.
- Next, the system deposits a thicker oxide film that further stabilizes the chip, creating a robust, uniform surface for the subsequent CMP polishing step.

- Avila™ 2 enables reliable stacking of 12, 16 or even more thin HBM dies.

- Next I'll discuss 3D chip stacking with hybrid bonding.
- Hybrid refers to the fact that both the copper wires and the dielectrics surrounding the wires of two chips are connected in the same operation.
- The first process step is CVD dielectric film deposition which electrically isolates the copper wires of the two chips.
- Next, etch defines spaces for copper bond pads which are aligned with each chip's copper wiring.
- Next, we use an integrated system that uses PVD to deposit a thin barrier liner that prevents copper from diffusing into the dielectric material and promotes good adhesion to the CVD film.
- Next, PVD deposits a thin copper seed layer that promotes high-quality copper fill.
- An electroplating step fills the lined bond pad trenches with copper.
- Next comes one of the most important steps in hybrid bonding.
- CMP removes excess copper and makes the copper wires and surrounding dielectrics perfectly flat and uniform.
- Applied's Opta™ CMP system dishes the copper wires, creating nanoscale divots in the wiring so it is recessed when the two chips are later bonded together.
- Next, a plasma preparation step activates the dielectric films for chemical adhesion and also cleans the two surfaces.
- Next, the chips are perfectly aligned and pressed together in a hybrid bonding machine.
- An anneal step strengthens the dielectric film bonds and expands the copper wiring so it fuses together.

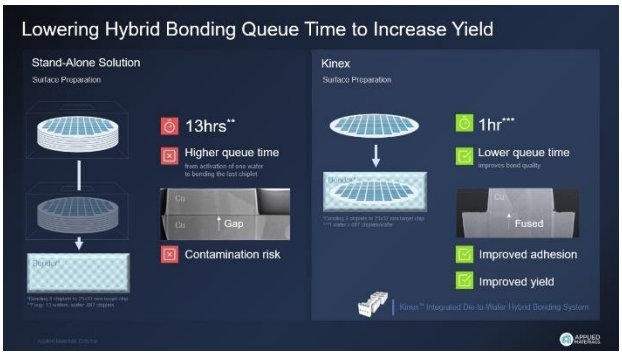

- Applied has introduced Kinex™, the industry's first integrated die-to-wafer hybrid bonding system.
- Kinex™ integrates plasma surface activation and cleaning, bonding, and metrology in one controlled flow.
- Integrated hybrid bonding reduces the queue time between the critical activation and bonding steps, which preserves bonding integrity and improves bond strength and yield.
- The enclosed system minimizes particle contamination, which is critical because hybrid bonding uses known good dies and needs to be nearly 100% defect free.
- Our integrated metrology makes sure the chips are perfectly aligned.
- Kinex™ is designed to help expand die-to-wafer hybrid bonding from a high-end niche application to the mainstream.

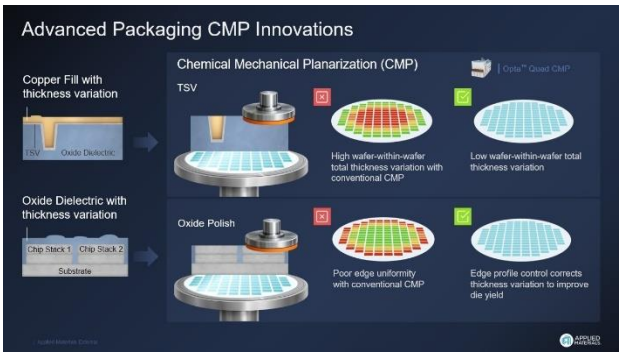

- I mentioned how hybrid bonding consists of several CMP steps.
- Compared to wafer processing CMP, advanced packaging applications introduce new challenges including thicker films, longer polish times, and tighter tolerances.
- Longer polishing times can cause variations from the perfectly level and smooth surfaces chipmakers need in applications like hybrid bonding where even tiny variations in wafer center to edge planarization can cause expensive yield losses.
- Opta™ Quad has a unique, real-time process control system that actively monitors and dynamically adjusts the planarization process on these long-duration, multi-step advanced packaging CMP flows.

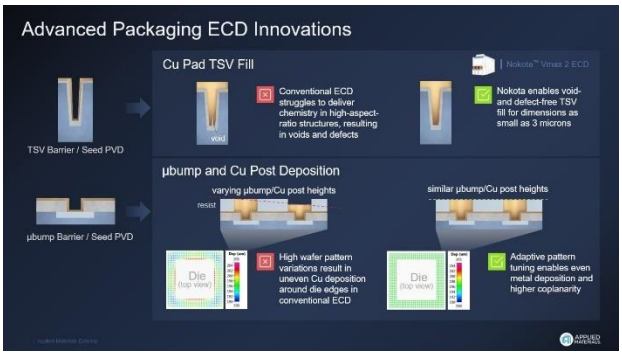

- Today, Applied introduced Nokota™ VMax™ 2, which is a next-generation electrochemical deposition system purpose-built for advanced packaging applications like HBM.
- Plating is a foundational step across interconnect schemes, from TSVs within dies to microbumps that connect stacked dies.
- As TSV dimensions shrink, it becomes increasingly difficult to deliver chemistry deep into narrow features while maintaining high-quality fill.
- Nokota™ VMax™ 2 is optimized for void-free, defect-free metal fill even below 3 microns.
- HBM stacking also raises the bar for plating tools, requiring tighter bump coplanarity and higher throughput.
- A key challenge is pattern density variation across the wafer, which distorts current distribution and leads to non-uniform metal deposition.
- Nokota™ VMax™ 2 addresses this with Adaptive Pattern Tuning, dynamically shaping the electric field to correct for layout-driven variation and improve uniformity and coplanarity.
- This control extends to both sides of each HBM layer, which is critical in multi-pass plating flows.

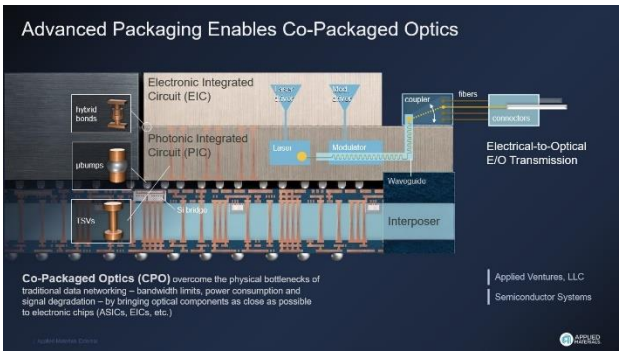

- My final topic today is co-packaged optics – or CPO.
- Optical fibers have long been used to transmit digital data at high speeds over longer distances, such as between racks and even between data centers.

- In co-packaged optics, electrical signals are converted to light waves that are transmitted through optical fibers.
- Co-packaged optics is enabled by advanced packaging technologies including hybrid bonding that bring electronic and photonic chips close together so they can operate as a single, tightly integrated system.
- Optical modulators convert electronic signals to light signals.
- Light-generating devices such as lasers are integrated close to or even on silicon chips.
- Photonic waveguides can guide light across chips.
- They are fabricated by depositing a silicon or silicon nitride layer that is etched to form narrow features that light can travel along.
- Fiber couplers transfer light between photonic chips to external optical fibers.
- To help develop the optical I/O ecosystem, Applied Ventures is partnering with optics and photonics startups along with manufacturers including GlobalFoundries.
- We also have R&D programs in the materials, process technologies and equipment needed to produce and integrate optical components.
- Our equipment opportunities include epitaxy, CVD, PVD, etch, CMP and process control.

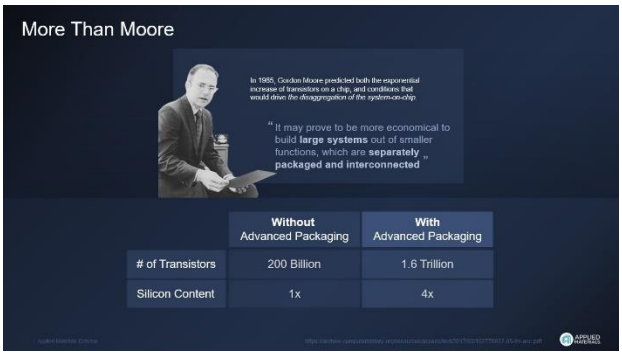

- Now I'll summarize by revisiting what I said at the beginning – that the future of chiplets that Gordon Moore predicted would be a key driver of AI today.
- If we had stayed with system-on-a-chip and Moore's Law scaling, then today's server chips would have around 200 billion transistors.
- Thanks to advanced packaging, today's AI server chips have over 1.6 trillion transistors, and pack 4 times the silicon content in every socket.
- By going beyond Moore's Law, we're enabling new applications like training, inferencing and agentic AI that depend on advanced packaging for performance, power and cost.
- Now I'll hand the meeting over to Lior Engel to discuss packaging process control. Lior?

- Thank you Jinho.

- The packaging market has long been served by low-cost optical and physical process control tools.
- But particles that were acceptable when large interconnects and bumps were used are turning into yield killers in hybrid bonding applications.
- Structures that could easily be measured using optical tools are becoming so small that eBeam resolution is needed.
- And delicate wiring can be damaged by wafer probers.
- So as packaging advances from 2D to 2.5D and 3D integration of known-good dies, packaging fabs increasingly need the same advanced process control technologies found in wafer fabs.

- I'll briefly describe four key benefits of eBeam in advanced packaging.
- One, eBeam columns produce a finer beam than optical tools.

- While optical systems struggle to image tiny defects and sub-micron features, eBeam can image single-digit-nanometer features and measure them with sub-nanometer precision.
- Second, eBeams can be tilted to profile sidewalls and measure high-aspect-ratio structures.
- Third, eBeam enables materials identification, which is critical to determining the corrective actions to take when yield issues arise.
- And fourth, eBeams can test the electrical functionality of delicate circuitry without making physical contact.
- These attributes make eBeam valuable in advanced packaging where proper manufacturing of features like TSVs, micro-bumps and hybrid bonding pads is critical to chip performance, power and yield.
- Applied has been in the eBeam business since the 1980s, and today we have the #1 market position.
- We are now transferring our eBeam expertise to the advanced packaging market with new products that will help our customers solve critical challenges

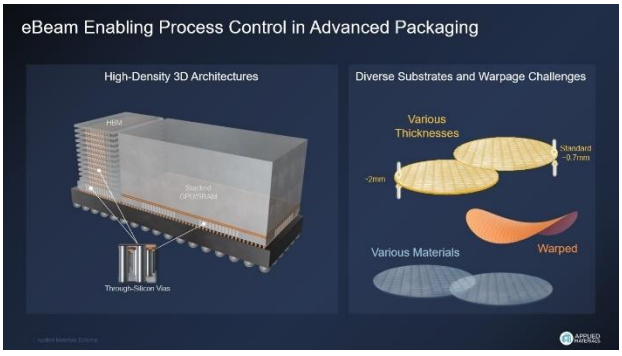

- Jinho's presentation has given us insights into how eBeam technology can help.
- For example, in stacking SRAM cache memory chips on top of processors.
- Also, in stacking HBM chips, transitioning from 12-high designs to 16-high and beyond.
- And in TSV critical dimension metrology to ensure full HBM stack connectivity.
- Another challenge in advanced packaging is handling the wide range of substrates Jinho described.
- This includes new materials like glass, as well as geometric considerations including thickness and warpage.
- Today, we are introducing two new eBeam products for our advanced packaging customers.
- Both were developed to bring the full benefit of eBeam technology to a wide variety of substrate geometries and materials, from silicon to organic and glass.

- Applied's VeritySEM™ eBeam system is used in wafer fabs to measure critical dimensions such as the height, width and slope of gate all around transistor structures.
- Our new VeritySEM™ 7AP is designed for advanced packaging applications including measuring hybrid bonding pads, TSV dimensions and microbump geometries.
- The system automatically configures itself to work with different substrate materials and geometries.
- It works in-line, delivering precise 2D, 3D, and high-aspect-ratio measurements including a unique tilted view for height and sidewall angle measurement.

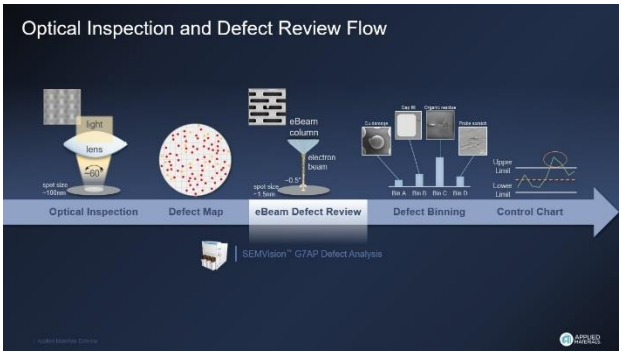

- Next, Applied's SEMVision™ is the #1 eBeam defect review system in wafer fabs with over 70% share of the global market.
- Our new SEMVision™ G7AP is designed for advanced packaging fabs, automatically handling the full variety of substrates.
- Customers use optical tools to quickly scan for potential yield-killing defects which are often too small for the optical technology to identify clearly.
- Our new SEMVision™ G7AP uses the superior resolution of eBeam to clearly identify and automatically classify these defects, enabling customers to rapidly address the root cause of advanced packaging yield issues.
- VeritySEM™ 7AP and SEMVision™ G7AP are already in production at multiple leading advanced packaging customers in both memory and logic.
- Applied has additional advanced packaging process control systems in development, including optical and x-ray products that we will discuss in the months ahead.

- And now I'll turn the meeting back over to Kevin Moraes. Kevin?

- Thank you Lior.
- In a few minutes, we'll begin the Q&A.
- But first I'll share a few examples of how our inflection-focused innovation strategy is driving growth for Applied in both DRAM and advanced packaging.

- Back in our 2021 Memory Master Class, we projected that scaling our Draco™ CVD patterning film, Sym3™ etch and PROVision™ eBeam technology could give us a $1 billion cumulative revenue opportunity between 2020 and 2024.
- In fact, we delivered more than a billion dollars in sales from these products over that time,
- And today, we believe we can generate more than 4 times that over the next 5 years.
- We also projected that improving DRAM periphery transistors and wiring with foundry-logic technologies could give us a $2-billion-dollar cumulative opportunity from 2020 through 2024.
- We delivered more than 2 billion dollars from these applications over that period.
- And we believe we can generate more than 3 times that over the next 5 years.
- Turning to advanced packaging, we projected in 2021 that our business would more than double between 2020 and 2024 to over $1 billion dollars.
- In fact, it more than tripled.

- And this year, we expect our advanced packaging revenue to increase by more than 50% to more than $2 billion dollars.
- We continue to use our inflection-focused innovation strategy to predict roadmap inflections and deliver new products that enable our customers to accelerate their roadmaps.

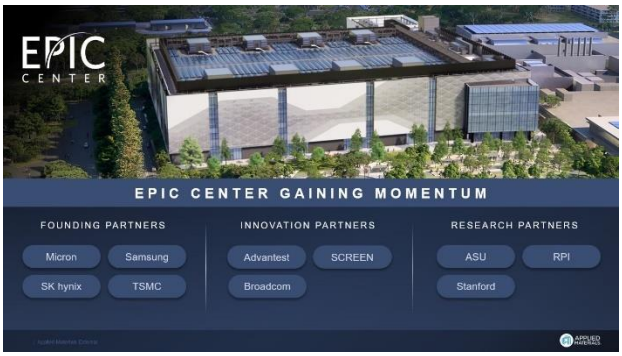

- We hope to see many of you at the unveiling of the new EPIC™ Center in Silicon Valley in October.
- At EPIC™, we will take inflection-focused innovation to a new level by co-innovating with our customers in the industry's largest collaborative research facility.
- Many of our leading customers and partners have already announced plans to join us at EPIC™, and there will be more announcements in the months ahead.
- In addition, we are completing the construction of a new EPIC™ advanced packaging center in Singapore.
- Now it's time for me to join the rest of the team for our Q&A session.