New Ways to Wire and Integrate Chips Master Class

PREPARED REMARKS | May 26, 2022

MICHAEL SULLIVAN | Corporate Vice President, Investor Relations

Hello and welcome to the second event in our 2022 Master Class series which is called, New Ways to Wire and Integrate Chips. The technologies we'll cover today will help the industry to accelerate improvements in chip power, performance and area-cost even as traditional Moore's Law slows.

Here is today's agenda. In a moment Dr. Kevin Moraes will explain how to improve chip performance and power with wiring innovations that solve the resistance challenges associated with EUV scaling, then Dr. Mehul Naik will show you how we can improve chip power, area and cost by powering transistors from the backside of the wafer. Next, Dr. Sundar Ramamurthy will update you on how we can improve all areas of PPACt for chips and systems using heterogenous design and integration, using hybrid bonding and advanced substrates. Finally, Dr. Raman Achutharaman will show you how Applied expects to grow faster than the WFE market in these areas. And after the presentations we'll host a Q&A session.

In the fall we plan to hold our Subscriptions and Services Master Class. We'll give you insights into the growth of Applied Global Services and show you how we're transitioning more of our services and parts business from transactions to subscriptions. Before handing the meeting over to Kevin, I'd like to invite Regina Freed to join me in the studio for a brief fireside chat to explain why these innovations are important to our customers.

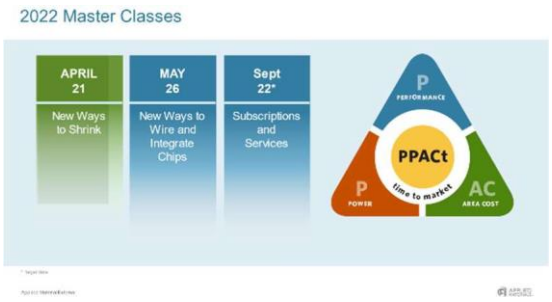

REGINA FREED | Vice President, Semiconductor Products Group

Mike Sullivan: **Welcome back, Regina.**

It's my pleasure. Thank you, Mike.

Mike Sullivan: **Regina, the topics we're covering in today's masterclass were actually raised at a technology event that you hosted a couple years ago. Tell us about that.**

Sure. One of the last industry events we had before COVID was the IEDM conference in San Francisco in December of 2019. I chaired a panel we called The Future of Logic, EUV is here, Now What?

Mike Sullivan: **Who attends IEDM and who was on your panel?**

IEDM is an IEEE conference that attracts the world's leading designers of logic and memory chips. Our panel had the top logic researchers from TSMC, Intel, Facebook, IBM and Stanford.

Mike Sullivan: **What did you guys mean by EUV is here, now what?**

In my mind it was time for celebration because EUV was really here. We could finally print features at the 25nm pitch in a single exposure and that could enable 2D scaling to continue well into the future. That's the good news.

Mike Sullivan: **If there's good news that means there was bad news. What was the bad news?**

At the same time, it was no time to relax, especially for us, because making things smaller doesn't make them better from an electrical point of view. We needed this panel to focus the industry on the new materials engineering challenges associated with EUV scaling so we could hope to solve them with new technologies.

Mike Sullivan: **In today's talk why don't you give us the top three takeaways from your panelists?**

Sure. My favorite quote was from one of our largest foundry/logic customers. He said, "Houston, we have a problem. The interconnect resistance is very high and if we don't solve the interconnect problem, we won't have any more transistors."

Mike Sullivan: **Did he literally mean that we are not going to be able to have more transistors?**

He said, "Even if we can develop a better transistor, you won't see the benefits of it because everything will be dissipated by the interconnect resistance." Another of our panelists said, "And the same thing goes for wiring. You have to have new materials and, more importantly, you need to have a new process that enables that, like selective deposition."

Mike Sullivan: **So basically you can spend a lot of money shrinking with EUV and making better transistors like Gate All Around, but the wiring resistance is going to eat away the performance benefits. So, what do we do about it?**

Fortunately, I think Kevin will give you the solution to this today and I don’t want to steal his thunder.

Mike Sullivan: **Great. So, what else came up at the panel?**

Well, another one of our panelists said that the third evolution of Moore's Law is design technology co-optimization, or DTCO, and she highlighted resources in backside power distribution.

Mike Sullivan: **Did everybody know what she was referring to two years ago?**

Funny enough, one of our panelists hadn't even heard about backside power yet and another said it wouldn't be ready any time soon. So here we are two years later, and backside power is now on everyone's roadmap. I think it's amazing how fast this trend is developing. I'm really looking forward to Mehul's presentation today because he's going to tell us how we can make it happen.

Mike Sullivan: **So, the third evolution of Moore's Law is DTCO and techniques like backside power. What's left?**

One of our panelists said there's a fourth evolution of Moore's Law coming and that's heterogeneous design. She said we can drive costs lower by designing the chips made on a variety of nodes and then stitching them together using advanced packaging.

Mike Sullivan: **And this was a designers' conference so did they know whether this was actually technically feasible?**

Actually, yes. Another one of our customers described this as bringing high-end fab technology into the packaging arena. And I'm looking forward to Sundar's presentation today because compared to two years ago we have a lot more technology that can do this.

Mike Sullivan: **Great. Well, thank you Regina for summarizing your panel. It turned out to be a real crystal ball for all the roadmap changes ahead, and I really hope we can get back to doing a lot more panels like it. And now let's turn the meeting over to Kevin to help us answer the question, EVU is here, now what. Kevin.**

KEVIN MORAES | Vice President, Products and Marketing, Semiconductor Products Group

Thank you, Mike and Regina.

My job today is to show you how we are reinventing wiring to solve the resistance challenges of EUV scaling so we can help give designers more transitions to work with. Then, Mehul will show us how we can make room for even more transistors by taking advantage of all the real estate on the backside of a wafer to deliver power. Finally, Sundar will show us how we can use similar wiring innovations to integrate multiple chiplets, all as one, giving designers an unlimited transistor budget. I've been working in these areas for 20 years, and I think the pace of innovation is beyond anything we've seen in the industry today. As consumers, we have some amazing experiences to look forward to.

Now, let's begin at the nano scale. The smallest wires in a chip are all the contacts that power transistors, gate, source and drain. The gate determines whether the transistor is on or off. And when it's on, electrons flow from the source to the drain. We need a little bit of power to turn the gate on and off, and more power at the source drain to enable the transistor to switch. Contacts are formed using one or two

metal layers, and their job is to connect the transistors to the surrounding interconnect wiring. Within each logic cell, the interconnect wiring connects transistors to help us perform simple operations like add, subtract and compare. We can create more complex logic cells or connect a number of cells together to perform more complex calculations and operations. The small cell interconnect wiring are what we call metal zero. Above are metal one, metal two and so on. The interconnects grow wider and longer as you proceed up through the stack. Vertical wires called vias connect each layer up to the next layer. Signals routinely travel up through the contact in vias, across to the interconnects, and back down again through the vias and contacts as we process data. The electrical challenges to EUV scaling are all found in these smallest contact vias and interconnects. And this is where materials engineering innovations are needed.

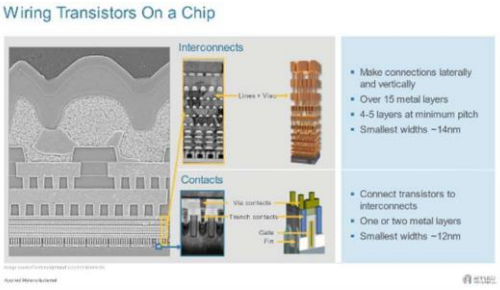

Next, let's focus on building the contacts. Since we are wiring at the atomic scale, we need to manage the physics of the materials we're using. The transistor is a semiconductor that is highly engineered to switch at particular voltages. Transitions are made with pure crystal and silicon, built with positive or negative ions, and stressed with epitaxy that pushes or pulls the silicon lattice in very specific ways to optimize performance. We can't just use any metal for the contact because the atoms of a metal like copper are mobile, and can diffuse into the silicon, which rules the switch. That's why we make the contacts using metals like tungsten and cobalt, which have larger atoms and diffuse more slowly. We can also deposit barrier layers between the insulating dielectric and the contact metal to further protect the channel from the contact metal and the processes used to deposit this metal.

The interconnect wiring is isolated from the transistor. Chips often contain over 50 miles of interconnect wiring. And we like to use copper because it has a very low resistance. To make each level of wiring, we first etch patterns into a layer of dielectric material. We line the trenches we create for the barrier metal that helps prevent copper from migrating into the dielectric and nearby transistors. We also deposit liners atop the barrier metal that are designed to make it easier to deposit copper into very small spaces. Unfortunately, the barrier and liner don't scale very much. As a result, as we scale with EUV, the trenches

become narrow and the proportional space consumed by the barrier and liner go up, while the proportion available for copper wiring goes down. And the smaller the copper wire, the higher the electrical resistance. So as Regina said, Houston, we have a problem. It's called interconnect resistance.

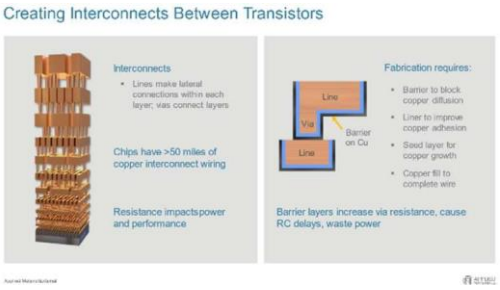

Here's some simple math. Imagine you use EUV to shrink the pitch of metal lines by 10%. Now the width of these lines are reduced by 0.9 times 0.9, which means that the area is reduced by nearly 20%. Also, the liner and barrier don't scale. So, the area might actually shrink by 40%. When that happens, resistance goes up exponentially. The resistance dissipates the chip's electrical signals, which reduces performance, wastes power and creates heat. That's the issue raised on Regina's panel. Even if you can develop better transistors, you won't see the benefits of it because everything will be dissipated by the interconnect resistance. Fortunately, we can attack this problem in three ways. First, we can identify where there is high resistance at the interface of two metals and develop substitutes. Second, we can look for alternate metals that exhibit lower wiring resistance. And third, we can develop new deposition methods that reduce or eliminate the use of barriers and liners.

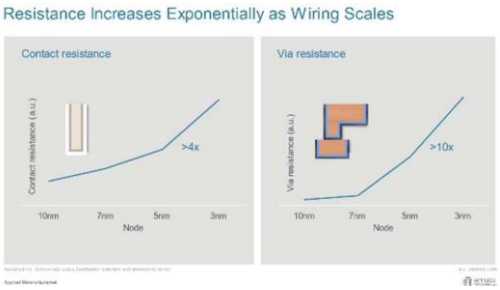

Applied has been leading the industry in developing new technologies that use these strategies to reduce contact resistance. For example, to help the industry scale to the 10nm node, we introduced a metal organic tungsten liner solution that became a substitute for a high-resistance titanium nitride liner. At the 7nm node, we created the industry's first cobalt contacts as an alternate to tungsten. We found a cobalt chemistry that allows us to thin down the barrier layers and still protect the transistor channel from the metal and the metallization process technologies. As EUV came in at the 7nm and 5nm nodes, we introduced the industry's first selective tungsten deposition process, which eliminated the barrier and liner, thereby maximizing the volume of the contacts and lowering resistance.

The selective tungsten deposition for contacts we introduced in 2020 is a major breakthrough. But it works best only in structures formed within a dense dielectric material, and when there's a pure metal in the layer below that, that tungsten can grow on. Today, we are announcing a new breakthrough we call Ioniq tungsten, that enables us to deposit pure tungsten in a much larger and less restrictive set of contact and via applications. Ioniq tungsten is an Integrated Materials Solution (IMS) that includes three processes in a single high-vacuum system. We initiate the deposition process with an integrated surface preparation step that reduces the interface resistance between the two metals. Next, we use a highly ionized PVD process that deposits a pure low-resistance tungsten conformer liner, which acts both as a barrier and as a seed layer upon which we can grow metal in the next step. Then we use a CVD process that is optimized to deposit the tungsten onto the thin liner and complete the metal fill. Our Ioniq IMS solution deposits pure tungsten in more applications, enabling our customers to lower resistance and scale with EUV.

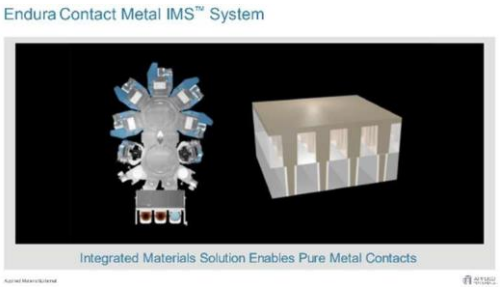

We are already seeing broad high-volume adoption of all our unique contact metallization technologies. Our customers can combine these solutions to achieve particular power and performance goals in their chip designs. Later, Raman will talk about the rapid growth we are seeing from these solutions as they help customers scale with EUV and manage contact resistance.

So now I'll discuss interconnect wiring, which is another leadership area for Applied Materials. At each new node, customers pattern the smallest interconnect wire that can be filled with copper. Lithography is not the limiting factor. Materials engineering and interconnect resistance are the challenges. I'll first explain how we create the critical interconnects at the 7nm node.

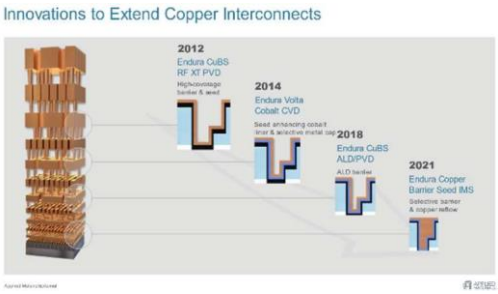

After patterning, we etch tiny trenches into a layer of dielectric material. The wafers are transferred to an Endura Integrated Materials Solution that integrates a series of steps in high vacuum. First, a surface treatment step removes any contamination. Next, we line the trenches with a diffusion barrier metal, like tantalum nitride, which keeps copper atoms from penetrating the surrounding dielectric and migrating within the chip. We then deposit a liner layer made of a pure metal like tantalum or cobalt that copper can adhere to. Then we deposit a thin copper seed layer and we fill the remaining space with bulk copper using a wet process called electrochemical deposition, or ECD.

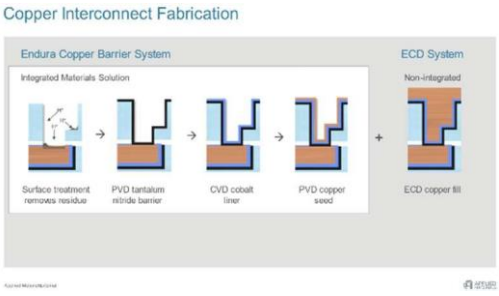

At the 5nm node, customers want to use EUV to create even smaller wires. There are two major challenges to scaling wires. First, the electrical resistance increases exponentially, and second, wet ECD processes cannot reliably deposit copper into spaces narrower than 20 nanometers, as is needed at 5nm and below nodes.

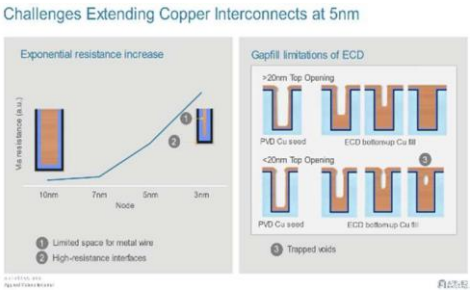

I'll explain how we are enabling the scaling to continue, using materials engineering beginning with the resistance issue. At last year's Logic Master Class, we unveiled a new Integrated Materials Solution that selectively deposits barrier metal only on the side walls of the vias that connect each layer of the critical metal layers. Our solution performs surface engineering and then selectively deposits the barrier to the via walls only using ALD. By eliminating the high-resistance barrier from the bottom of the vias, we reduce the via resistance by up to 50%. When you consider that there are four to five layers of critical wiring at this pitch, and billions of vias, you can see how materials engineering helps us compensate for the resistance challenges of EUV scaling.

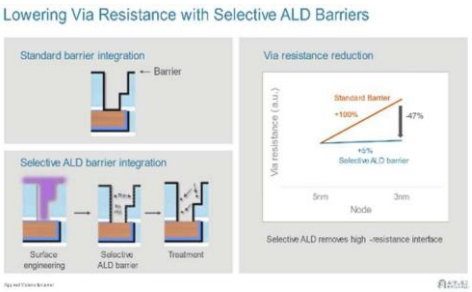

Now I'll address the copper fill challenge. After we selectively deposit the liner and barrier to the sidewalls, we move the wafer to the next in-vacuum chamber to initiate a unique dry copper deposition sequence. First, we deposit a thin layer of copper using PVD. Next, within the same chamber, we heat the wafer, which causes the copper layer to reflow like honey, partially filling the small feature. We need to repeat the copper deposition and reflow process to reliably fill all the remaining space with pure copper, leaving no gaps or voids. To further scale with EUV, and extend copper to future nodes, we can add more of these dry PVD deposition and reflow cycles to create pure copper wires in even smaller spaces.

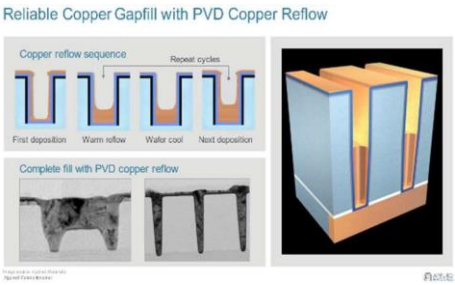

To summarize, at the beginning of today's meeting, Regina asked a question, EUV is here, now what? The answer is, the new IMS solutions we've created to solve the contact interconnect resistance issues of EUV scaling. These wiring innovations are having a dramatic impact on our business. Later in the meeting, Raman will show you how our wiring process steps and revenue are expanding as we continue to scale with EUV. Thank you for listening. And now I'll hand the meeting over to Mehul.

MEHUL NAIK, PH.D. | Managing Director & Principal Member of Technical Staff, Semiconductor Products Group

Thank you, Kevin.

Earlier, Regina discussed how our major customers were just beginning to talk about Backside Power Delivery two years ago. Today, Backside Power is firmly on the roadmap and likely to arrive soon after the first Gate-All-Around transistors begin to appear. Customers will have unique timetables and design strategies. Backside Power is important for power and area-cost reasons. At our New Ways to Shrink Master Class, Uday explained that there are two ways to increase the logic density. One is intrinsic scaling or classic 2D scaling with EUV. The other is DTCO, which involves clever changes in logic cell layouts that increase density, independent of the lithography pitch. Backside Power is a form of DTCO.

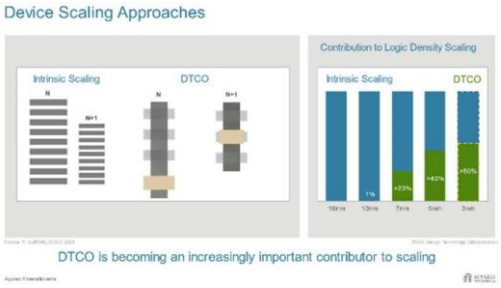

This chart summarizes the PPACt improvements the industry is achieving with the progression from planar to the high-k metal gate transistors to multiple generations of FinFET, and soon Gate-All-Around. The combination of intrinsic scaling, new transistor architectures and DTCO techniques is improving performance

by 3.5X, area by up to 5.5X and power by more than 6X over this period. Gate-All-Around should arrive in high volume within the next two years and give us several generations of improvements in power and performance. The question is what will give us the next big opportunity for logic density scaling as Moore's Law slows? I believe the answer will be from rethinking the power delivery network.

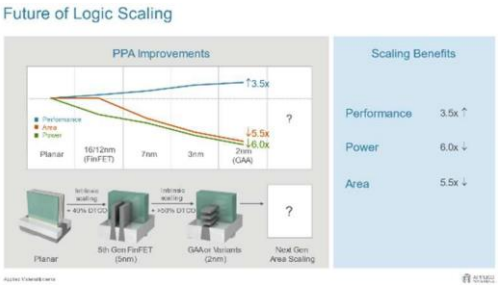

The network's job is to route power from the off-chip power regulator to the transistors, delivering a specific voltage so that the transistors can switch quickly and reliably when the gates are turned down. Today, the network routes power from the front side of the chip. The power wires pass through each of the 12 or more metal layers of the chip, and at each level there is wiring resistance that reduces the supply voltage. Typically, the network has a design margin that can tolerate a 10% voltage drop between the regulator and the transistors. As we continue to scale with EUV however, we encounter high resistance, along with routing congestion. As a result, we may not be able to scale beyond 3nm using existing techniques without experiencing voltage drops as high as 50% and serious transistor reliability issues. We are hitting the physical scaling limits of the traditional wiring approach.

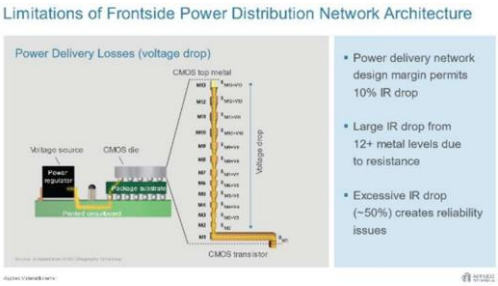

A second issue with today's power delivery networks is logic cell area scaling. As a reminder, a logic cell is a minimum set of transistors and wiring needed to perform simple operations. Within the logic cells are power lines called rails that need to be a certain size to deliver enough voltage for the transistors to switch. They do not scale as well as the transistors and signal wires. As a result, the power rails are now about three times wider than the other elements. If we could remove the power rails from the cells, we could significantly increase logic density, even without changing lithography.

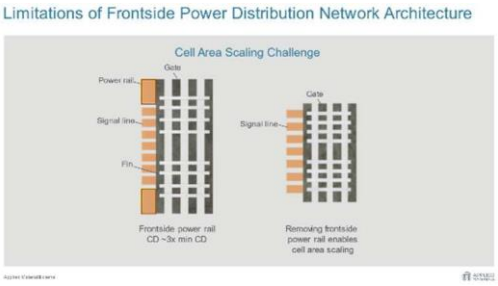

A Backside Power Delivery Network would circumvent the 12 or more wiring levels of the chip. The copper wires would be built on the backside of the wafer, where there is no routing congestion. Resistance would be lower and voltage drop between the external power regulator and the logic cells would be reduced by as much as 7X. Ideally, the power rails would be moved outside of the logic cells, enabling logic density to be scaled by up to 30%, which is equivalent to as much as two generations of lithography scaling.

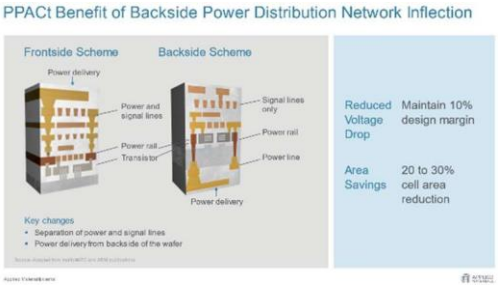

Returning to our roadmap chart, the Backside Power distribution network completes the picture. It complements the performance and power benefits of Gate-All-Around transistors with the mechanism to scale the area-cost as well. Together, Gate-All-Around and Backside Power networks will allow the industry to continue to advance PPACt.

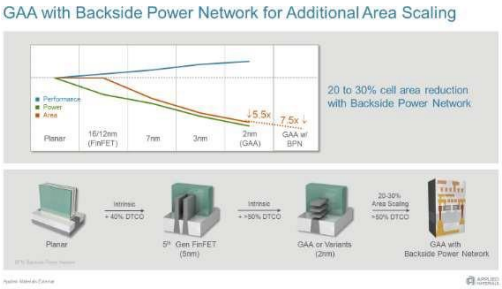

Now let's talk about the elements of a Backside Power distribution network. We will need an external voltage regulator that can supply power to the backside of the wafer. A network of copper lines will route power to areas just below the logic cells. Finally, and most importantly, power will be delivered to the logic cells and individual transistors. There will be design trade-offs. The simplest approaches will be easier to manufacture. More complex approaches can improve power efficiency and area scaling. And the most complex approaches can maximize power and area scaling.

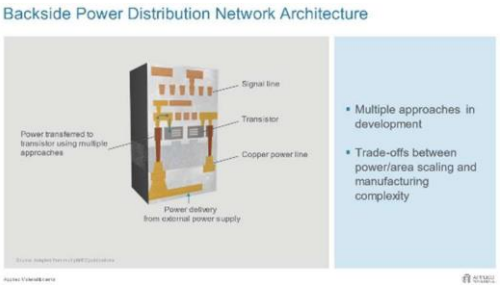

Based on publicly available information, designers are evaluating three distinct architectures. In the first approach, the logic cells will retain a power rail. The Backside Power distribution network will be connected to the power rail by a nano-scale through-silicon via or a nano-TSV. This power routing scheme consumes some of the available power and the area scaling is modest because the logic cell needs space for the buried power rail and a nano-TSV. In the second approach, there is no power rail in the logic cell. Instead, a power via directly transfers power from the backside network to the cell or the transistor contact. This approach is more complex, but it improves power efficiency and increases cell area scaling. In the third approach, power from the backside network is connected directly to each transistor's source drain. This approach is the most complex to manufacture, but it offers maximum power efficiency and scaling. It'll also require major innovations, which I'll describe shortly. We expect different companies to pursue different Backside Power strategies over time. And the industry may not converge on any one particular approach.

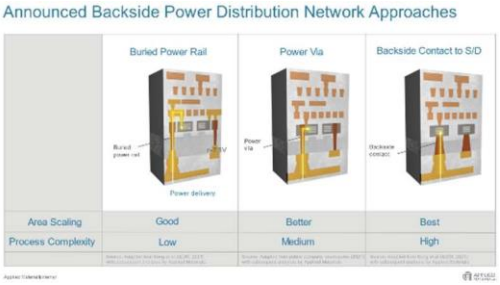

From an equipment and materials engineering perspective, Backside Power can use many of the technologies used today in areas like transistor contact engineering and through-silicon vias. The more basic approaches can use unit process tools. The most advanced approaches will need co-optimized solutions and Integrated Materials Solutions where multiple existing steps take place in the same system under vacuum. Next, I'll provide a high-level overview of how we might build the most optimal Backside Power scheme. First, the front side of the wafer is fully processed, beginning with the transistors and then all of the contacts and interconnect wiring that Kevin described. Next, we bond a carrier wafer to the top of the wafer. The carrier wafer adds stability, enabling us to make radical changes to the backside of the wafer without harming all of the delicate structures we have created. To contact the transistors from underneath, we need to get close. We flip the wafer over and use grinding and etch techniques to remove all but about one micron of the backside. But we still have too much silicon thickness and thickness variability to create transistor contacts. So next, we use a precise CMP step to get even closer to the

transistors and reduce the thickness variability to just a few nanometers. Applied's Reflexion LK Prime CMP enables the step with in-situ metrology and real-time process control. Next, we etch into the backside of the wafer near the transistors to create trenches that we fill with a dielectric material to isolate the transistors from the Backside Power. This is followed by another CMP step. Applied's Eterna Flowable CVD and Reflexion LK Prime Dielectric CMP platforms are being optimized for these steps. Next, we etch even deeper into the backside of the wafer, at very precise locations, to create vias that reach the transistor contacts. Applied's Sym3 etch for nano-TSVs is well suited to this delicate application. Now we create the backside contact to the transistor source, which is key to transistor power and switching performance. This is much like the contact engineering Kevin described, except that we have a thermal constraint. We need to engineer the contact to have the lowest possible resistance. And this normally requires high-temperature epi and anneal processes. However, the backside contacts are being manufactured with the frontside transistors and interconnects in place, and they would be degraded by these high temperatures. To manage this, Applied is building on our leadership in epitaxy, implant and anneals to develop a co-optimized, low-temperature solution for backside contact engineering. Next, we engineer the metal interfaces to the backside contacts. To minimize resistance, Applied is developing an Integrated Materials Solution that combines up to seven steps in high vacuum, including chambers for pre-clean, selective silicide deposition, ALD or PVD liner deposition and a new metal fill. We are also developing a co-optimized CMP step that leaves us with a perfectly uniform backside contact layer. Now that the transistors have backside contacts, we can build a copper backside power distribution network using techniques similar to the ones Kevin described earlier.

In summary, Backside Power Delivery Networks are an exciting new development in the industry's PPACt roadmap, promising to increase logic density by up to 30%, which is equivalent to around two nodes of lithography scaling. The simplest approaches will deliver some of these benefits, while delivering the full potential will require new innovations in materials engineering. Applied Materials is already developing co-optimized and integrated materials engineering and methodology solutions to help accelerate our customers' progress. Thank you for listening. And now I'd like to hand over the meeting to Sundar.

SUNDAR RAMAMURTHY | Group Vice President, GM, Semiconductor Products Group

Thank you Mehul.

As one of our customers said, there is a fourth evolution of Moore's Law coming, and it's heterogeneous design. That's a very apt description because heterogeneous design and integration will be like Moore's Law in the Dennard scaling era when a single strategy gave us simultaneous improvements in power, performance, and area-cost. We will soon be able to recover this benefit as an industry, even as traditional Moore's Law slows.

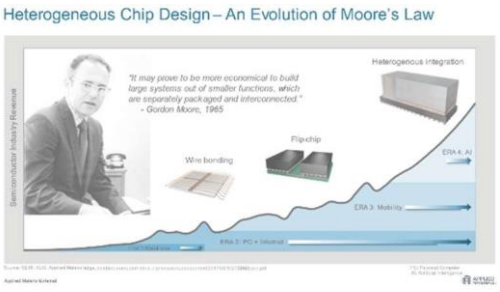

I'll begin by summarizing why we need heterogeneous design and integration. In the Dennard scaling era through the early 2000s, logic die sizes were relatively constant and every two years, we got twice the number of transistors per area to work with. There was a single CPU core, shown here in black, and we ratcheted up the power and frequency, shown here in red and green, to generate more computing operations per second. This in turn made existing software programs and data processing run faster. The trend was sufficient to propel the PC industry and enable a large server industry based on PC architectures. When Dennard scaling ended, we couldn't get more performance without excessive power

and heat. So, we adopted multi-core CPU architectures and parallelized the software to distribute the computing workloads over multiple low-power processors. On the chart, you can see how the number of cores grew while the power and frequency remained constant. After 2010, general purpose workloads fell victim to Amdahl's Law and hit the limits of parallelism, causing performance to plateau, which you can see in blue. But other workloads like 3D graphics and machine learning continued to accelerate as we added more cores. In recent years, we have seen a bifurcation of computing architectures with general-purpose workloads running on CPUs and specialized workloads running on application accelerators based on GPUs and TPUs. Today, transistor counts continue to increase as you can see in yellow. But with Moore's Law slowing and the cost of the most advanced nodes increasing, we are reaching the limits of cost-effective monolithic scaling across most computing architectures.

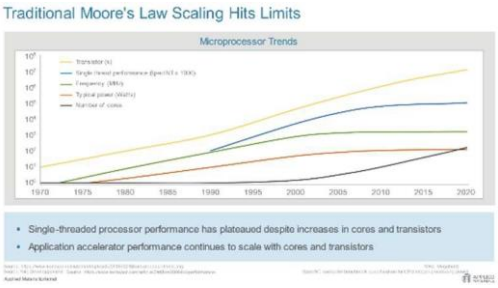

A major challenge is the reticle limit. The largest mask pattern we can print on a wafer is around 26 by 33 millimeters, or just about 860 square millimeters. When Moore's Law was working well, designers could put a large number of high-performance PC and server chips in that space, or a small number of extremely high-performance chips. But with transistor counts now growing faster than 2D scaling, chip sizes are increasing, and many companies are hitting what we call the reticle limit. This forces the industry's transition to heterogeneous design and integration.

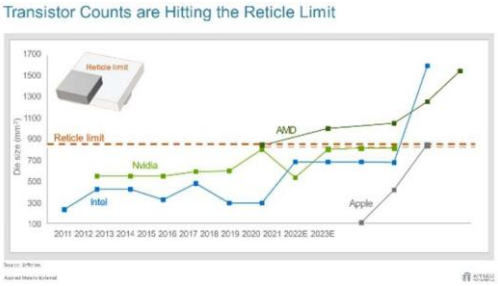

One particular problem is SRAM scaling. SRAM cache memory can easily consume a third of a CPU's die area. The cache is critical to system performance because it keeps the data and program instructions the processor needs right next to the processor. Each SRAM requires six transistors and the SRAM designs have become highly optimized to the point they barely scale with the rest of the logic circuitry.

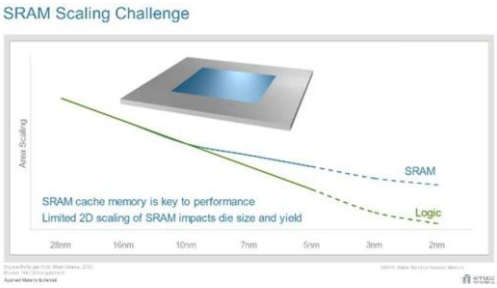

In Kevin's class, you saw that each chip includes back-end-of-the-line interconnect wiring. Conceptually, if we can directly connect two chips using these back-end-of-line interconnects, we can make the heterogeneous die perform as one. A new and very promising example is that we can bond a large cache memory chip to a CPU chip, enabling us to overcome the reticle limit, increase performance, and reduce die size and cost.

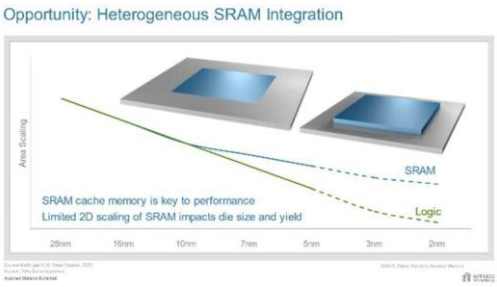

There are other technologies that don't scale well or can't easily be integrated on the same silicon used by CPUs and accelerators. These include DRAM and flash memory, analog, power and optical. Using advanced substrates and advanced packaging technologies such as Through-Silicon Vias, we can bring all of these technologies closer to the logic and improve the entire system. As an industry, we can transition from the system-on-a-chip era to the system-on-a-package era.

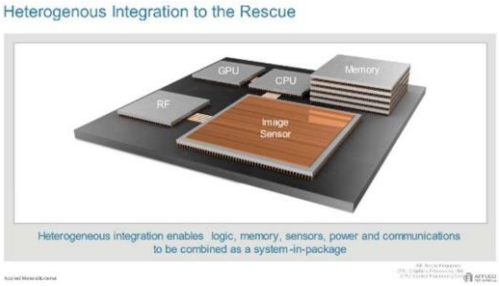

To accomplish this, two major breakthroughs in heterogeneous integration are needed. First, we need a new technology that can efficiently connect individual logic and cache die to one another using their

copper interconnect wiring. This is called hybrid bonding. Second, we need new types of advanced substrates for heterogeneous chip integration. I'll now discuss each of these innovations.

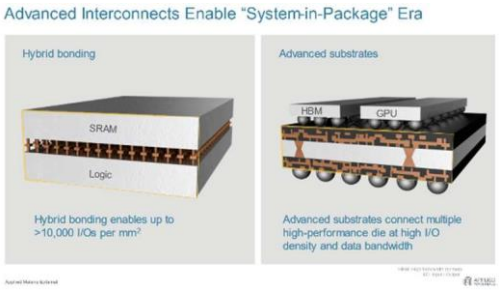

Hybrid bonding connects die that are produced on two or more different wafers. The final copper interconnect layers of each chip are united using new process technologies. Hybrid refers to the fact that we are connecting both the copper pads that connect power and signals and the surrounding dielectric material. The bonded layers can have as many as 10,000 connections per square millimeter. There are two approaches. One is called wafer-to-wafer hybrid bonding. We match up and bond two wafers and then singulate the wafers to produce individual chips. The second approach is called chip-to-wafer hybrid bonding. We start with one unsingulated wafer and attach individual die with the help of advanced pick-and-place technology that aligns each die with its mate with high speed and precision. After the die are bonded, we singulate the wafer. Both approaches require extremely precise alignment of the copper interconnects and surrounding dielectric. Specialized materials engineering is needed to create perfectly bonded contacts, which can number in the hundreds of millions per wafer. The technology to do this resembles traditional back-end-of-line interconnect engineering, which is a strong leadership area for Applied Materials.

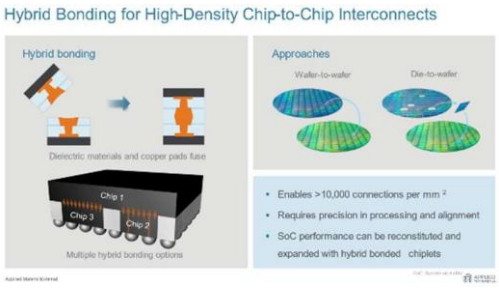

Now, I'll take you through the hybrid bonding process flow using wafer-to-wafer bonding as an example. First, we process both wafers to a particular metal interconnect level. Then we deposit a dielectric layer that has been optimized for the hybrid bonding process. We etch the dielectric to create vias that are perfectly aligned with the underlying metal interconnects. We align the vias using a copper barrier seed step and then fill the vias with copper using wet ECD. Next, we use a customized CMP process that polishes the wafers in a way that leaves the copper contact slightly below the height of the dielectric. This is called dishing and ensures that nothing impedes the bonding of the dielectric areas of the two surfaces. In hybrid bonding, the surfaces need to be exceptionally uniform, exceeding the requirements of conventional interconnect processing. Next, we clean the wafers to remove any remaining particles. We

treat the surfaces with plasma, which cleans the metal pads and forms reactive areas within the dielectric layers, which chemically activates them for the bonding process. The two wafers are precisely aligned and as they are brought together, the reactive dielectric surfaces form molecular bonds. Finally, the wafers are annealed to stabilize and cure the dielectric bonds. The heat from the annealing process expands the copper pads and causes them to fuse, which completes the electrical connections. Die-to-wafer hybrid bonding is similar except that we bond individual chips to the host wafer. Over the past several years, Applied Materials has developed a comprehensive product line that performs most of the required steps for hybrid bonding, including metrology and process control. We're also collaborating with development partners to provide our customers with all the technologies they need to develop and ramp complete end-to-end hybrid bonding solutions.

Today, I'm pleased to announce that we are working on a unique Integrated Materials Solution for hybrid bonding. It includes all the process technologies our customers need, including industry-leading pick-and-place technology from Besi, one of our key development partners. Our integrated approach addresses one of the key challenges to making hybrid bonding successful at an industry scale. Earlier, I described how we treat the dielectric materials to create reactive sites prior to bonding. These treatments are short lived and easily damaged by atmospheric chemicals and particles. Any delays between activation and bonding can weaken the integrity of the bonds and lead to failures well after the chips have been shipped into the market. Our integrated solution keeps contaminants away from the delicate surfaces. It also includes robotics, software and automation that our customers need to stage materials and optimize process sequences to minimize the time between activation and bonding, and thereby achieve the best outcomes. Our solution is now running at Applied's Advanced Packaging Development Center in Singapore where we are engaged with a number of customers.

Next, I will discuss advanced substrates. Most of us have seen the green fiberglass motherboards used to connect individually packaged chips using legacy technologies like wire bonding and flip chip. These interconnect technologies are still in use and growing, but they're not sufficient for advanced heterogeneous integration. The I/O densities are limited to around 10 per square millimeter for wire bonding and around 100 for bump. The power consumption per data transfer is too high, ranging from 10 picojoules per bit for wire bonding to 1.5 picojoules per bit for bump. But using advanced substrates, we can increase I/O density by 100 to 10,000 times and reduce power per bit by 30 to 200 times.

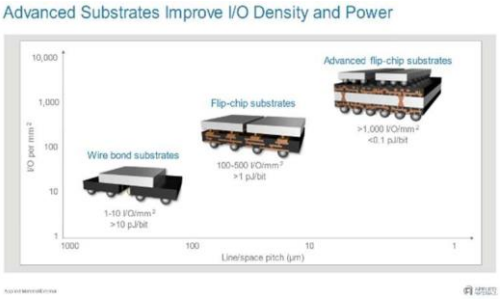

Customers use advanced substrates to integrate chips side by side in 2D fashion and to stack chips in 3D to further improve power, performance and density. Some designs combine lateral and vertical integration and are called 2.5D. 3D stacking often uses Through-Silicon Via technology which Applied pioneered years ago and the TSV market is now growing rapidly. Once populated with the key chips and supporting components such as resistors and inductors, the advanced substrates are connected to system boards.

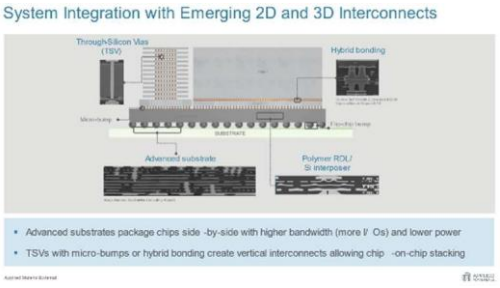

From today's earlier presentations, you can guess why silicon is becoming a popular material for advanced substrates. After decades of experience, the industry has become very adept at using silicon to create dense contacts and high-performance, low-power interconnect wiring. But there is a problem. As customers integrate dozens of chiplets to create high-performance heterogeneous processors, chip packages are growing larger. Packages as large as 10,000 square millimeters will be common over the next several years. Silicon wafers offer 70,000 square millimeters of surface area, but they are round, which means there is edge loss. To keep pace with the heterogeneous design inflection, the industry needs to go bigger.

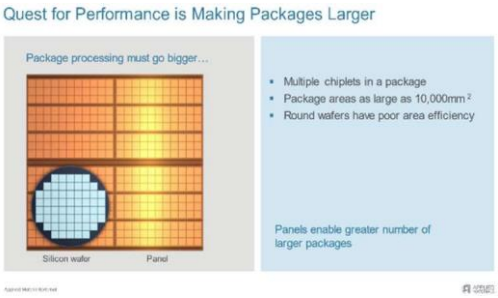

Applied Materials is working to accelerate the ecosystem for large semiconductor-grade advanced substrates based on a variety of materials. These new substrates will enable designers to use larger packages and integrate more chips at a competitive cost. In 2020, we acquired Tango Systems which has years of experience depositing wiring on substrates as large as 600 by 600 millimeters. Our Tango Systems integrate multiple process steps in a high-throughput cluster architecture and can deposit films on both sides of each panel.

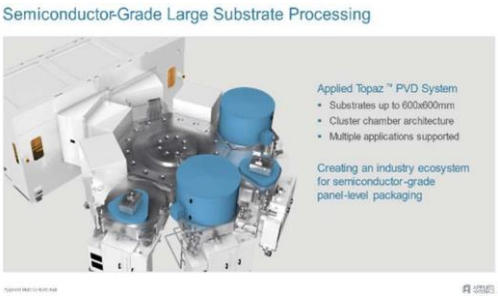

To further enable the ecosystem, Applied's Packaging Group is teaming up with our Display team, which has decades of experience engineering large substrates. Our teams are working on new methods to pattern even smaller lines, spaces and vias to create high-density contact pads and wiring. In fact, the teams are already providing our customers with ebeam test and inspection systems that have been optimized for advanced substrates. The ebeam technology can see through the layers of the materials, giving customers a non-destructive way to test designs and ensure quality. In short, Applied is working broadly with our customers and partners to accelerate the heterogeneous integration ecosystem. We will help the industry transition from system-on-a-chip to system-on-a-package and supplement 2D scaling with 3D volumetric scaling.

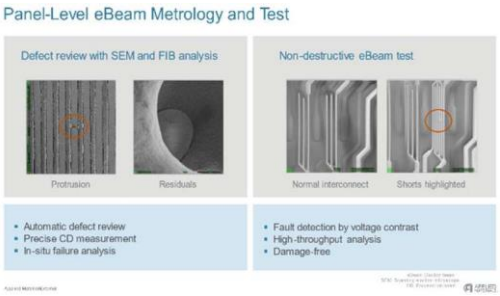

So, in summary, it's a very exciting time in the industry. For decades, PPACt was driven by 2D scaling and integrating more transistors on a single chip. Today, as Kevin described, we are continuing to enable 2D scaling with EUV by solving the resistance challenges. But as Mehul showed us, intrinsic 2D scaling is slowing and DTCO techniques like backside power distribution give us newer ways to increase logic density. I hope I've left you equally excited about heterogeneous design and integration with hybrid bonding and larger advanced substrates. These technologies are helping designers overcome the reticle limit, giving them virtually unlimited transistor budgets to improve PPACt at the chip and system level. Applied's strategy is to accelerate these innovations and ecosystems to be the PPACt enablement company for our customers. Now I'd like to hand the meeting over to Raman who will discuss our growth opportunities. Raman.

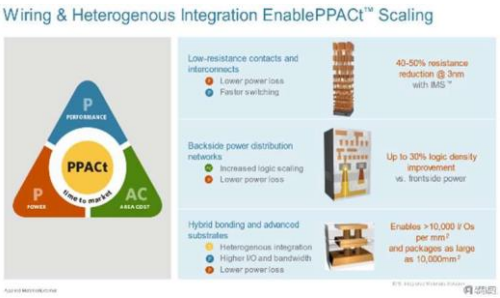

RAMAN ACHUTHARAMAN, PH.D. | Group Vice President, Semiconductor Products Group

Thank you, Sundar.

Our strategy to be the PPACt enablement company shapes how we invest our R&D dollars. It also informs where we can expect to grow faster than the WFE market through our 2024 financial model horizon and beyond. Our goal is to more than double our revenue between 2020 and 2024 in our PPACt enablement areas.

Here are the new playbook inflections we covered in today's Master Class. Kevin focused on the new ways to wire transistor contacts and interconnects, Mehul detailed the backside power inflection, and Sundar explained heterogeneous design and integration with hybrid bonding and advanced substrates.

These inflections will play out at different times. The wiring inflections are here today. In packaging, we're already seeing healthy growth in established technologies, while hybrid bonding is an emerging growth opportunity. Backside power is also an emerging inflection that'll give us new product and growth opportunities well beyond 2024.

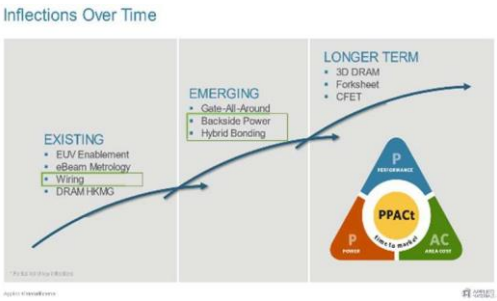

Let's now take a closer look at our growth opportunities in wiring, where we expect to grow at three times the rate of the WFE market between 2020 and 2024. As Kevin described, EUV scaling creates exponential increases in wiring resistance and our customers need new materials innovations to enable EUV scaling to continue. Beginning with the contact wiring, you can see that the contact metallization steps increase by more than 50% from 7nm, which was a partial EUV node, to 3nm. Our contact revenue opportunity per wafer start increases by more than 80% from 7nm to 3nm. Now, let's

look at the interconnect wiring. Our process steps and revenue opportunity approximately triple from 7nm to 3nm, and the number of critical interconnect layers is more than twice the number of contact layers.

Next, let's talk about packaging. As we showed at our last Investor Meeting, we expect to double our packaging revenue between 2020 and 2024.

In fact, our growth in packaging has been accelerating. From 2015 to 2020, we grew by a CAGR of around 10%. But, in 2021, we grew by over 50%, achieving over $800 million in the calendar year. Today, we believe we are on track to hit our packaging revenue goal one year early.

Packaging is another leadership business for applied where we have over 60% share of the equipment markets we serve. Our portfolio includes PVD, plating, CMP, CVD and etch, and we are number one in four of these five areas.

An important observation I'd like to make is that as the industry transitions from system-on-a-chip to heterogeneous design and integration, the packaging technology becomes more advanced and our

opportunity in wiring grows. We are bringing our leadership technologies for on-chip wiring to this new and growing area and helping customers improve power and performance as they transition from system-on-a-chip to system-in-a-package. As of 2021, most of our packaging revenue was in legacy packaging applications, including aluminum bond pad formation, bump and TSV. When hybrid bonding arrives at industry scale, we expect to expand our opportunity with Integrated Materials Solutions for hybrid bonding, along with metrology and process control. We are very excited about heterogeneous design and integration as a growth driver for Applied well beyond 2024.

Finally, I'll comment on backside power distribution. As 2D scaling slows, our customers will move the power wiring to the backside of the wafer, which frees up plenty of real estate for more transistors on the front. This is a great example of DTCO, where we increase logic density independent of lithography pitch. This is another emerging inflection where Applied has strong leadership technologies for our customers. Mehul showed us three possible implementations based on public papers and roadmaps. The first approach is the least complex and can use unit process equipment. The second offers greater logic density but is more complex to manufacture. We are developing co-optimized solutions for this approach to help our customers speed time to market. The third approach is significantly more complex. We are developing Integrated Materials Solutions as well as co-optimized solutions to enable this third approach. The approach increases logic density by around 30%, which is equivalent to two nodes of EUV scaling. These solutions are increasingly valuable to our customers, and we expect our growth opportunity to scale with the complexity and benefits.

In summary, we hope you have enjoyed today's Master Class and that it has given you insights into the roadmap challenges, technology inflections and the unique solutions we are bringing to our customers today and well into the future. And now I invite you to join us for the Q&A session, which begins in a moment.

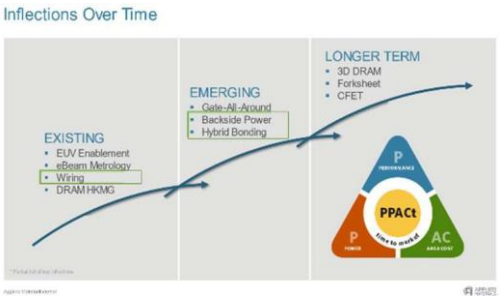
Inflections Over Time
EXISTING
EUV Enablement
eBeam Metrology
Wiring
DRAM HKMG
EMERGING
Gate-All-Around
Backside Power
Hybrid Bonding
LONGER TERM
3D DRAM
Forksheet
CFET
PPACt
P
AC