David Perlmutter - 9-13-10

[Beginning of recorded material.]

Dadi Perlmutter: Good morning, folks. Thank you. Good morning and welcome to IDF. Well, the issues that Paul makes, all the promises, I have to go make them happen. But I will start with a dream.

Everything that's part of the greatness of mankind, let's say, humans that for thousands of years, dream about something and then go make it happen. And whatever is unobtainable today may be obtainable tomorrow or the day after tomorrow. And I'm going to start with the dream my daughter had, because I talked to her several years ago, my youngest daughter. She's now close to 18. She was about 15. And I asked her, you know, when you grow up, what is that you want to invent? Do you want to do anything? She said, "You know, dad, I'd like to have a way to connect something to my brain so when I wake up in the morning I could watch on the video my dreams." I scratched my head and said, "This is going to be quite complex." May going to take a little while.

But I have all the reason to believe that maybe not in five years, maybe a bit more than that, that may be achievable. Why not? I bet there's a lot of progress in brain research. There's a lot of progress in the way we understand the signaling of the brain and definitely great progress on our ability to process complex signals and get them onto a video screen. And I'm told that Justin Rattner, our CTO, is going to talk about capabilities in development today to read brain waves. So maybe my daughter, maybe even in my lifetime, will be able to see your dreams.

And that's what we're all about – innovative ideas, innovative thoughts and capabilities and making it real. The extraordinary of yesterday is the ordinary of today, and the extraordinary of today is the ordinary of the future. That's the way we are. That's the way we like it to be. That's basically how our purpose in life, making these things happen. That's the pride that we take.

The computing experience has evolved significantly. You know, I always thought when I talk about this one internally to Intel, I still remember the good old days that I connected to my computer with a punch card reader. When I told it to my son, he kind of scratched his head and said, "I don't believe that it really ever existed." I don't think many people in the room even remember this. Now, it's all about graphics interfaces and moving to more sophisticated stuff.

The way we used to input things into our computer, not just too many years ago, was all about number crunching. We fed numbers one way or the other to the computer, and we got some results on the outside. We want to process today much more sophisticated data types. They're all eventually translated into numbers so the computer can use them, but there are now on photos and video capabilities and about 2.5 billion photos on Facebook today and 30 billion videos really viewed per month.

And creating media, sophisticated media, is not an issue of a hundreds-of-millions-dollars investment in Hollywood studios.

Every one of us create very sophisticated media and want to put it on another device and consume it and send it. And it's not just transcoding and storing and playing. It becomes into analyzing the data and getting some information, some intelligence out of it. And the trend will just continue to go.

We'd like to get more interactive with our computers. We'd like our computers to understand us better. Today we use mouse and GUI, going into touch, but would like our interaction with our computer to be much the way we interact with each other. Big portion of our language is the gestures that we do with our hands, with the look of our face, the tone of the voice that we are using, and saying exactly same sentence will be interpreted by you very differently if I'm smiling, my body language, the way I do things. And this is the way we'd like the future of computing to be. Much more interactive.

So let me see if I can do that right away and use my hands and make sure the computer understands me. So this is a bunch of pictures, and I'm told that it should move. Well, the dog is at the right place. And I'll do this. Okay. Maybe the guy I'll call in will try to make it better. Oh, here it is. Okay. My wife said that I'm exactly the same. I don't understand what she means. So I'll try to call into the demo to make sure that we really understand gestures.

So I'd like to call Francis MacDougall, the CTO and co-founder of GestureTek. So I presume you'll be way better than I am on this, right?

Mr. MacDougall: We'll do our best.

Dadi Perlmutter: Thank you. So.

Mr. MacDougall: So GestureTek has been doing work with human tracking for the purposes of interaction for about 20 years. Historically, we've done it with 2D cameras in everything from entertainment, digital signage, health and fitness – we're even on cell phones. So what we're showing you here with Dadi, we're excited about a new evolution of 3D cameras coming into the consumer space. So there's a camera in front of this one that just let's us do very simple . . .

Mr. Perlmutter: Ah, it looks easier for you, you know? You have the magic hand? You have something in your hand?

Mr. MacDougall: That's right. It's tuned for me. And so you can also do the resize, so you're really getting multitouch control in a standard setup. But let's go behind the demo and see what's really driving this. You can see a 3D camera image being tracked on screen there. The yellow parts are my body. As I walk closer to the screen, you can see that I'm actually getting the depth for every pixel of motion here. As I raise my hand, we're able to track that hand in free space. We can do it with two hands. You could actually put ten hands in there and track those. But what do we do with that?

So one of the things we can do is actually guide the mouse, and you could imagine controlling a TV from the couch or any kind of control app. I'm going to launch a little application here by going over the top of it. Now, you see that it's kind of getting my body image out of the background, and you're wondering, "Well why doesn't it respond to all gestures?" Everybody wonders, how can you tell an intentional gesture? So what we suggest is the idea of an engagement gesture, so when I put my hand up in front of here, it'll detect the five fingers, and now I can be led into the center.

And we have a little gesture wheel here that lets you launch applications. Now, in a living room environment, this could be your six favorite TV shows. It could be YouTube, G-mail, whatever you like. In this case, we're just going to launch one with a single swipe, and we've got a little avatar control here, so a little cute guy that I can do my leaning around. So I've got head tracking, center torso, as well as the two hands being tracked. Okay? So what I'll do, though, now is show you where we really want to go with this, and we'll go back to the desktop with the thumbs down gesture. Okay?

So where this can go in the future? Full avatar tracking, so here we have a full skeleton model being tracked in real-time with floating point, but this is on a very meaty machine. What we're really excited about is, with AVX, being able to do this kind of thing – motion capture – in the home.

Adam: What's up, Dadi?

Mr. Perlmutter: What's this look?

Adam: We're doing the motion capture demo next, right? What do you guys think? Do I have everything in the right place?

Dadi Perlmutter: Well, I don't think we need it. I think we have enough performance on the CPU, so it doesn't have all these sensors on your body to do it anymore, so I think you understand this gesture.

Francis MacDougall: Thanks. Poor Adam. So, I think Dadi, we're going to end on a little demo of a game. So this is an off-the-shelf game. I'm going to use my hand position in order to control this. And as I make it wider, I get higher speed, okay, and then standard steering control is what we use. So this is an unmodified game. We're just running it directly, and I should be able to catch up to these guys pretty quick because I just keep the accelerator on all the time. Of course, the most fun in any racing game is to try to destroy the car.

Dadi Perlmutter: Okay. So what really makes this now doable versus before? You said you worked on this one for 10 years, so what makes it available now?

Francis MacDougall: So the big change is that 3D cameras have now dropped in price hugely. So they were $10,000 two years ago. They're coming out in the $150 price range. The second is, of course, the increase in speed that we're able to process floating point data with on the new architectures.

Dadi Perlmutter: Okay, thank you very much.

Francis MacDougall: My pleasure.

Dadi Perlmutter: Thank you. This technology will become even better, because the more processing power you have in your machine, the cheaper the 3D camera can be made. And in the future, we'll make it even more common, more available, and making sure that our machines are going to understand us better. But the importance of high-performance is not just on the computers to understand us and our gestures. There are a lot of complex data analytics analysis happening today and are huge opportunities for business.

This huge amount of data which is being collected using video cameras, sensors source on bridges to make sure that the bridges do not collapse when there are a lot of cars driving all over it, and, of course, in the future, there's going to be a lot of sensors into our body, with a lot of communication to the doctors or the HMOs to tell them about our health and our situation, so we can do a lot of prevention.

This requires a huge amount of analysis that is going to be required, and I'd like to invite Paul Rodwick, a VP from Oracle Business

Intelligence, to go talk about what Oracle is doing on business intelligence. Hi, Paul.

Paul Rodwick: Good morning, Dadi. Well, you're right, not everything is as exciting as a video game, but really, to businesses, it's just as important or even more important. It's not just the size of data that is increasing, but really the user community for business intelligence. What we've seen is the days from where it was just a few handfuls or maybe hundreds of analysts needing to get fact-based insight about their organizations in order to decide what to do and why to do it, and now we're moving to hundreds of thousands.

One deployed customer of Oracle Business Intelligence in the retail industry is moving to 300,000 live users, and one is already out with 1 million users over the public internet.

Dadi Perlmutter: Wow.

Paul Rodwick: Now, what this means is there's much greater concurrency requirements, because if you put up a pretty BI dashboard, sometimes that fires off 30 or 40 different concurrent queries. You multiply a million users by 40 concurrent queries by a lot more that people are asking, and you really need a lot of processing power.

Dadi Perlmutter: Okay.

Paul Rodwick: So what Oracle and Intel have been working on is to optimize system performance all the way up from the hardware through the highest tiers of the software stack. And an example of that is the Oracle Exadata Database Machine that was delivered over the past year. And now we're working with the Westmere-EX. Already have Oracle Database and Oracle Business Intelligence running on the machine, and that promises to give far greater performance and scalability for business users to move beyond just the walls of their own building, but to also offer fact-based insight also to their customers, their business partners, and a vast network of people.

Dadi Perlmutter: Okay, thank you very much, Paul. That will be excellent.

Paul Rodwick: Thank you.

Dadi Perlmutter: Thank you. And we have a Westmere-EX box over here. This kind of application used to require much bigger computers, and now we can make it available to much more businesses that could afford buying these boxes.

And this is enabled by our new Westmere-EX, which brings new capabilities enabled by our great scalable architecture and 32-nanometer technology. It increases the number of cores from 8 to 10, which is a 25 percent increase, which means with 10 cores you have 20 threads could run in parallel on this piece of silicon. It increases the memory from the previous Xeon 5500 by 2X to 32 gigabytes per DIMM, which means now a four-socket system could have up to 2 terabytes of memory, which is, as I mentioned, 2X more memory capacity, and memory capacity is a huge deal to handle these kind of databases that Paul from Oracle was talking about.

It is socket-compatible, which reduces the development time for the developers to make this one happen. And, of course, it comes with the newest security technologies, which are very important to protect the data in the analysis and everything that we do. The secure technologies like the encryption instructions Paul mentioned in his talk and the TXT feature that is extremely important. And the availability of this kind of machine is going to be in the first half of next year.

Performance is really very much about real-time and high-quality computing. I mentioned my daughter earlier, and she, of course, her life is around Facebook. And what do people this age want to do when they get from a party? They have the race. Who is going to post the video or the picture first on Facebook? So when other people log in, the pictures are going to be there, too. So it's all about the speed of processing.

I don't think my daughter wants to wait five minutes. She will want a faster computer that may be able to afford it to go be ahead of everyone else. So no one wants to wait minutes. People want to really do it in seconds. And that's what people tell us. What people really tell us is that they want all these capabilities. And we talk to thousands of people around the world from seven different countries on every continent, and the message is quite consistent. It's about responsiveness, it's about connectivity, it's about visual computing and visual experience, and, of course, being power efficient. And users want it in real time, not to wait for ages.

And these kinds of things don't just happen. The consumer, the people that interact with the computer, want simple interaction. They want to do the thing they naturally do, they don't have to get trained to do that, and just get it done. But it requires a very sophisticated technology, a very complex microprocessor, and a very complex system, and a very complex solution to go and make it all happen

And this is our calling. This is what we do. Year after year after year, we deliver a new technology. Four years ago, I stood on this stage and we announced the first generation of our converge-call technology called Merom. And then we went every year and introduced another one. Two years ago, we had introduced the great technology of Nehalem, which really fits up our EX and EP servers and our clients. And, of course, we came with 32 nanometer just at the beginning of this year, and it goes on and on and on.

Today, I'm extremely proud to go talk about Sandy Bridge, which is our next-generation microarchitecture. This is a unique one, because this one we are very much putting together basically all that is required, the whole system of a PC, on a single piece of silicon.

This is enabled by our leadership 32 nanometer in a second-generation high-k gate, because it not only allows us to put around a billion transistors, but it allows us to do that in even lower power and better battery life than the previous generation.

And just to give you a perspective of what a billion transistors really means, first, human beings are really enamored by huge buildings. And the Empire State Building has about 10 million bricks, just 10 million. Ten million is an interesting number, because I checked the Intel history. And we used to have 10 million transistors on Pentium 3, somewhere around '99 timeframe. And we had about 10 million – a little bit – on Pentium 4, something around 2004.

You know, I talked about dreams and ideas. The human brain has about 100 billion neurons. Well, a neuron is a little bit more sophisticated than a transistor, but nevertheless numbers are getting close. And if we are able to continue our pace – which I believe we will continue – 100 billion transistors are not much out. About 10 years, 11 years out. So we are getting close.

And you could imagine what kind of capabilities would be enabled. We're going to ship this product and put it on the shelves early next year. And it has a lot of architectural advances I would like to share with you.

The first one is really – we said putting graphics onto single chip of physical. Beforehand the CPU used to communicate to the memory controller, and the graphics were talking just to the memory. And usually – especially on graphics and media – the bandwidths of the memory interface is extremely important.

So what you really want to do in a smart integration is to really get the graphics onto the CPU, but not just continue to connect to the memory controller to the memory, but really use the huge cache that exists on the CPU. This one gives about four to five X faster throughput when the data is in the cache, which is a huge reason why our performance is getting better.

On top of that, we have of course significantly improved and reoverhauled the architecture of both the graphics, the media, and the CPU. But when you put the things together, you could make things really smart. And one of the smart things you could make when you integrate things together is power management. At the end of the day, it's all about power; it's all about battery life.

Last year, we just brought you with the core technology, the innovation method called Turbo. And the Turbo is an extremely good method because what it uses – it gets the CPU, uses the thermal headroom that exists, and just boosts the performance when you need it. Because it's all about the response time.

But Sandy Bridge brings Turbo to the next level. And what you see in the white color is basically the extra benefit you get from the Turbo of Sandy Bridge because it uses more capabilities and goes even beyond the TDP limit. But the real beauty of this second generation of Turbo is when you get to more cores. Because one of the limitations that existed on the first generation of core was that the more cores you added to the game, the less benefits you get from the Turbo because you already exhausted the thermal limits.

Well, it's exactly the opposite with Sandy Bridge because of the architecture and circuit design techniques as well as the ability to go beyond the TDP limit. As you can see, other than adding one band of frequency, we get many more bands of frequency, which is exactly what you need to get to fast response.

Now, if you put the graphics on this CPU, you're going to get same phenomena done on the graphics as well. So the graphics would be able to use the headroom of the thermals on the system and boost its frequency, what we call the dynamic frequency feature. And it all plays together because in many cases when you run the graphics at full steam, the CPU is not that active and vice versa. So you could really use and get the best performance and the best capabilities by the time you need it. This is what we call designing for fast responsiveness.

So this is something we are extremely proud of, and we believe it's going to bring new capabilities to the market and make sure that we do a great job. And this is going to be the logo, and this is why we believe that Sandy Bridge really deserves the visibly smart computing solution.

And it changes very much the way we look at a PC. Because now it's become – and I'm going to show that later – a center of many things you do. Anywhere from media processing to support stereoscopic, three-dimensional, Blu-ray playback to high quality and increasing the quality of the video, making it better. Paul talked about the wireless display. Anything from high performance 3-D gaming and very much the vector extensions that is going to make many things capable. And I'm going to talk about that next in my talk.

The first one I want to really show the media and the dynamic range of photo editing, and I would like to call Craig to go show us what this machine can do. Hi, Craig.

Craig: How you doing, Dadi? All right, let's take a look at some things for Sandy Bridge and this next generation of our second-generation core. So what we're going to show today is we really want to give you guys a general feel and impression of how second generation of core is really going to improve not only just the incredible performance uptake that we get from generation to generation, but literally every corner of the system is going to improved for performance for literally all the applications that you're processing.

So what I have here is a little photo program. And this has actually been in the news a lot lately, so let me go ahead and kick these off. But if you guys are familiar with high-definition or high-dynamic-range photography, what we're taking is some of our existing photos – and the way this works, Dadi, is I'm either able to take multiple photos or, in this take, use a whole bunch of algorithms to dynamically change how a single photo looks to bring out all the vibrant colors and all the great detail that we want to take a look at.

And we're really going to show the muscle of this second-generation core, versus our current generation of the core, i7. And you can see here we're burning through these pictures, and literally we're looking at all the detail. We're literally going through and taking multiple exposures and dragging out all those great parts of the picture. And, well, it looks like we're already done on this side. So let me go ahead and I'm going to bring up a couple of quick examples about what we did.

This guy is going to take a little bit longer to finish over here, but I just want you guys to get a feel of just an in-general, overall performance. We're getting a huge performance boost and, well, look at the overall output that we get from just making our media. So from one picture to the next I can see all those bright colors are brought out – and we'd like to thank our guys from HDR Studio, in order to make this demo possible.

So some truly great stuff. And how it's going to be able to process all of your media even faster than ever. So some pretty great stuff, Dadi.

Dadi Perlmutter: It's good even for me because I'm a lousy photographer.

Craig: Well, this will make your pictures look gorgeous no matter how many times you're shooting at the ground. But I want to give you one other kind of a sneak peek. Actually, maybe we'll do a little quicker demo.

Dadi Perlmutter: Sure. Go ahead.

Craig: I have one for the next one. So let's move on down, and what we want to process on this one is we're going to take a look at encoding some video. And what I have here is Media Show Espresso from our friends over at Cyberlink, and we're going to go ahead and take a clip and encode this clip to maybe something I could send to my mom on her iPhone. So I have two of my clips here.

And as I get these guys ready, I have a beautiful, high-definition source file at 1080P. And as I go ahead and bring up my Apple profile -- also want to tell you that this is also running at 30 megabits per second. So we're huge, high-definition video, and I don't want you guys to blink while we kick this off because you just may miss it. So take a look at our second-generation core versus core i7 processor – oh, man. It's really taking off.

And we're looking at just a couple of seconds where a file that big is now done. So we're talking about literally dedicated media processing built inside the chip itself. So as any sort of user, even in a mobile session, I'm able to get this incredible performance and encode all this video on the fly.

Of course, we want to talk about some new advances in our existing features, so let me go ahead and show the audience what we're talking about –

Dadi Perlmutter: The other one is still working, I presume?

Craig: He's going to take a while. So in the meantime, I'm going to go ahead and have this guy connected via WiDi, and let's show the audience what we had in our beautiful, high-definition clip here as I'm sending it up to our TV here.

So we're able to show – we have a few more tricks that you'll see in WiDi by the end of this year, but some truly amazing stuff and an overall great experience for all of your media, accelerating it all across the way for all your applications, Dadi. So I can't wait to put this one out.

Dadi Perlmutter: Thank you very much, Craig.

Craig: Thank you.

Dadi Perlmutter: So Paul was right when he said it's going to take 4 minutes, 2 seconds. The other one is about to finish. But there are more things we'd like to do with our computers. And the beauty of putting everything together on a microprocessor, now every transistor computes. And you can use these billion transistors to make something for you. And we invented something called AVX, which makes vector processing way better. And this is being used also to do a lot of media analysis. So I'd like to call Adam now. I hope he's not offended too much. Or you're not. And you look better now.

Adam: Thanks for giving me another chance, Dadi.

Dadi Perlmutter: Sure.

Adam: I get a lot of compliments on that suit.

Dadi Perlmutter: Oh you did? Okay.

Adam: Yeah, I might have to put it back on for a photo shoot for my Facebook page.

Dadi Perlmutter: Okay.

Adam: Anyways, let me start by introducing this. We're going to do a bakeoff demonstration here. So I went ahead and put two of these great all-in-ones on stage. This is our current generation that we're shipping. It's an all-in-one. And what I have on over here by you, Dadi, is an all-in-one from Acer, and this features our second-generation core. What's great about our second-generation core, we do get that great AVX, but we also are able to bring down the power usage of it as well, so we're announcing 35- and 45-watt desktop parts for these.

So as we can see, we can continue to improve the performance. And with the great floating-point performance you get with AVX, you can use it on, as you said, a lot of media demonstrations.

Here we have an application from Yuan Fang, and this is an interior design application you see here. And you can imagine this being in the hands of somebody who does maybe professional interior design or something. You can choose different wallpapers and stuff. But I did put together a benchmark, so let me go ahead and kick it off on both of these. And I have a workload that is going to crank through. I'm going to kick them off at the same time.

Pay close attention to the bottom of the screens, and you see a blue bar racing across on the Sandy Bridge system here, whereas our last generation system -- it'll get it done using SSE, which is great, but because with AVX -- we can double the width of the extensions that we get with SSE using AVX, and so we get a terrific up to 2X throughput for these SIMD extensions, and we've got a dozen new instructions that our software companies can write code to this, and as you said, multimedia, but also financial and other things that are a heavy floating point can really see this performance. And what we get is what was a workstation-class professional thing, everyday people like you and I can do these things real time on our PCs. We have AVX from laptop all the way up to server, so it's great to see these things move into real time for real users.

Dadi Perlmutter: Okay, thank you very much, Adam.

Adam: Thanks.

Dadi Perlmutter: And my other daughter, and I have many of those, she's an architect. I bet she'd love to have these capabilities on her desktop.

So let's move on. Paul talked about security and trust. The thing that we're going to bring with Sandy Bridge and future microprocessors are more and more capabilities. Either using our new instructions, new capabilities we put, like a DXE on the microprocessor, or capabilities that we put on our manageability engine, which is very much an engine that would run under the OS, a secured kind of application.

And we're going to bring things to life that really are important for us as consumers or people running businesses, anything with protection of identity, anything that has to do with detection and prevention of malware. We all know that sort of thing that it will do. And these are big, wonderful capabilities we are going to continue to bring on and to make sure that the software and the hardware features are working together to make it better and more secured.

And, of course, features like securing data and assets, and we already introduced the antitheft technology, and we're going to continue to improve it to make sure that your assets, if you lose your computer or your computer has been stolen, the data in it and your identity in it is well-protected.

And, of course, last but not least, when something bad happens, and things do happen, a fast recovery and enhanced patching that will be able to do via vPro, and we're going to continue to improve these capabilities going forward.

It's all about the continuum. So you bet that Sandy Bridge is not going to just stay into the PC space. We are going to bring it to other computing devices to be able to use these wonderful features and capabilities.

So let me start with something which is not usually talked about as much. Everybody knows that Intel has a huge amount of presence in the data center and in the cloud, but very much the world communication is run on IA. And up to tonight was about application processing and control processing being delivered to the telecommunication equipment manufacturers or service providers via our embedded group, but with Sandy Bridge, with the new capabilities, better performance and, in many of it, the vector processing allows us to continue to replace things that used to be done in the past by dedicated DSP, digital signal processing, devices.

And the packet processing is increased in a significant manner. What it does, if you read what K. T. and Wawi said about this one, it allows them to go move from dedicated architecture that will require a huge amount of development time and software and maintenance to really move to general-purpose computing, writing the application, writing the software, moving software from one generation to the other, for one application to the other, and really get to the situation that their time to money and ability to sustain, maintain, and evolve these capabilities is improved.

We talk about analytics. And many of these analytics require real time, because, especially when things that you want a real fast reaction, you do not want to wait until the data is sent to the data center and coming back.

Now these are tasks – most of them could be done by a 5-year-old. If you look at your top right, you could ask a 5-year-old to say, "You know what? Tell me when you see some strange behavior." Strange behavior is really defined by one person going back and forth multiple times. "This may be suspicious, call us." Our computers today are not as smart. It is getting smarter.

And I could go on and on. Simple tasks for human beings, not simple tasks for machines. But you would like to make more and more of the things, especially when you talk about terrorism or fire hazards, traffic detection. You have all these capabilities. And a computer is very good. We have good enough performance to go detect that things are changing. And I would like to call Craig now to go show us real things what could be done with a good computer on video analytics. Craig?

Craig: Yeah, let's talk about some real-time usage models, huh, Dadi?

Dadi Perlmutter: Yeah, sure.

Craig: So what we have here is this development system. It's built on our Sandy Bridge, second-generation core. And actually for the first time, what we are able to do as far as in this real-time DBR and high-definition camera setup as well as video analytics, first time, an industry-leading competitive advantage as far as eight full 1080P streams from cameras just like this that we're able to bring in all at the same time and literally on top of that, as well as processing all that video, which previously, now that we're doing it all in CPU, could only be done before in a dedicated DSP. But now we're able to do all this real-time video analytics right on top of the feeds that we're capturing.

So let's go ahead and take a look at a couple of examples. And you showed this one off in your foil there, Dadi, but what we're taking a look at here is actually some flow control. So with the high-definition feed, not only do I get this incredible amount of detail, but I can draw literally in all these virtual fences or how I can go ahead and do counters and that type of control as far as collecting a bunch of data for however traffic is coming in and out, whether it's actually taking a look at real-time traffic motion or I can go ahead and maybe use for something like an event like this at Moscone Center.

So really the . . . that's really loud. It looks like we have some sort of a problem here on our real-time analysis. Dadi, what I'm showing here is that little box there is what we call a digital fence. And that's where I've put the executive donuts. Now the executive donuts are now completely secure, but it looks like maybe we had a little – uh-oh. It looks like Adam is back in the suit again.

Dadi Perlmutter: He was offended.

Craig: He's back in the suit again and really abusing his diet. I'm not sure if he knew those were yours, Dadi, but our apologies. Now at least we know who to punish after this one.

Dadi Perlmutter: Okay.

Craig: So as you can see, we can collect all this great data and process it on the fly, and that's really what we're promising with the second-generation of core in Sandy Bridge is, well, we can all get it done on the fly and have the performance to really kick it up to the next level.

Dadi Perlmutter: Thank you very much, Craig.

Craig: Thanks a lot, Dadi.

Dadi Perlmutter: I can now write a book, Who Stole my Donut? Of course we'll take this technology to the data center, and next year we'll put Sandy Bridge into our entry and efficient performance, and in the future into our expandable performance, and later on, of course, we'll bring our next-generation microarchitecture, code name Ivy Bridge, into the server lane. But what is in it in the future? The possibilities opened by an architecture like a Sandy Bridge, and of course it's going to be even better going forward, opens up possibilities for new types of platforms, platforms that we could only imagine once. And I would like to look at what the combination of sensors, smart sensors, and high performance can do. So I would like to invite Jeff Bellinghausen, founder and CTO of Sixense, to talk to us about sensing.

Jeff Bellinghausen: Hi, Dadi.

Dadi Perlmutter: Hi, how are you doing?

Jeff Bellinghausen: Good.

Dadi Perlmutter: So what are you going to show us?

Jeff Bellinghausen: So I'm going to give you a demonstration of the Razer Sixense controller running on the Sandy Bridge platform. So what the Sixense technology does is give you absolute one-to-one position and orientation tracking. So that means that I can pick up this object with the controller, move it left and right, up and down, all the way in and out. I can even take it and put it behind my back, and the tracking is always absolute.

Here, I'm using two controllers in this environment, so I can pick up one object with one hand and one with the other and move them around. You can pass objects back and forth between different hands. You can grab it with both hands and grow it to change the scale of the object. This application has a very high-performance CAD engine built into it. So I can take these objects and do different kinds of Boolean operations to cut the objects and do different types of modeling operations. And the interesting thing here is that this is very floating-point intensive, these cutting operations. So on the Sandy Bridge platform, it really utilizes both the full CPU and GPU performance.

If I just grab on to the world, I can navigate the view through the scene with either hand. And if you just grab on to the world and rotate, you can change the view or translate or even grab with both hands and zoom in to change your scale. So it's a very fluid, one-to-one, 3-D user interface.

Dadi Perlmutter: I bet you could use it for very nice gaming capabilities.

Jeff Bellinghausen: Absolutely. So this particular application is just a 3-D modeling environment for building different shapes. But, of course, this software is part of our Sixense TrueMotion SDK that you can embed into any kind of an application. So you can imaging something like a Google Earth style map viewer where you can zoom out and zoom in on the whole terrain view, or you can have maybe a SimCity type application, where these could be buildings that you're planning and drawing out roads or editing a track.

One last feature I want to show you here is that I can bring up this GUI panel, which is kind of like a virtual PDA, and I can use that to load different objects into the world.

Dadi Perlmutter: Okay, great. But, you know, I made this public that I'm not a big gamer. So I'm looking for some other reasons to use this capability.

Jeff Bellinghausen: Sure. So our initial market's going to be gaming, of course, and we're going to be coming to market early next year with a Razer-branded controller system, and we're working with ValSoftware to have support in all their games, including Portal 2. But of course there are lots of opportunities for other applications as well. One we're really excited about is the 3-D television technology. It really matches perfectly with the 3-D control system. To be able to control your TV menus or surf the web is a great application.

We're also working with our athletic partners to miniaturize the sensors, to make a wearable system, so you can have them on your wrists and ankles, and have absolute tracking of your full body for health and fitness. And of course for rehabilitation, you could use the sensors to log your rehab process and then upload that to your physician later.

Dadi Perlmutter: Okay, thank you very much. Very exciting.

Jeff Bellinghausen: Thank you, Dadi.

Dadi Perlmutter: If you could think about these capabilities will allow people really to make remote work, operating machinery, surgery, and all now becoming available using very much day-to-day, mundane technology like a PC technology.

So the future ahead of us is very much revolving around the themes. It's about getting what consumers really want. It's about fast responsiveness, great visuals, improved security and trust, and of course it will come in ever-shrinking form factors and will enable connectivity, as well as a huge amount of types of different sensors, different capabilities, all enhancing the PC and, of course, as Paul showed, the whole continuum delivered by Intel.

So what you used to think about what a microprocessor is, what a PC is, is not true anymore. The dreams of yesterday are becoming available today. And the combination of smart sensing, but very much a technology that enables us to put a billion transistors on one dye, that gives all the performance, all the capabilities, that every transistor really computes and gets something done, and it allows to put a huge amount of new capabilities, where the computer you have, which is the day-to-day PC, could be the center of gaming, displaying your best video on a TV, downloading, transcoding, doing all kinds of analysis on media.

It could be the center of video analytics, or any other business intelligence, because the opportunities are really endless.

So the opportunities are definitely not just ours. We're enabling. And you, the developers, either you do software using all the capabilities that exist on this microprocessor, or you build systems, or use, like the examples I showed you, connecting sensors into these devices to create great solutions. So thank you very much.

Female Voice: Ladies and gentlemen, technical sessions will begin in 20 minutes. Press and analysts, we invite you on stage for a photo opportunity. Thank you.

[End of recorded material.]