Bob Baker; Sr. VP. Technology and Manufacturing Group

Sean Maloney: – integration. We've come a long way from the 4004. We will be integrating new instructions, taking it further over the next two to three years. Dadi will be talking about that through the Sandy Bridge generation. But we've come a long way. We think the technology in the next 12 to 18 months is very much post-recessionary technology. We believe collectively we're going to have a really tremendous next 12 months, as Paul said.

Thanks very much for coming along to IDF. We're looking forward to a great next three days. Thank you.

[Applause]

Female Voice: Ladies and gentlemen, please welcome Bob Baker.

[Applause]

Bob Baker: Thank you. It's really exciting to be here this afternoon to be able to share with you where we see semiconductor technology going in the future, and the promise of continuing Moore's Law out in the future. I also want to touch on and provide you a few more details on 32-nanometer and 22-nanometer technologies that Paul talked about this morning. And lastly, I want to talk to you about some of the changes and transformations that are required for us to be successful in system on a chip and other technologies to fulfill all of the opportunities that we have with our products in the future.

The first part is our continuing relentless pursuit of Moore's Law. And just as a refresher, I want to touch on: Is Moore's Law ending? It

comes up -- I see it repeated in the popular press, and people question it. And I think let's first start why we like it. It delivers more transistors that we can turn into performance and higher-powered products, or it allows us to cost-reduce and deliver things more economically.

The price per transistors continue to drop as we've been able to shrink transistors, fit more on, incorporate more functions, and value what Sean talked about for integration, and the benefits of that in enabling new products. More transistors, more cost effective.

But I think we may want to take another look at this. One way to look at it is the growth in mips per watts that we've had. From 1970, the growth in computing power has continued to climb. And I think if we put that in a practical manner, a way that we can all relate to it is in gas mileage.

The price of oil has gone up. If the auto industry had developed their technologies and been able to have Moore's Law, we'd be getting 100,000 miles per gallon. Quite an extraordinary feat. But it shows you relative expansion of what we've been able to do in semiconductors, and bringing those to market. Really amazing what Moore's Law can do.

This is a pictograph of the two different options that you have. One is continuing to shrink through the nodes the dye size, getting the same amount of circuitry with a node, with a 70 percent linear shrink, giving you about half the area, delivering the same performance.

The other option is to add architectural innovations, integrating other components or adding transistors for cash to produce more powerful microprocessors. So you have the option to design where your operating point and where your business is going to need to be. So let's see what that looks like on 45 nanometer, or High-K/Metal Gate transistor technology, and what we've been able to do.

Five different 45-nanometer microprocessors. From the tiny little atom that Sean showed you that you throw 1,000 of them in a little glass jar, all the way up to an eight-core Nehalem EX running 16 threads in a 24 megabyte cache, a tremendous range of computing power all enabled for its operating point by Moore's Law, and shrinking the functionality of our products for the right market and the right feature set.

So this is what we've done. This morning Paul talked about it, and we've done it in a big way. Two hundred million units of CPUs shipped with 45-nanometer with high k metal gate transistors, giving better performance, lower power, and better features for the marketplace.

Where will this keep going? Where will it take us? This is a good example of what we've been able to do. And this is just a short two-year transition of the ramp we've had on 45 nanometer CPUs. But will Moore's Law continue? And how do we continue to create this pipeline of technology that feeds our industry and feeds the different types of products and capabilities that we want to bring to market?

This is an example of our technology pipeline all the way back to 65 nanometers when we introduced it in 2005, through our 32 nanometer that's ramping in production today, and out into the future. You can see options for 22 nanometer that we'll talk about later. But there are advances needed to get to 15 nanometer and sub-15 nanometer by 2015 and beyond.

There are a range of different options. People talk about carbon nanotubes and how you could use that to change the wiring on a circuit, and how do you take advantage of that? There are lots of different techniques and capabilities that are in research today.

But it really requires two things. One, it requires a continued dedicated investment in research and development. You have to invest the dollars to create the innovation. The other is investigation, reviewing alternatives, and then selecting the right options to incorporate into innovative technologies. It's moved from just a technology where you're scaling lithography and the materials are a minor part of the process scaling.

It's really turned into a materials game, first with copper and low-k inner dielectrics to advance wiring to keep pace of the transistors, and then a few years ago with strained silicon. How do you get more performance out of the transistors at a given size at lower power? Lastly, two years ago when we introduced our high-volume, High-k/Metal Gate process, it was about changing the material of the Gate stack.

It really is how you incorporate new materials to continue Moore's Law and continue to shrink transistors to get that performance benefit. It's also the ability to do this cost effectively. There are lots of technologies out there that are in research, but they'll never be capable of running 200 million microprocessors in [volume] for a range of applications. Determining the right points to integrate cost effectively, to deliver that to the marketplace is key to our ongoing innovation.

Going beyond High-k/Metal Gate, what's next? Let's spend a minute talking about an area called "three-five semiconductors." It's not new to the industry. It's been used for communications. But it's a promising area of research for how we incorporate that into silicon technologies. These are materials we're combining groups from Group Three and Group Five into new materials, and incorporating those on a silicon process.

What I'd like to do at this point is invite out to the stage Professor Jesus del Alamo from MIT, who's a leading researcher in this area of three-five semiconductors. Professor, if I could ask you to come out onstage?

[Music]

Thank you for joining me. Professor del Alamo is one of the leading researchers in this field, and Intel works with him in looking at research alternatives for how we can take his research and look out in that 2015 and beyond timeframe for new types of capabilities that we can build into our semiconductors. So let's get to it. Could you first of

all please tell us about the research you've been doing, and why this might be a promising area for our future?

Jesus del Alamo: Yes. Thank you for inviting me to be here. I have been working for over 20 years at MIT on three-five semiconductors. In fact, I learned my trade in Japan, where I spent two years working at NTT Labs after graduate school. In fact, I have been working on three-five transistors, which have been of interest for a while on communications and defense.

At MIT, my emphasis has been on power amplification, very high frequencies all the way up to millimeter-wave regime. More recently, we have taken an interest in exploring three-fives for addressing the challenges that silicon is facing as it continues to scale down.

Bob Baker: Okay. And why is this an exciting area of research as it might apply to semiconductors and silicon technology for advancing microprocessors?

Jesus del Alamo: Well, three-fives are really amazing materials. They have really unique optical properties. We use them to do lasers on light-emitting diodes. But they also have really excellent transfer properties. Electrons move very fast in certain three-five semiconductors. And that allows us to make transistors that operate at very high frequencies with very high efficiency.

So, for example, smart phones have several three-five transistors in them. And in fact, satellite links and fiber optic communication links

also employ a number of three-five transistors. So this high electron velocity can actually be exploited in a different way, which is to make transistors that operate on very low voltages while there's still this plain outstanding performance. And that is very promising to address the power challenges that we see in future high-density generations of integrated circuits.

Bob Baker: Great, great, that's wonderful. And what level of progress have you made in applying this to common day technologies for, you know, mainstream semiconductors?

Jesus del Alamo: Yeah, so in the last few years we have been exploring the use of three-fives for logic applications. And, in fact, this was not an obvious thing to do. It was not clear that these materials were suitable for logic applications. And, in fact, in our experiments at MIT, we have found that these devices really have excellent characteristics operating at very low voltages.

For example, our transistors operate at 0.5 volts, which is less than half of the voltage of mainstream silicon technology. We have also been emphasizing scaled-down transistors. We have made transistors that are 30 nanometers in [Gate] length. And that's about the size of the 45, the mainstream 45-nanometer silicon generation of today. So this is a technology that is looking very good. We have been making great progress. And, in fact, recently our transistors have broken several world records in terms of high-frequency response, even though that is not our main emphasis today.

Bob Baker: Great, great. But this is a different kind of material. It's used in, you know, different types of semiconductors for communications applications. What's the likelihood of it actually being incorporated into a silicon-based technology that we can ramp a couple 100 million microprocessors on?

Jesus del Alamo: Yeah, so I think combining three-fives on silicon really represents a very exciting proposition. I think the vision is to replace the silicon channel by a three-five channel, hopefully without breaking anything else. And our experiments suggest that that alone should speed up the electrons by a factor of two while operating at half the voltage. Now, if industry focuses its attention on this technology, the results that I have discussed, the [just] university results, then we should expect much bigger returns in performance on power.

And so this is not to say that there are not great challenges ahead. We still have many technical problems to solve. But there is now a community out there that is growing, that is spending a lot of effort, and that is going to tackle these challenges.

Bob Baker: Outstanding. And if just for a minute you could, what would Nirvana look like in five years? As a three-five silicon researcher, what would Nirvana look like in, say, the near future?

Jesus del Alamo: So my five-year dream scenario is that we will have a vision for a future CMOS technology that incorporates three-fives, and that we also have a vision for scaled-down technologies beyond that. Also, it's important to realize that three-fives is actually a family of materials,

not just one material. And so that means that once we manage to incorporate three-fives in silicon, there are many new exciting opportunities that open up.

And so I think in five years from now I hope that we are able to tell the world that the microelectronics revolution is here to continue and is going to continue to bring us the exponential progress that we have come to expect. And it will do that through the incorporation of three-fives onto silicon.

Bob Baker: Outstanding. Well, I really do appreciate you coming all the way from back east just to share with us the possibility for incorporating three-five and this very promising research work that you're doing, and how we might make it extend Moore's Law for 2015 and beyond. So thank you very much for coming to visit us.

Jesus del Alamo: Yeah, thank you very much.

[Applause]

Bob Baker: So there are number of technology opportunities to continue Moore's Law. It takes innovation. It takes investment and research and development. It takes creative engineers to be able to incorporate these in our base silicon process to give us the cost, volume, and capability that we need for future devices. It gives us scaling.

We are very confident that research work like this, while we haven't selected which technologies will be able to work out in that timeframe

for 15 nanometer and beyond, but there are a number of promising options for us to continue that. It takes that innovative research work and dedicated professionals like Professor del Alamo to continue to help us do that.

And Sean showed a number of great demos out here. I kind of get jealous because I can't bring those demos that I'd like to. They just don't fit in the carry on compartment of Southwest Airlines. But I would like to give you a demo via video of some of the things that it looks like in my world. What are the things that power these, that produce and make these products that realize Moore's Law? So if we could just run that video and give the crowd a sense of what our demo would look like.

[Video]

So I hope that gives you a good idea of some of the manufacturing capabilities and the types of manufacturing, equipment, and process technologies that go into manufacturing these amazing products.

And now I'd like to shift our attention to a little more near term. Let's talk about 32 nanometer, let's talk about 22 nanometer, and let's talk about this two-year cycle. Again, the goal is to power this two-year cycle that you see up here where we're delivering a new capability every two years. And now we're extending that out.

Paul and Sean talked about Westmere, and we're in production on 32 nanometers. But what I want to get a sense across to you is 32

nanometers is in production. Two years ago we showed test chips. It shows the planning that we've had to prepare four factories and invest $7 billion to get that capability in place. It takes equipment design. It takes equipment development. It takes new device modeling characterization and validation tools to make sure that these products work with the incredible integration that's going on today.

And it takes time to understand the new design methodologies to bring these out so that Paul can show a Sandy Bridge demo on stage while we're just beginning the Westmere ramp. It takes the investments that we make today. And it means you'll have the confidence as you develop our platforms that you'll know that those nodes continue in a two-year cycle. And you can depend upon your products and your developments to take advantage of them.

So two years ago on this stage we talked about this SRAM. It was a whopping 291 million bits and, you know, 1.9 billion transistors. And like clockwork it's in production today. I wish I could describe the incredible effort that's gone into it by our manufacturing technologists and the people who develop these processes. Because as you develop these, you need to cut your defect levels in half to be able to get the yields that you need and not have extraneous problems impact your designs. It's a 50 percent area reduction. It's our second generation of High-k/Metal Gate where we've already produced 200 million and understand a great deal about how that technology works.

And the other point to understand is this is better than any other 32 nanometer or 28-nanometer transistor that you might read about. The

characteristics of this 32-nanometer technology are far beyond that which has been reported anywhere in the industry. So it really is a phenomenal 32-nanometer technology.

The shrink of the Nehalem, all this promise of these products for Nehalem, we're working on in the factories today to produce it for Q1 products that you'll see on the marketplace, this will enhance to shrink the dye or add new features like Sean talked about, the encryption engine that's going to be added to Westmere and hardware. And we'll be able to bring it in volume to many, many different price points and the many different applications that we're looking at for our technologies.

Early this year we announced our $7 billion investment in four different factories. Some you've heard of before where development's taking place, like D1D. D1D is taking a more prominent role, and I'll get to that in a minute in terms of what we're doing for SOC technology development to enhance our product line. And the other two fabs are Fab 32 in Arizona and then Fab 11X. Tremendous amount of logic production experience in all four of these factories. They aren't new. We're either expanding or revising these factories to incorporate the new 32-nanometer equipment, and they're ready to go whenever the product demand is required.

It's the people who have the capability to deliver this and deliver it on a new cycle every two years. Sean talked about the 22-nanometer technology and I'd like to just hold up the way for one more time. I have to use this demo twice because the boss gets to use it once, but it

really is phenomenal. It deserves to be held up here twice; 22-nanometer technology on functioning, extremely dense; 2.9 billion transistor product.

If we could shift this to the next slide. What I want to go into is a little bit more of things that Paul didn't say. The work to refine the process goes on because we have test vehicles in addition to the [SRAM] to refine the yields, to take advantage of the circuit characterization needed to bring this on for microprocessors in the future.

It also helps us debug the new process equipment. Because you get a very large array, you can understand the capability across a large dimension. It's more difficult than manufacturing just a few transistors. So this capability is here today.

If we could go to the next slide, I want to talk about one other piece. The smallest transistor is .092 square microns. But there is another one for the first time. Usually we show you one SRAM and this one shows you a second one. It's a lower voltage and slightly larger because we're working in parallel on a lower power transistor at the same time we are the higher power transistor for our high-end microprocessors.

Not only that, there are test circuitry on here for analog, for I/Os, and for other key features that we need, not only for the super fast CPUs, but for integrated SOCs that will come to market when this technology is ramped in production. So those are the circuits that we're working on today, not just for CPUs. It's really an outstanding capability.

So we've talked about this SOC business and what it means to us; the wider range of this continuum of devices from high-end to low-end and different applications and what's required. And what I want to talk about is what it takes in the factory and the planning systems and in design and capabilities within the company to make us more than CPUs.

To expand what we've talked about in our consumer electronics, we've talked about embedded and we've talked about handhelds. In these three market segments, Intel has been pretty clear that we're focusing on for opportunities in the future.

What I want to go back to is 45 nanometer, though. People write about Intel as new to SOCs, but if you look in each one of the categories we're producing, and you'll see at the show during these three days, different examples from Sodaville in consumer electronics, which our system monitor features for the consumer market. You heard Sean talk about Jasper Forest, which is two Nehalems with, again, a different level or system level integration to target that for communication and other markets.

And lastly Lincroft, which is part of the Moorestown platform to bring that SOC to market for small, low power, very, very efficient handhelds. All of these are based upon the technology and the learning that we've had on our high-volume CPU technology processes.

We've talked about in the past at other forums, and you've seen this from Intel, is how the optimization of being an integrated device

manufacturer is so critical to the success of delivering this two-year cadence technology node. It takes a process optimization, but it also couples with our design tools and how we design and build our masks to be able to do this in a cost-effective manner.

It also takes alignment with manufacturing capability to respond to demand upsizes products transition from, one technology node to the next. And lastly, it takes high-technology packaging. And all of these are optimized together. You can't design a high-end, High-K/Metal Gate transistor with strain in it and have the package induce other strains that you can't control. So it's got to be integrated as a system.

The other part is as we get into an expanded area and to SOCs we have to look at how we're going to support customers. It isn't just one or two large CPU design teams at Intel. It's many smaller design teams who are enhancing the cores and bringing those to market.

So it takes a different kind of customer support, even to support an interface with our divisions internally, so that we're able to handle a Westmere first production while we're doing a new Sandy Bridge on top of the unannounced SOCs that we're working on for 32 nanometer.

It takes a view of platform; what platform quality is required, how the package is integrated. And it's becoming more and more critical how we interface with the software to bring that entire package together. Intel has these unique capabilities all the way from mask to [silicon] and design tools, to software and packaging and platforms to bring it to life. So these are really outstanding.

This is an overview of our SOC processes. Back on our 45 nanometer we had a process called 1266. That was the first High-K/Metal Gate process and we made a dot process. We added features for SOC, but it wasn't a full-blown process development target at the SOC market. But as I showed, we do have 45-nanometer products based upon a derivative of that technology.

32 nanometer is the first time that we've developed unique process technology features for SOC product design and the capabilities that they require. And as you can see, on 22 nanometer we have 1271 on the drawing boards looking at how to design 22-nanometer technology to take advantage of that in the future; capabilities that incorporate a number of different features. It's very, very different doing it at 32 nanometers when you incorporate capacitors and inductors.

We're also on the right-hand side providing options to our design teams for different metal stacks. Different densities that you need for cost-effective SOCs may have different characteristics in the metal stack to get optimal performance. So 1269 is designed for those capabilities for our divisions to choose for their product designs optimized for the capability they need. And that includes special analog circuitry that's required for these types of designs. All of these are built into this 32-nanometer technology that we've talked about.

Not only does it take the features on the [die] and in circuit, but it takes other features. It takes leadership designed for manufacturing and test tools. How do you cost-effectively test these types of products and

what do you invest when you're reusing a core, yet have low test times and you want low failure rates and high quality products?

We're also working on a modular IP database that our designers can share across the company, both internal and external IP where we're able to incorporate that on our process technologies. And it goes to converge tools and methodology so the design teams across the world can reuse and use the same design methodology and get the results that they're expecting off of our high-volume microprocessor and other SOC developments. So it provides that entire layer across our company to address these different markets.

It also takes different packaging. SOC is not a high-power microprocessor. We need small packages that go in thin and light. We need different types of packages for handhelds. So we're developing a range of technologies integrated to our process so that we know that they'll be reliable and high quality, as well as the cost points that we need for certain types of SOC product designs. Because you know that it takes a different operating point for a handheld than it does a laptop than it does a server.

The other key thing that that will do is we just start running one or two high-volume microprocessor designs. It's going to take better responsiveness from the factory to address these markets as well as we do the CPU market. So it's requiring support across a number of different product lines, number of NPIs and delivering this when our divisions need it to get it into customer designs and customers' hands.

Over the past couple of years, we focused on factory throughput time. That's kind of the core; how fast can you turn those new designs? How fast can you put product through the factory once you get an order? And even though we've introduced 45-nanometer technology with High-K/Metal Gate, we've been able to reduce throughput time by 50 percent. And in the past year we've cut it some more, for an overall reduction of 62 percent in factory throughput time.

Again, for our customers it takes more than that. The second part is how do we commit better? This is a graph of our focused area on Just Say Yes. When a customer calls for an order, what percent of the time are we able to answer that day that we can deliver that product when they need it or give them a commitment for when we can deliver it? You think about Intel being complicated with multiple factories around the world and multiple assembly tests on a number of products. How fast are we able to respond to the market? So we focused on how can we make that product more available? Just say yes and make sure that we're delivering that product on time.

And not only that, how do you deliver it faster? The last is from the time you get an order to the time you can respond to a customer. You need faster factories, which enable better ability to say yes to tell a customer when you can take that order and then deliver it in a shorter amount of time. All of these things will be required to address the product ranges that both Sean and Paul talked about today.

The next area is not just in SOC design or new packaging or how our factories support but it’s in how do we use different types of

technology, and [NAN] is one. Here's the technology progression of our NAN product over the past short while.

And what I'd like to do is talk about how we're looking at how our customers use our products, how we feed that back into not only our designs but our manufacturing to come up with new innovations. And I have a short video that shows one customer with what it's been able to do for them applying NAN technology in a new way.

[Video presentation]

Bob Baker: So this gives you an idea of how to use a new technology. But there's much more in the future for how we apply this technology and how we get it to create an optimal platform delivered to our end customer so it gives them value.

But I'm not the expert in this. What I'd really like to do is invite Rick Coulson who's an Intel Senior Fellow and director of our storage and I/O technology operation and our technology and manufacturing group. If could have Rick come out, I'd like him to be able to explain to you how we're looking at these types of NAN technologies and how they marry up to create the optimal platform. So, Rick?

[Applause]

Richard Coulson: Thanks, Bob.

Bob Baker: Great to see you.

Richard Coulson: Yeah.

Bob Baker: So I'm glad that I have an expert out here because as you know I am not the world's expert on I/O and storage technology, but I know you are. So let me just turn it over to you not to talk about Intel SSDs that are in the market today but what are we looking at and what are we researching with our platform groups to develop and apply technology in new and creative ways.

Richard Coulson: Intel SSDs are delivering great gains to things like the game development and the game play you just saw but with existing platforms, existing software just the way it is. What we're focusing on is co-optimization of the SSD in the platform so that as we move forward we can get even more gains in power, in performance and in cost of the total platform as we make these optimizations.

My group looks at SSDs and how to improve SSDs. We do research in what we can do to-to make Intel future SSDs better. But what we also do and probably more importantly is we look at the platform and see how the SSD and the platform play together and what kinds of things come up, and we look at the interface between the SSD and the system, the software, overheads, those sorts of things and find platform bottlenecks [unintelligible].

So what we've been looking at, you know, the current [source] stack was designed with [disk strength] in mind. And so SSDs are about 100 times faster in terms of latency than a disk drive. And so everything

about it, the physical interface, the protocols on that interface, the software driver and the chipset interface all have that kind of thing in mind. And while we do get the great gains now with SSDs, as we look at optimizing some of those things like interrupts and driver speculation, improving the physical interface between SSDs and the system, we expect great gains in power and performance and cost so that's good.

Bob Baker: Great.

Richard Coulson: So we can expect a lot there. But what's exciting to me is to think of it as what used to be the storage performance of a workstation, now it just becomes the same thing now in a handheld device. It becomes easy. And the performance that you used to expect out of a server, now you can get that on your notebook. And the performance you used to get out of the biggest data centers that have the most racks and racks and racks of disk drives is what you can expect out of a – just a normal, small-medium business kind of server.

In fact, I brought one for you to see.

Bob Barker: Good, we do have a demo. It's more than a factory. Okay.

Richard Coulson: Yeah, yeah. This is an Intel 5520 AC server running two Xeon 5500 series processors. And here we're doing 1,076,000 IOPs, 4K [2:1 Read-Right] IOPs. We're doing that with seven Intel SSD technology prototypes right here. And there are three of them in this expansion PCIE expansion box, and there are four of them in this box.

And so with that pretty realistic workload we're getting a huge number of IOPs. And in this demo we've done a number of things to emulate and implement some of the things we've been talking about in terms of platform optimizations and the like.

This many I/Os per second is about 4 gigabytes a second of storage bandwidth – random storage bandwidth which for – I know for you it doesn't mean that much, but for me as a storage guy that's a huge number, a very huge number. So I'm really excited about that.

And the [CPU] utilization, just in the small-medium business server doing more than a million I/Os is about 50 percent, so really nice.

Because I know this isn't your thing I want to put this in perspective a little bit. This is about 5000 disk drives worth of random performance. And 5000 disk drives when you think about it, that's a lot of floor space. Maybe more importantly, that's about 50 kilowatts of power. And this whole setup here is about 500. So it's about 100x for random I/O.

And if you think about managing 5000 of anything that's really difficult. Managing seven, that's practical. So what was kind of impractical becomes totally practical in the small-medium business server. So that's what I'm excited about. A million I/Os is a nice round number. So . . .

Bob Baker: Right. Well, Rick, thanks for coming to talk about how we're applying new technology in different ways to enhance our platforms and bring real balance and real capability to not only large data centers but also the small and medium enterprises. It's really exciting to see the work that you're doing.

Richard Coulson: Thanks.

Bob Baker: Thanks for coming out and sharing with us.

Richard Coulson: Okay. [Laughs] Here we go.

Bob Baker: Thanks, Rick.

[Applause]

Bob Baker: So what we've talked about today is the relentless pursuit of Moore's Law, that there are technology options to keep this two-year cadence and bring additional material features to the gate of the transistor, to continue our scaling, to continue to operate at lower voltages yet with higher performance. Absolutely critical to bring the promise of that technology to you, our developers, over the next 6, 8, 12 years, as long as it goes. But there are promising technologies that we're looking at for incorporation at the middle – the later half of this decade that we believe absolutely will continue Moore's Law.

Not only that, we've talked about, you know, our 45-nanometer technology that two years ago was an industry trendsetter of

incorporating a High-K/metal Gate that other people talked about but aren't shipping. We're not only in production with our second generation, it's on two major CPU products: the Westmere which we'll introduce shortly and then the future products which are in development like [Sandy Bridge] and others that you'll hear about. These are absolutely critical technologies while we're developing 22 nanometers, and in two years we can come and talk about how we're in production on that technology with another 50 percent area shrink.

Not only that, for the market opportunities that Paul and Sean talk about, and you're going to hear about for the rest of the two days of the conference -- How do we expand our capabilities for system on a chip? Not only to have the technology features and silicon that optimize it for our products, but how do we have the manufacturing, the packaging, and other unique technical capabilities that are going to be required to extend the IA architecture into new areas where we can all benefit from writing that code and proliferating it into different systems throughout the industry.

So thank you very, very much for your attention, and I look forward to Moore's Law's bringing all of us opportunities for our businesses in the future. Thank you very much.

Female Voice: Ladies and gentlemen, this concludes the keynote session. Press and analysts, we invite you to join us now onstage for a photo opportunity. Have a fabulous afternoon, everybody.

[End of session.]