2021 Logic Master Class

PREPARED REMARKS | JUNE 16, 2021

MICHAEL SULLIVAN | Corporate Vice President, Investor Relations

[Page 3]

Hello everyone and welcome to today's Master Class. I'm Mike Sullivan, the head of Investor Relations at Applied Materials.

[Page 4]

Back in our investor day in April, we talked about our idea of holding Master Classes to give you a deep dive on technology trends that are going to be very important to the semiconductor industry, but also very important to the growth of Applied Materials. We held our Memory master class in May and that content is still available on the investor page of our website. Today, we'll do logic and then in the second half we'll cover specialty semiconductors, heterogeneous design and advanced packaging and, inspection and process control.

[Page 5]

And in fact I have an update for you today, we've decided to hold our ICAPS and Advanced Packaging class on September 8th and that will be followed by our class on Process Control as well as our AI^{xTM} platform technology on October 18th, so we hope you'll join us for those events.

[Page 6]

Now, what I'd like to do is run you through today's agenda. I'll talk to you in a moment about our logic thesis and then I'll be joined here for a fireside chat with Chidi Chidambaram from Qualcomm. Then we'll hand things over to Uday Mitra and he'll introduce the technical presenters. We're going to have three technical presentations today. One on transistors, one on interconnects and another on patterning including something new called DTCO. When we're done with that, Raman Achutharaman will be here again to summarize how the technology trends are going to be impacting our business in the years ahead. And when he's done, Uday, Raman, and I will be here to take your questions.

[Page 7]

So what are today's key messages? I hope you'll take away three things. One, the AI era is driving secular growth in logic capacity. There's going to be a lot of spending. Number two, where's that spending going to go? Our belief is that PPACt enablement with the new playbook is going to capture a growing proportion of the WFE market growth over the next several years. And number three, although the term's been around for a while, DTCO is getting a lot more attention today. And the reason why is DTCO is a great way to use materials engineering to drive continued growth in 2D scaling at existing line widths.

And now what I'd like to do is remind you of our demand thesis. We believe that we are in the early stages of a transition. In years past through the year 2018, we were in an application-centric world where humans generated most of the data.

[Page 8]

That changed: in 2018 there was a crossover. And now we're in the early part of a data-centric world where machines are generating most of the data. We think that by the year 2025 machines will generate 99% of the data in the world, and that's having a profound change on the semiconductor industry and also on all of the products downstream of us.

[Page 9]

In fact, what I'm showing you here is that we're seeing increasing amounts of silicon content being designed into all kinds of high-volume devices in order to enable this new data world. So for example, in high-end smartphones, between 2020 and 2025, we believe that silicon content is going up by over 60%. In automobiles, it's going up by more than 50% and in servers nearly 100%. And as we all know around us just looking at the world, there are shortages everywhere for components that go into these devices.

[Page 10]

And what does that mean? Well, one thing that's new, which we haven't seen in years past, is that the major customers are announcing strong multiyear capital investments in order to keep up with this demand. And what they're saying is that these are not speculative investments. These are investments they're making based on true demand that they see over the next several years.

[Page 11]

And it's not just the customers. In addition, what we're seeing are incentives being announced for domestic production by countries and regions around the world. You're seeing this in the US, the EU, South Korea, Japan, China. Everyone understands that having a healthy, vibrant economy today means having semiconductor supply. Now we don't believe that these investments necessarily significantly raise wafer fab equipment spending, but we do think it's a positive indicator. So what does this mean for Applied Materials? Now we certainly enjoy having these tailwinds at our back but growing with the market is not our plan.

[Page 12]

What we laid out the investor meeting was a plan to grow 50% faster than the market. And how do we do that? Our belief is that the PPACt playbook that we enable -- more of the growth is going to come in these areas, and that's going to give us a chance to outgrow. And it would be one thing if Applied was the only company talking about PPACt.

[Page 13]

But in fact, we're seeing lots of people in the industry talking about this playbook, talking about new themes like DTCO. So when you look at some of the conference proceedings recently, like SPIE and ISSCC, you see our major customers and their customers all talking about new materials and structures, new architectures, new packaging, all of these things. So really what you're seeing is the whole industry is beginning to converge and agree that this is the path going forward.

And that brings us to our fireside chat with Dr. Chidi Chidambaram, an IEEE fellow who serves as vice president of engineering at Qualcomm where he leads the company's process technology and foundry engineering team. Thank you for joining us today Chidi.

DR. PR CHIDI CHIDAMBARAM | Vice President of Engineering, Qualcomm

It's my pleasure Mike, thank you.

Michael Sullivan: ***Chidi, what trends will drive logic demand over the next 3-10 years?***

We at Qualcomm believe we are right in the midst of a fourth industrial revolution. The coming together of key technologies like 5G, AI-artificial intelligence, and IoT devices that are pervasive, are going to enable a combination of user experience and customized demand that's near true industrial revolutions. We are excited about the next decade.

Michael Sullivan: ***Excellent. How would you describe the state of traditional Moore's Law scaling. And while we're covering heterogeneous design and packaging in a future Master Class, what are your thoughts on these topics as well?***

I think it's a pretty much industry known secret that Moore's Law has been slowing. Clearly, if you look at the recent technologies from the foundries and IDMs, we don't get as much performance or power scaling from the new technologies. So we have to look for enabling the user experience through other means. And in that context, we look at the 2.5D and 3D integrations as real opportunities for us to engineer user experiences that are not coming from the traditional process technology itself.

Michael Sullivan: ***I see. Applied talks about enabling PPAC-t with "t" being time to market. How do you define and think about PPAC-t?***

The PPACt is an acronym that the industry typically uses. It stands for performance, power, area, cost, and time to market. We also use these benchmarks, but the one that I really like to think of as the "t" to stand for time to manufacture. Not really time to market, because in the new technologies, particularly when the scale has gone up so much, we are shipping so many million of units in the first couple of quarters of production in these technologies. The time to manufacture is a pretty critical element. It's become much more significant than even the power, performance -- and they may take some time to reach. But if you don't get enough of our production volumes in time, I think all of us pay a big price.

Michael Sullivan: ***OK, thanks. Now, let's dive into the technology, beginning with the transistor roadmap. How do you see FinFETs evolving, and what comes after FinFET?***

In the last few generations, the power scaling has been accomplished by what's known as fin reduction. We went from -- three generations ago -- four fin devices to three fin devices to two fins. Unfortunately, we can't scale beyond the one fin. So to accomplish the power reduction, we have to switch the architecture. And that's where I see things like Gate All Around coming into play. It could be a sheet, you know people call it different names, but essentially it is getting the gate on all four sides. So going from a planar to a fin, you put the gate on three sides and you got like 150 millivolts reduction, and

when you go to Gate All Around you get one more side so going from three to four will give you not quite as much voltage reduction, but 50 millivolts or so. We hope to catch the power reduction from these architectures.

Michael Sullivan: ***Quick follow up on that. In about what time frame do you think we might see gate all around coming into the marketplace?***

I think you know, '24, '23 productions. The earliest could be '23, but '24 production is reasonable, I think.

Michael Sullivan: ***We talked about the transistor; can you define what the transistor contact does and discuss some of the problems and maybe some of the solutions that you see in this area.***

The contacts are really the elements that touch down into the silicon. They are the first things that contact the silicon so everything that we can get out of this having to come through a contact. They play a very significant role in making sure whatever the juice the silicon generates, reaches the end user finally. And it's become a bigger challenge because the resistance of these contacts are limiting how much current we can get out of the silicon. And you know clearly from a design point of view, trying to find new materials and metals that can have lower resistances, and also trying to thin the liner -- that's kind of a barrier between the metal and the insulator that sits there -- are two approaches that you can think of to get better contact scaling going forward.

Michael Sullivan: ***We talked about transistors, contacts. Next, now can you define what the interconnects are, and talk about the innovations you think the industry needs with the interconnects.***

If silicon is the ground floor on a multistory building, the interconnects are the remaining 15 floors of it. It's that large, and if you laid it out on a straight line, it can run into miles. So you can imagine how, as we thin these metal lines, the increase in the resistance is really going to affect your final performance -- the frequency that you get -- and the thermal heating that these produce are all very, very significant. And in this context, I think backside power is a good, interesting solution that's going to come in to help in the future. So we need to do all these things to get the resistances down.

Michael Sullivan: ***We've been hearing about something called "buried power rail" or "backside rail." Can you describe what that is in a nutshell?***

Traditional chips, if you look at it, we put the contacts on the top -- bumps on the top -- and power them through the top. And they go through all these many layers of metal and the contact to reach the final silicon, where you're actually processing the signal. And in the process of taking the power from outside to the silicon, you lose a lot of voltage. Instead, if you can bring fat wires through the back and directly contact the silicon, you end up not losing so much voltage in accessing the silicon. I think that's capable of giving at least 15% to 20% performance improvement. It's a neat technology that we are looking forward to.

Michael Sullivan: ***Does it also reduce power as well as increase performance?***

It enables lower capacitances, so therefore you will see some dynamic power reduction as well.

Michael Sullivan: ***What about scaling. Does that make the chips smaller or does it not affect that?***

It will help scaling also because majority of the standard cell area is used up for power delivery itself. If you can efficiently bring the power through the bottom of the cell, it'll free up area to shrink the cell size itself.

Michael Sullivan: ***OK, so finally, how do you see patterning and scaling evolving? For example, at ISSCC, Dr. Mark Liu of TSMC said that with each advanced node, DTCO is playing a bigger and bigger role as compared to intrinsic scaling. What are your thoughts on this?***

The traditional scaling, if you think about a construction building and the scaffolding sitting right outside it, it has a rectangular XY grid, and if you just bring the X's together and the Y's together, you scale the rectangle and that's how traditional scaling or pitch scaling happened. But starting a few years ago, we could not rely on just the pitch scaling itself because it was not resulting in a0.7X linear and 0.5X area scaling. So we had to do other things to achieve this end result, and those are collectively called design technology co-optimizations. Actually this is very near and dear to me, because in my IEEE Fellow citation, I was cited for DTCO, as bringing those innovations to production. So that has become the larger fraction of what enables scaling going forward. Therefore it's pretty significant that we leverage all the new ideas to explore those possibilities for upcoming technologies.

Michael Sullivan: ***It sounds like the term DTCO has been around for quite a while but it's getting a lot more attention now. Why is that happening?***

The general offering from the process technology scaling itself has declined a little bit. And the main reason is, the pitch scaling has slowed significantly and since this scaling has slowed, we've started innovating on DTCO, and that's been going on for a couple of generations. But we're at a crossover point where the contribution from the pitch scaling is decreased in magnitude compared to what the DTCO can do. But even the DTCO today cannot achieve what it could five years ago. So in general both are slowing, but today the larger fraction is coming from DTCO.

Michael Sullivan: ***And in closing, what's one idea that you'd especially like to impress on us today?***

Manufacturability is a key idea that I would like the tool community and the foundry community to kind of partner together in achieving what we've realized. Because to take innovation and create a one-of-a-kind device is kind of easy, and it's being done and actually a lot of places are working on that. But what we really need is how to -- when you make many tools and multiple process steps, and a wafer that varies from center to edge -- how to get all of it together. So when I ship my cell phone, every user gets the same battery life out of it all the time. That's really manufacturability and that really needs the industry to come together.

Michael Sullivan: ***Fantastic. Well thank you so much, Chidi for joining us today.***

Thank you, Mike. I really appreciate it. I enjoyed the opportunity

Michael Sullivan: ***And now I'd like to turn it over to Dr. Uday Mitra to introduce the technical portion of the meeting.***

DR. UDAY MITRA | Vice President, Semiconductor Products Group

[Page 14]

Thank you, Chidi and Mike, and hello to everyone on the Webcast.

I'm Uday Mitra, Vice of President of Technology Roadmap and Strategies and it is a pleasure to introduce this technical part of this program. We've had an interesting time putting together the technical portion of the master class.

Our goal is to help you understand the most meaningful and exciting inflections that are developing in the advanced logic roadmap over the next 3 to 5 years and beyond.

[Page 15]

FinFET scaling for improving Power and Performance continues through materials engineering improvements in the channel, Source Drain and Gate. The major inflection in the transistor roadmap though, is GAA which is totally enabled by materials, and processes, and Applied Materials, with our products, IMS strategies and deep customer engagements is extremely well positioned to benefit from this inflection.

[Page 16]

In a moment you'll meet Dr. Mike Chudzik who is a Vice President of Technology and one of our top materials and device experts. Mike joined Applied 7 years ago and was at IBM for nearly 14 years before that. He'll talk about the transistor roadmap, from FinFET scaling to Gate All Around.

[Page 15]

Apart from the transistors there are also major changes in the Back End as the Interconnects are a critical scaling barrier for logic performance, with a big challenge being via and contact resistance – driving the need for new ways of building devices, new materials and new processes such as Selective Metal Fills and Barriers.

[Page 16]

Addressing these will be Dr. Mehul Naik who is a Managing Director and Principal Member of our Technical Staff. Mehul has been at Applied for 26 years and holds over 70 U.S. patents. He'll help you understand why the transistor contacts and interconnects need to scale along with the transistors and how we're going to enable that to happen.

[Page 15]

Finally, area reduction in a logic standard cell which is an essentially group of transistor and interconnect structures that provides a Boolean logic function today is no longer done solely by pitch scaling alone. In fact, over the nodes DTCO or Design and Materials Technology Co-optimization is playing a much more predominant role, examples are Single Diffusion Break, Contact overactive gate and in the future Buried or Backside Power Rails.

[Page 16]

To talk about this as well as how Materials Engineering is enabling traditional pitch scaling by enhancing EUV patterning as well as Multi-patterning, I would like to introduce Regina Freed who is our Managing Director of Patterning Technology. Regina has been in the industry more than 20 years, and today she'll explain DTCO along with advanced patterning applications that enable customers to scale faster and reduce area and cost.

After Raman's presentation later today, I look forward to joining Raman and Mike for the Q&A session.

And now Mike, it's time to begin the transistor session.

DR. MICHAEL CHUDZIK | Vice President, Technology

[Page 17]

Thank you, Uday.

My name is Mike Chudzik, I manage an integrated process module solutions group here at Applied.

I'll now take you through the transistor challenges and roadmap as the industry continues to push the boundaries of power, performance and area-cost scaling.

[Page 18]

The CPU is heart of the smartphone. It determines our user experience, and performance and battery life are paramount. So chip manufacturers are all pushing the boundaries of advanced transistor design and process technology. Drilling down into the process of the latest iPhone processor, we can see it's based on 5nm node technology, with around 12 billion transistors and 24 billion transistor contacts.

The flexible silicon design includes over 7 threshold voltages to enable both high performance and long-life mobile computing.

[Page 19]

This 5nm technology is the culmination of a long string of process innovations that have of occurred over many decades. We are now in the advanced FinFET era and looking to a new device architecture called Gate All Around.

At Applied, we help enable these innovations with our leadership in equipment and process technology. Key innovations to enable the transistor roadmap include epitaxy, ion implant, metal gate formation, rapid thermal processing, Chemical Mechanical Planarization or CMP, and a special kind of etch called selective materials removal.

We can combine these leadership technologies to enable integrated materials solutions where adjacent steps can be co-optimized for the best possible performance and power results.

[Page 20]

The transistor operates as a switch, and to enable the best performance, we focus primarily on reducing the switching delay.

Minimizing delay is a function of maximizing drive current and reducing capacitance and resistance. In the FinFET transistor, speed can be optimized by tuning various physical parameters including the fin height, the gate length of the channel, the mobility of electrons running across the channel, the threshold voltage used in switching, and the thickness of the gate oxide that helps control the on-off state of the switch.

Typical node to node drive current increases are on the order of 20 to 30%. We reduce electrical resistance by implanting to achieve higher activated ion dopants. Another critical performance knob is design variability because performance is gated by the slowest transistors in a circuit. By tightening the distribution to reduce variability we enable faster circuits. This is true even if the average speed of the transistors decreases. In this example, the red design is faster despite lower average speeds.

[Page 21]

As mentioned on the previous slide, the transistor's physical attributes determine its performance. FinFET construction can be broken down into three main modules: one, the channel and shallow trench isolation -- two, the high-K metal gate module -- and three, the transistor source drain resistance module. Listed beside these three modules are the key scaling levers that enable higher performance:

In the channel module we optimize channel isolation using taller and thinner fins. We optimize the channel for high mobility. In the High-k metal gate module, we optimize tunnel oxide scaling and tune the threshold voltages of the metal gate.

In the transistor source drain external resistance module, we optimize placement of the embedded source-drain and tune the NMOS source drain using a Silicon Arsenide layer.

[Page 22]

Now we'll talk about some of the challenging solutions to pushing FinFET designs even further. In the Channel and Shallow Trench Isolation module, we have been increasing fin height and reducing fin width over several technology nodes. These fins are becoming more fragile and tend to bend during the manufacturing process.

Applied is mitigating the bending with new approaches in trench isolation involving precision materials engineering of the oxide. Fins made with Silicon Germanium enable higher drive current but are more susceptible to oxidation. Applied has solutions for this including Silicon Germanium epitaxy, nitridation and Shallow trench isolation oxidation.

[Page 23]

We'll take a closer look at this. Taller, narrower fins are more susceptible to bending due to strains caused by the isolation oxide that needs to be placed between the fins

This bending causes process variability and counteractive strain which degrades mobility and impacts threshold voltages. Applied has developed co-optimized materials engineering solutions for fin bending

including flowable oxides, ion implantation and anneals steps, which are monitored using our PROVision® eBeam metrology and inspection technology.

We are able to mitigate the issue to improve threshold voltage variability and increase drive current upwards of 5%.

[Page 24]

Now we'll take a closer look at the High-k and Metal gate module which is the heart of the transistor.

These stacks are highly complex and can contain upwards of 7 layers. These include the Interface and high-k layers and the metal gate layers. Interface and high-k scaling is critical to gate oxide reduction which boosts the drive current. The metal gate is tuned to ensure the transistor has the correct work function which determines the threshold voltage.

[Page 25]

We can see that since 14nm, the interface layer and high-k haven't scaled at the same rate as the other physical parameters that lead to higher drive current. The physical dimensions of the interface layer at 8 angstroms and Hafnium oxide at 15 angstroms illustrate how difficult this problem is to solve.

Applied has a new integrated approach that enables scaling to continue. We're combining process steps in vacuum to engineer interfaces and tune the process. This unique integration prevents harmful exposure to the atmosphere which creates moisture and prevents scaling.

We've demonstrated a new integrated gate stack that enables an optimized interface and improves drive current by 8 to 10%. In fact we are presenting this new technology at the VLSI technology conference this week.

[Page 26]

Designers can use this ability to tune the threshold voltages over a greater range with a higher degree of accuracy to tune their chips for high performance or mobile computing using a single process flow. There are two main ways to tune the threshold voltages using process technology.

The first is thickness modulation where the metal gate thickness and materials are tuned using deposition and etch to provide specific work functions. The second is dipole engineering which uses dopants to change the work function of metals to tune the threshold voltages. This approach relies on ALD, anneals and selective etch.

Applied is enabling the threshold voltage tuning with our leadership portfolio of ALD technology for work function metals, in-vacuum processing, and anneals.

[Page 27]

The third major module we'll take a look at is the transistor source/drain resistance module. The major contributors to the resistance inside the transistor is the interfacial resistance between the metal contact and Silicon transistor, and the external resistance of the source and drain regions.

Generational node scaling has reduced the contact area by roughly 25% per node. This reduced area drives up the resistance. The interfacial resistance is scaled by co-optimizing epitaxy, implant, metals

and anneals. The external source/drain resistance is scaled by the use of sculpting etches, epitaxy, implant and anneals.

[Page 28]

Applied is optimizing the external resistance with a new way to laterally etch beneath the spacers in a conformal fashion. This brings the embedded source-drain stressors closer to the channel to provide higher drive current.

In addition to this conformal etch process, Applied has developed a novel selective Silicon Arsenide epitaxial layer. The Silicon Arsenide lowers the resistance and allows the Silicon Phosphorous stressor film to be deposited closer to the channel which improves drive current up to 8%.

[Page 29]

The finFET architecture is being pushed to the limit. Narrowing the fins helps produce shorter gates. Fin width is defined by traditional lithography and etch. Controlling the width is becoming harder, and the variability in threshold voltages impacts performance. Gate All Around gives the industry a new way to solve the fin width variability. We will transition away from using lithography and etch to define the fin width and instead use epitaxy which enables very precise thickness control.

In the Gate All Around approach, the finfets are essentially turned sideways. The channels are formed with a superlattice of Silicon and Silicon Germanium. The Silicon Germanium is eventually removed, leaving behind a stack of silicon channels. Let's take a closer look at the Gate All Around architecture to understand its value, construction and unique requirements.

[Page 30]

From a performance perspective, Gate All Around enables lower variability while enabling gate length scaling which increases drive current 10-15% and reduces power consumption.

As I mentioned, the process borrows heavily from the finfet flow. But the channel is made with a superlattice of Silicon and Silicon Germanium, and the silicon germanium is selectively removed at a later step to reveal the horizontal silicon channels. The key to well-formed channels are sharp interfaces and highly selective etches.

Forming the embedded source drains requires better etch techniques and growth control. Gate all around devices require a new feature called an inner spacer that is used to reduce capacitance. The spacer requires very controlled, selective etching and gap fill processes.

The high-k metal gate wraps completely around the channel. Forming the uniform gate requires super-conformal interface, high-K and metal ALD processes. Additionally, we need new metrology steps to monitor and measure the formation of the super lattice.

[Page 31]

Now we'll take a closer look at the channel module. Two new process steps work in concert to create the well-defined channel "slabs" or channel. The process starts with a Silicon and Silicon Germanium alternating super lattice that is deposited with epitaxy. The layer thickness ends up defining the channel

width. In the channel reveal process, quickly transitioning from the Silicon Germanium to pure Silicon is critical to forming sharp channels.

Applied has demonstrated the ability to transition greater than 30% Germanium per nanometer of thickness. We have also developed a subsequent selective Silicon Germanium etch process which reveals the Silicon slabs or channels with selectivity of greater than 500-to-1.

[Page 32]

Another unique module is used to form an inner spacer that isolates the high-k metal gate from the embedded source drain spacers. The spacer is formed by slightly notching the Silicon Germanium inwards before the reveal process which allows for a low-k spacer to be deposited using ALD.

Applied's leadership in selective etch enables this flow. Finally, the high-k metal gate module presents unique challenges due to the shape of the channel. Because the high-k metal gate wraps around the four sides of the silicon channel, the interfacial layer, high-K, and metal gate need to be conformal on all sides, and particularly underneath the channels.

Applied's technologies for these structures and the new integrated solutions I talked about earlier, provide highly conformal processes that evenly coat the undersides. We can see the uniformity of the elemental concentrations around the silicon channels.

Thank you for your time today, and now I'll hand it over to Mehul to discuss the contact and interconnect modules.

DR. MEHUL NAIK | Managing Director, Semiconductor Products Group

[Page 33]

Thank you, Mike. My name is Mehul Naik, and I am responsible for the Logic Module group in the semiconductor products team.

To continue scaling logic, we need to do more than shrink the transistors. We also need major innovations in the transistor contact and chip interconnect modules. These will lower resistance and capacitance to maximize drive current, increase performance and reduce power.

Today, we will review some of the key contact and interconnect scaling challenges the industry is facing and how Applied Materials is helping solve them. Let me start with some context.

[Page 34]

As Mike showed, a leading-edge smartphone processor has over 12 billion transistors and over 24 billion via and contact connections. It also has over 15 layers of metals of varying dimensions, with the ones closest to the transistors being the narrowest. So, with all of this progress, why is there so much discussion in the industry about a scaling wall?

[Page 35]

When transistor dimensions scale, the devices get faster. It's like having a free lunch.

But this is not the case for metals. As dimensions scale, metal resistance increases, and power-performance scaling is negatively impacted. Left unchecked, this issue can offset the benefits of transistor improvements. The contacts that connect the transistor to the outside world also contribute to transistor resistance and influence the transistor's drive current.

The interconnects are the current-carrying lines that connect all of the individual devices within the chip. Since the interconnects route the signals and distribute the power, they control the RC delay -- or the speed of the circuit -- as well as the power consumption. As dimensions shrink, contact resistance increases by a greater than factor of 4. Because it can contribute 10% of transistor resistance, the contact also needs to be improved each node.

The interconnect is susceptible to a 10X increase in via resistance. Interconnects consume close to one-third of device power, and they account for more than 75% of delay. So improving interconnect resistance is the best way to improve device performance.

[Page 36]

The key to improving contact and interconnect resistance is developing new metal fill processes that minimize or entirely eliminate high resistivity liner barriers. Thinning or eliminating the liner barriers reduces interface resistance and maximizes the volume available for metal conductors. Now that we've worked through the fundamental ideas, let's look at the details in reducing contact resistance to improve drive current.

[Page 37]

Through the 10nm node, the transistor contact and its middle-of-line via were filled with tungsten that was deposited with a liner that provided adhesion between the via and the insulating dielectric. If we tried to use the same approach at the 3nm node, we would have a 4X increase in resistance and negate the benefits of transistor scaling.

Applied delivered low-resistance contact engineering solutions which shifted the resistance curve. We were able to maintain the 10nm resistance characteristics at the 5nm node. This 2X resistance reduction was achieved by replacing tungsten with Cobalt as the contact metal, and by developing the industry's first Selective Tungsten fill process for the middle of line via. We plan to further lower resistance by eliminating the liner that's been needed for the contact metal fill.

[Page 38]

Applied was first to market with a robust selective deposition process for metal fill. This new technique eliminates the need for volume consuming liners, and this enables scaling to continue. It also reduces via resistance by more than 40%. At the device level, this can reduce power consumption by close to 6%.

This technique is integrated in an Integrated Materials Solution platform where multiple surface preparation and deposition technologies are integrated in one system under high vacuum.

[Page 39]

Now let's discuss interconnect performance.

As we learned earlier, shrinking the dimensions increases resistance which in turn increases signal delays and power consumption. Conceptually, we want to emulate the contact module improvements by eliminating high resistivity interfaces and maximizing the metal volume. But the implementation is very different compared to what we saw in the contact area. One additional consideration in the interconnect module is that the low-k dielectric films that insulate the metal lines need to get stronger with scaling to eliminate patterning issues.

A key focus here is creating new low-K films with better mechanical properties.

[Page 40]

A typical Copper interconnect structure is composed of three different films. A Tantalum Nitride barrier is deposited between the conductor and the dielectric to maintain good adhesion. A Cobalt liner adheres to the Tantalum Nitride barrier and facilitates Copper fill. Finally, Copper is deposited into the remaining volume. So what is needed to further scale the interconnect?

The Tantalum Nitride Copper interface is the highest resistivity interface and controls the via's resistance. The best way to reduce via resistance would be to eliminate this interface altogether. But this could only be accomplished by developing a selective barrier process. Another increasing challenge is filling these ever-smaller structures with Copper. We would need an entirely new Copper reflow technology.

[Page 41]

Today, we are really excited to announce our new Endura® Copper Barrier Seed IMS™ solution. This is an Integrated Materials Solution with new selective barrier and copper reflow technologies. This truly remarkable system combines advanced technologies for ALD, PVD, CVD, interface engineering, surface treatment, reflow, and on-board metrology, all integrated in vacuum.

This solution reduces the via resistance by close to 50% and provides a path to scale Copper to the 3nm node.

Here is an animation that will run for about a minute and a half and demonstrate how this remarkable solution works. Let's now look at how Integrated Materials Solutions are changing the way interconnects are built.

[Page 42]

In the past, we had lots of space to work with and could use conventional techniques like liners and barriers that provided good adhesion. The larger spaces were significantly easier to fill. As we transition to selective processes to enable scaling to continue, we need to carefully engineer the interfaces between materials.

We need to deploy new materials; we need to combine multiple technologies in vacuum to protect materials from the environment; and we need to prepare surfaces to make sure they are pristine so that selective processes work. It is now common for us to integrate 7 or more process steps including

surface preparation, CVD, PVD and ALD in a single system. On the right-hand side of the chart, you can see the results.

With IMS™, we've been able to significantly lower the resistance to scale dimensions which match the resistance levels of the previous node. This is a huge PPACt improvement because each chip has over one billion vias per layer and at least five critical layers. Lowering the resistance in this way can have a huge positive impact on power consumption.

[Page 43]

Now let's talk about dielectric isolation. The dielectric constant of these films determines the capacitance of the circuit. Therefore, it impacts both RC delay and power consumption. We want to lower the dielectric constant to simultaneously increase performance and reduce power.

A typical interconnect stack can have as many as 15 metal levels, and 3 different dielectric films can be used. An oxide with a κ value of 4.0 may be used at the thicker upper levels. A porous film with a κ value of 2.6 may be used in the intermediate levels, and a dense film with a κ value of 3.0 may be used at the tightest and lowest metal levels.

A key challenge to continued scaling is improving the mechanical strength of the low-k film used at the tightest pitch, thereby avoiding pattern collapse and at the same time making sure that the power consumption remains the same.

Applied has developed a new generation of its Black Diamond® family of CVD films that increases mechanical strength by 40% at the same dielectric constant. These new films are currently being proven at 3nm node and expected to extend to 2nm node.

[Page 44]

To scale beyond the 2nm node, we may need major breakthroughs in interconnect design. First, let's talk about power delivery. Each logic chip is made up of standard cells, which are groups of transistors and interconnect structures that provide specific logical functions. Each cell needs space for the signal lines and a power rail.

The power rail delivers power from an external power supply to the transistor. The power rail is typically 3 times larger than the smallest interconnect wire. Therefore, it is a major impediment to cell scaling. In addition, the power delivery network must be routed through connections at each of the metal levels of the chip, and this can easily be up to 12+ levels.

At each level, there is a large drop in voltage caused by the resistance of the metal. Designers can cope with total voltage losses of around 10%. But metal resistance is increasing with each node shrink. It is projected that without a new architecture, the power distribution network could consume 50% of the incoming supply voltage. To enable scaling to continue, a new architecture is being proposed.

[Page 45]

The new "buried power rail" architecture moves the power rail into the back side of the silicon wafer, beneath the transistors. This allows the transistor cell areas to be scaled by up to 33%. It also allows the signal lines to remain larger and have lower resistance. The buried power rail also eliminates the

voltage drops that occur in conventional routing schemes. It is projected that the voltage drops can be reduced by a factor of 7. Applied Materials plans to help drive this interconnect inflection with our expertise in metals, isolating dielectrics, etch and CMP processes.

In closing, the PPACt roadmap in logic needs concurrent innovations in transistors, contacts and interconnects. The conventional approaches are being pushed to the limit, but we already see new architectures that will be enabled by new materials and materials engineering techniques. Many of these are in the sweet spot of what Applied does in unit processes, in integrated solutions, and with the help of eBeam metrology.

I believe we will see tremendous improvements in PPACt spanning high performance computing to ultra-mobility.

Thank you for joining us today, and now it's over to you, Regina!

REGINA FREED | Managing Director, Semiconductor Products Group

[Page 46]

Thank you, Mehul.

My name is Regina Freed and I am Managing Director of Patterning Technology at Applied Materials.

As you heard from Mehul and Mike, today's device features are already so small that as we scale them further to reduce area and cost, we risk negatively impacting power and performance.

Applied Materials works closely with our customers to anticipate and understand these unwanted impacts, so we can develop new materials and materials engineering solutions like the ones shared by Mike and Mehul.

In addition, we've developed a new suite of materials and materials shaping solutions that enable creative ways to simultaneously scale area-cost while improving power and performance.

[Page 47]

Let's look at what determines the size of a logic device. Basically, a logic device is a set of individual logic cells. Each of these cells has several gates that are connected to one another by metal lines. In the vertical direction, we have the gates, which are separated from one another by a distance we call the Contacted Gate Pitch. In the horizontal direction we have the metal wires that connect the gates to one another, separated by a distance we call the Metal Pitch. By multiplying these two pitches together we get a good measure of the area occupied by a single unit cell. Traditionally, the industry has focused on scaling the cells by making the gates and wires thinner and bringing them closer together.

We call this pitch scaling or intrinsic scaling, and it has given us incredible gains in logic scaling. A new method is gaining attention and will provide an increasing proportion of the scaling gains in the future. It's called design technology co-optimization or DTCO. Since logic cell designs are relatively complicated and have many functions to perform, optimizing their layouts opens up an additional

degree of freedom for scaling. What if we could re-arrange the elements -- without making them any smaller -- to reduce the overall footprint?

A good analogy is building a house. When lot size is limited, instead of shrinking the bedrooms to make space for an office or a game room, we can add a second story. Or a cellar. To do this, we need some engineering. We'll add supports for structural integrity. And use excavating equipment to dig out space for the cellar.

DTCO is similar. We can move an element, like a transistor contact, from the side of the device to on top of the active area. This is called contact over active gate, and it allows us to scale the cell area without pitch scaling. We see an increasing set of opportunities to use DTCO concepts and materials engineering to continue to 2D scaling, and I'll give you some examples later on. But now let's shift gears and take a look at the opportunities and challenges facing intrinsic scaling today.

[Page 48]

EUV is already in production, and going forward, we can further shrink dimensions by combining EUV with multi-patterning techniques. Extremely small features can be patterned this way, if we can solve for EUV variability.

I'll explain. We want the edges of features to be straight and smooth. But in reality, there is roughness and non-uniformity at the edges of every feature. In the past, this wasn't a big concern because the features were much bigger than the edge roughness. But as we continue to scale with EUV, the features and the edge roughness are becoming more equivalent.

There is a trade-off between litho resolution and Line Edge Roughness, and we add non-uniformity when we increase the number of multi-patterning steps. The consequence is that variability is sharply increasing as a percentage of feature sizes. This results in patterning defects. For example, in some places, the width of a metal line becomes so small that we have a pinch off, which creates an open circuit.

In other places, two neighboring can become wider than intended, and they merge into one another, creating a short circuit. Another name for this variation is stochastics. When you hear us talking about "stochastic" defects, we're talking about the variations that result in these electrical opens and shorts.

[Page 49]

Fortunately, we can use materials engineering to reduce these defects. Traditionally, spin-on dielectrics and furnace steps have been used to transfer the litho pattern into the device layer. We are replacing the spin-on dielectric with a high-quality CVD material that is co-optimized with our Sym3® etcher.

We integrate the deposition technology into our etch chambers. We selectively deposit the thin CVD material on top of rough features, tuning the deposition to deposit more on wide openings and less on small openings, thereby correcting the differences between adjacent features. After the deposition, we use a specially tuned etch mode that etches small features faster than large features, once again reducing the differences. So by co-optimizing the CVD with our most advanced etch products, we can smooth the lines and eliminate many of the stochastic defects. In addition, we're taking advantage of our eBeam metrology capability.

Optical defect inspection is fast at scanning entire wafers. But to find stochastic defects, we need to measure size variations in tiny features. Our eBeam technology has the resolution to measure this variation and is the fastest of its kind. Using smart sampling, we can get meaningful patterning insights on several wafers within an hour.

[Page 50]

In summary, because variability is becoming the key scaling limiter and causes roughness on the sidewalls as well as opens and shorts, we have worked with our customers to optimize the patterning process for variability reduction. The result is a 50% reduction in local variation of the size of the features, what we call LCDU. A 30% reduction in LER, the roughness of the sidewall of a device feature. And a very significant reduction in defects, enabling continued scaling with good device yield. Next let's talk about EUV multi-patterning.

[Page 51]

In the past, multi-patterning used extra layers of spin-on patterning films to compensate for poor etch selectivity which rounded the spacers used in multi-patterning. By using the higher-quality, conformal CVD films, we can create square spacers that do a better job of doubling the EUV patterns and producing straight features.

DRAM customers are already using our new CVD material together with our Sym3® etcher to create better spacers and reduce variability. Now we are applying the same technology to multi-patterning in logic with EUV. Our EUV multi-patterning process reduces the number of patterning steps by 30%, and this helps to reduce patterning time, cost, variability and defects.

[Page 52]

Looking into the future, we will see a combination of intrinsic scaling and DTCO, with DTCO providing more of the incremental scaling benefit. I'll discuss one example called Single Diffusion break, already adopted into production, where one barrier is used instead of two to isolate individual transistors.

Replacing a double diffusion break with a single break is an obvious way to reduce cell size. This is where materials engineering comes in. At today's small scale, isolating transistors with a single feature is only possible using a special material with high-quality dielectric properties that can be deposited at high aspect ratios with no seams and gaps. The size of the feature we will need to fill has an opening less than one thousandth of a diameter of a human hair. Applied has developed the Eterna® CVD film and co-optimized it with our CMP technology. The single diffusion break improves horizontal scaling by about 15%.

[Page 53]

Do you remember the buried power rail that Mehul discussed earlier? This is another great example of DTCO. Simple pitch scaling has a negative impact on performance of these metals – yet we must reduce area and cost to scale. Materials engineering will be key to the roadmap, enabling our customers to improve power and performance while reducing area and cost. Moving the power rail from inside the logic cell to the silicon underneath enables 2D scaling of over 20%.

[Page 54]

In summary, advanced logic will continue to scale into the future, both using pitch scaling and DTCO. The contribution of DTCO to scaling is increasing. Materials engineering techniques are going to be vital to making transistors, contacts and interconnects work at a smaller scale. Materials engineering has evolved from a handful of materials, blanket deposition and line-of-sight etching to conformal deposition and selective gapfill and etching. Applied Materials will make greater use of co-optimization, integrated materials solutions and eBeam metrology and inspection to enable new devices and 3D structures.

We are already working closely with our customers to accelerate the next several logic nodes and provide chip designers with simultaneous improvements in power, performance, area, and cost.

Now Raman, I'll hand it over to you.

DR. RAMAN ACHUTHARAMAN | Group Vice President, Semiconductor Products Group

[Page 55]

Thank you, Regina, and thank you Mike and Mehul!

It's great to meet with you again in the second of our four-part master-class series. I hope that after today's presentations, you recognize how materials engineering and the New Playbook are becoming more important to enabling the roadmap of our customers.

[Page 56]

What I would like to do next is to help connect these technology trends to our business growth plan. As we showed you at the Investor Meeting in April, we feel very good about the secular growth drivers in computing and semiconductors. In our base case, we see these drivers fueling equipment investments at an 8-to-9% CAGR.

Applied's opportunities are greater than this and we believe our unique and broad capabilities will enable us to outgrow the market by around 50%. I hope these master classes are answering questions about how we can do this. As a quick reminder, we held the memory master class on May 5th.

[Page 57]

We described how in 2020 through 2024, Applied can generate $1 billion in cumulative revenue from DRAM capacitor scaling. And how we can generate $2 billion in cumulative revenue from DRAM periphery scaling. As Mike, Mehul and Regina described, similarly we see large growth opportunities in foundry-logic as well. Let me start with a quick overview of the market.

[Page 58]

In the past 10 and 20 years on average, foundry-logic spending has been over 55% of equipment spending. So high exposure to foundry-logic spending is valuable. Enabling the foundry-logic roadmap with leadership technology -- and accelerating PPACt -- is even more valuable. We are already seeing

this. As Gary said on our earnings call in May, “We are expecting CMP, Epi, Thermal and Implant to all grow more than 50% this year.”

[Page 59]

Over the past 10 years on average, specialty and trailing node foundry-logic has driven about a third of the foundry-logic spending. We expect this overall ratio to hold going forward. There is a lot of business opportunity for Applied in those specialty nodes and markets.

[Page 60]

This is why around two years ago, we formed the ICAPS business group focused on enabling our customers with our broad portfolio of technologies. As Dan said on the May earnings call, our revenue opportunity in ICAPS is already much greater than $3 billion this year. We’ll discuss the ICAPS business with you in September.

[Page 61]

Another topic that’s been gaining attention recently is advanced packaging. Companies are showing how new techniques like hybrid bonding allow silicon to be combined in new ways to simultaneously improve power, performance, area, and cost. Applied has a key role to play in these efforts, and we will do over $800 million in packaging revenue this year, which is up around 60% from last year. We’ll also detail our packaging opportunity for you in September.

[Page 62]

And today we focused on the two-thirds of the foundry-logic market that comes from investments in the three most advanced nodes, which today are 3nm, 5nm and 7nm technologies.

[Page 63]

Foundry-logic spending is much higher today than it was just a few years ago. There are two big reasons for this. First, the new demand drivers associated with the Big Data and AI era. The data explosion is driving up silicon content in new and existing high-volume devices. Second, the technology is becoming more complicated. It takes more engineering to scale these devices for continued area-cost reduction -- and at the same time to improve power and performance.

Between 2016 and 2024, we see the number of process steps doubling and wafer fab equipment spending increasing by 2.7X. Many of these new steps are materials engineering steps from Applied. This line of sight to new applications explains why we expect to outgrow the market. I’ll highlight out some of the specific growth opportunities from the presentations you saw earlier.

[Page 64]

Mike described to you the Gate All Around transistor inflection. It’s great for PPAC, simultaneously improving power by 25-30% and performance by 10-15%. Much of this benefit depends on properly engineering the fin width. In finFET, this is controlled by lithography and etch. In Gate All Around, it will be defined by transitioning to epitaxy and selective materials removal. These are leadership businesses for Applied.

We plan to co-optimize our technologies and deliver integrated materials solutions for Gate All Around inflection. We also plan to use eBeam metrology and inspection along with our Applied AIx™ platform to give customers the actionable insights they need to bring Gate All Around to market faster. As compared to FinFET, Gate All Around will provide an incremental TAM opportunity of $1 billion in our leadership areas for 100K wafer starts per month.

[Page 65]

In Mehul's presentation today, you learned how scaling the transistor is meaningless unless you can scale the contacts and interconnects with new materials and technologies. We officially introduced a new Integrated Materials Solution called the Endura® Copper Barrier IMS™ system which has selective barrier deposition and reflow copper technologies.

These are critical technologies that enable scaling to 3nm and beyond. At the investor meeting we showed you how our interconnect revenue opportunity is tripling from 20nm to 3nm. In fact, our interconnect opportunity grows by over 30% in the transition from 5nm to 3nm.

[Page 66]

Patterning has been a big growth opportunity for us in recent years. In 2012, our share in patterning process steps was in the single digits. Since then, we put a lot of R&D to work in developing new CVD films and our Sym3® etch product, co-optimizing them and improving our CMP products to accelerate patterning breakthroughs for customers.

Regina showed you some amazing technology where we co-optimize CVD and etch to fix EUV stochastic errors and allow scaling to continue. We've great momentum in this area and have doubled our patterning revenue over the past 4 years. We expect to double again by 2024. Our cumulative patterning revenue opportunity over the next 4 years is greater than $3.5 billion.

[Page 67]

Regina also discussed how our eBeam metrology and inspection technologies are enabling these new patterning solutions to work. This is driving our process control revenue to grow over 50% this year, which is twice the growth we had last year and eBeam revenue will be around $900 million. I hope today's class helped you see why we have so much momentum and customer pull for our unique capabilities.

We'll hold another master class in October to give you more insights into our Applied AIx™ and process control technologies.

So thank you for joining us today. We hope you are excited about the roadmap and can see the important role our materials engineering will play in accelerating PPACt roadmap for our customers.

Now, Uday, Mike and I would like to address any of the questions you have about today's presentations.

Mike's let's begin the Q&A.