Keynote Address, September 13, 2012

Tigerfish®
Transcribing·Editing

203 Columbus Avenue · San Francisco 94133
toll-free 877-TIGERFISH

www.tigerfish.com

Keynote Address, September 13, 2012

[Start of recorded material]

Male Voice: Ladies and gentlemen, today's presentations contain forward-looking statements. All statements made that are not historical facts are subject to a number of risks and uncertainties, and actual results may differ materially. Please refer to our most recent earnings release form 10Q and 10K filing, available on our Web site, for more information on the risk factors that could cause actual results to differ. Thank you.

Male Voice: Ladies and gentlemen, please welcome Intel Chief Technology Officer and Managing Director of Intel Labs, Justin Rattner. [Applause]

Justin Rattner: Well, IDF is kind of like summer camp, huh? Summer camp for geeks. I hope you enjoyed that. I guess since this is the second year in a row I've come out with some kind of crazy headgear – last year I had Mooly Eden's beret, so I figured I had to top that this year.

Dadi Perlmutter on Tuesday talked about all the activity at Intel around perceptual computing, and I just wanted to make sure that you understood at Intel Labs we're part of that perceptual computing effort. But we decided, of course, since we're research people, we would go all the way to mindreading. And here's the prototype. [Laughs] Uh-oh. I have to tell you, I actually did an entire presentation with the ears on. And after a while, it's really a pain in the neck. You know, because nobody is listening to you, just as right now you're not listening to me, you're just looking at the ears and waiting for them to do something completely ridiculous. The instructions actually say that if you're very calm, the ears will fold down. But I'll be going full key, and the ears will fall down, so I don't know. Anyway, I'm not going to wear them for the whole keynote this morning, but this is our contribution to the perceptual computing revolution. Okay.

Well, good morning. I know it's day three, but I'll add my welcome to all the many welcomes I'm sure you've received already. Everybody hanging in there? I know it's been a busy two days, and day three promises to be just as interesting and just as informative as the first two days. Let me also express my personal appreciation for you showing up this morning. You know, day three, people are going, "Gee, should I go to the keynote? What's Justin going to say?" So I'm really glad you're here.

You know, I'm often asked how I choose my keynote topics. And in fact the PR department at Intel seems somewhat annoyed with the fact that regardless of what the messages are for any IDF, I tend to follow my own drummer here. Well, I think the observation I would make, having done this for I guess this is my seventh IDF. I may be off. In engineering terms, plus or minus one. But the topics tend to choose me. You know, it's not like I'm sitting here going, "Okay, what would really be a great IDF topic?" They sort of find me.

And that's so true for today's keynote. It's one that I have been thinking about for at least the last three years. And I get really excited, and I start checking with all the different teams and all the different research efforts. You know, "Are we ready with this? Are we ready with this? Are we ready with that?" And inevitably somebody will go, "No ready." And I kind of have a rule that I set after an incident at IDF a decade ago. You're going to see that particular incident. I set a rule that said I wasn't going to give a keynote where I just predicted some technology or forecasted some technology. Dammit, if I can't come here on this stage in this hall and show you that technology actually working, not going to talk about it. And that was the case with today's topic. All right?

So they're ready now, so we can have the keynote. It's great, and it's really exciting, and you'll understand why here in just a moment. I think it's pretty clear now that with mobile computing, the mobile computing revolution, all these new mobile devices, tablets, phones, Ultrabooks, all of this stuff, that in the future, everything that computes will connect – connect to one another, connect to the Internet. Everything will be connected, and increasingly always connected. And we'll talk more about being always connected and the challenges that come with that.

And I'm really talking about everything – a lot of talk about Internet of things. That's not the topic for today's keynote. I'm saving that topic until that technology is ready too. We'll talk about that maybe in a year or two. But literally from the simplest embedded sensor with a tiny processor, to literally the most sophisticated high-end supercomputer you can imagine, I think nearly all of them will connect in one way or another wirelessly. So that really has set the bar for us at Intel Labs in terms of where we need to focus. And we have a very, very strong focus on many aspects of mobility, and particularly the communication aspects of mobility, which we're going to talk about today – communicating wirelessly.

So it came as quite a challenge for us when one of our executives was onstage ten years ago, and started talking about a capability, a technology that really none of us knew how to do, how to achieve it. And as we sat there listening to the keynote, we were reminder of this classic Dilbert carton, where Dilbert gets the specs, which require the speed of light to be violated, gravity to be neutralized, human cloning to be practical. And if you're a developer, you know exactly what I'm talking about. You've lived this cartoon probably more than once in your career, where you're suddenly asked to do the impossible.

It turns out that our quest in wireless began right here, probably on this stage. I don't know if they were in this hall or not, but probably on this stage ten years ago. And I want to play that little video clip for you because you're going to hear it for yourself. I mean some of you may have actually been here for this.

[Video begins]

Male Voice: Imagine every chip set that we build would come with a radio. Imagine every processor that we might build comes with an integrated radio as part of it. And when I say radio free Intel, this is what I mean. That we literally could get to the point where radios are integrated with every product that we build, able to operate across PAN, across LAN, across WAN environments, seamlessly roaming, connecting to all of them.

[Video ends]

Justin Rattner: Okay. Well, when Pat finished describing radio free Intel, you know, we just sat there in disbelief. Especially the free radio part. And I have to tell you, I'll let you in on a little inside dirt at Intel. The mobile wireless group was particularly unhappy with the free part. You know, they said, "We don't give away our radios. We sell them." So Pat Gelsinger got quite a bit of flack about the free part of radio free Intel.

But the technical challenge that Pat gave us to fulfill this dream of radios anywhere and everywhere – maybe I should say it was more of a fantasy at that point than a dream. We sort of get to dreams after the fantasy part – was really too powerful for us to ignore. And thus began our quest to do the impossible. And as I said, that was 2002. It's taken us ten years, but we think we did it. And you're going to see that in just a moment.

So if you know anything about radios, if you're in your 50s or your 60s – I think it's probably less true today – you probably played with radio as a kid. You know, you started with a crystal set. You know, I added a tuba amplifier to my crystal set. It was a very high-end crystal set. And some of you either were or are ham radio operators, so you probably touched radio directly that way. And very clearly, if you grew up with radio in the '50s and '60s, it's an analog technology, like most of what goes on in the real world, in the physical world. Or you could say God preferred continuous functions.

Unfortunately, digital technology, which is what powers most everything we do today, has really surpassed analog technology in most areas – just about every area except in radio. And before you say, "Well, wait a minute. There are digital baseband radios." Yes, and there's software defined radios. I'm talking about the RF section of the radio. The trouble with analog is that it doesn't scale with the technology. In particular, it actually gets worse as the transistors get smaller. So I'm talking about the classic elements for analog radio design, like inductors and resistors and op amps and all the analog machinery that we know today. And if you look at analog integrated electronics, it's actually much better sort of around 100 nanometers in terms of feature size than it is at 10 nanometers, let's say, where we'll be in the middle of the decade or slightly later. When you try to shrink analog elements, as I said, they actually get worse, and what happens – and you can see this here in the graphic – that tends to make the analog portion of the chip be larger relative to the digital portion. The digital portion is shrinking, and the analog portion is sort of locked at the scale required by those analog elements. Of course, the digital part is also, not only is it getting smaller, it's actually getting faster and lower-powered. It's just getting better in every respect.

So, if you want to create a radio of the sort Pat described, you really have to cast it in digital terms. And you know, you know radio, that's not a very obvious thing. It's not obvious that that can actually be done. So, first of all, it shouldn't surprise you that Pat now runs a software company. I guess he decided, you know, can't stand the heat, get out of the kitchen. He runs DMware if you haven't been reading the news of late. So, he's not doing digital radios anymore, but we're doing digital radios. But, the challenge was too great, and you know, like all good engineers, students of technology, we took that problem and we hit the books. And what we realized when we started looking at communication theory, particularly radio theory, is that radio is expressed as a series of equations. It's actually a mathematical problem.

And as soon as you have that, as soon as you make that observation, you have that insight, you go well, if it's a mathematical problem, we can compute that. You know, and all we have to do – I don't want to make it sound trivial – what you have to do is build a computer that can run those equations in real-time. So, it's a very special computer, but it is a computer, and we're using the mathematical foundations of radio to really crack this digital radio challenge.

So with that, I want to bring out Yorgos Palaskas. Come on out, Yorgos. Yorgos is the research leader in the radio integration lab at Intel. Good morning, how you doing?

Yorgos Palaskas: Good.

Justin Rattner: You nervous?

Yorgos Palaskas: A little.

Justin Rattner: Well, they haven't thrown anything at me yet, so there's always a chance, but let's see if we can get through it. So, you know, I should probably warn everybody before we get into this section of the keynote that this, we're going to move very fast through, you know, a lot of very technical stuff. So as they say in television, now is the time to get your geek on, because we're going to be in a geek moment here very, very shortly.

So Yorgos, it might be best, I think, since not everybody here, you know, is a radio expert familiar with radio technology to start with an analog radio, typical today contemporary analog radio design, and just get everybody grounded. Particularly the RF section, because I think we'll all agree that the base band is now for the most part a solved problem in the digital domain.

Yorgos Palaskas: Correct. So, this slide shows a [unintelligible] analog traditional idea. And you can see that it is primarily analog [unintelligible]. It has a lot of analog [unintelligible] that use a lot of [vectors], a lot of analog devices, and that's why it doesn't scale, as we were saying earlier. It also usually takes a lot of tweaking to get it right. It might take multiple [unintelligible] to do it right, and that's why people typically do traditional radios in all their processors when they're mature and stable, and therefore we don't integrate typically with a microprocessors.

Justin Rattner: That's right, today most radios are separate chips, right. And those passes through the [fab] to make the tiny tweaks to the analog portions take a lot of time. I mean, months can go by for a pass. So, we're clearly motivated to do it. So, when you started thinking about how to do this in the digital domain, was it just, okay, we've got this analog block, we just have to build the digital equipment, or did you take a different approach to it?

Yorgos Palaskas: No, actually had to completely basically rethink radio operation from a digital perspective. We wanted to attack this problem as a computational problem, as I said. In the main cases, it was not just converting an analog block to a digital equivalent. In many cases, we really had to invent new things that basically would work optimally with fast transistors and would scale well.

Justin Rattner: So, I mean, you did literally go back to the books. You went back to those fundamental equations and imagined how you would build a special computer to solve them?

Yorgos Palaskas: Something like that, but when we actually had to implement the things, though, we actually were not able to find a lot of things, pretty much nothing.

Justin Rattner: Well, your job is to make the impossible possible, so that's what you do. So, you know, break this down for us. I mean, this was a 10-year quest from the time Pat issued that challenge – he may not have imagined it as a challenge, but it was a challenge for us – you know, how did you, tell us how you went about doing it. Tell us the story. I mean, this is really an exciting moment.

Yorgos Palaskas: Yeah, it was a monumental effort, because as I said, we really had to build this capability from the ground up and really rethink the fundamentals of radio operation. Now, I will discuss a couple of blocks here. We don't have time for too many blocks, but we'll specifically focus on two of these blocks. We'll start with the digital frequency synthesizer, which is one of our best, first important results. So, when we did the digital frequency synthesizer, that was the only integrated frequency synthesizer that could do Wi-Fi and 4F specifications. So, at that point, we realized that we had something. It was digital, and it was the only one that could do it. So, that was an exciting . . .

Justin Rattner: I remember these great internal debates. I remember making a trip to Israel where I got yelled at by engineers of all races, colors, and creeds, you know, you guys are crazy and this will never work, and on and on.

Yorgos Palaskas: [Unintelligible] is so different from what we are used to, but it takes time to actually.

Justin Rattner: Okay, so when we, you know, when we look at the synthesizer, you know, which was this first important block, did that give you confidence that these circuits would scale? This is sort of the fundamental problem with analog.

Yorgos Palaskas: Yes, we had the feeling and the intuition that things would scale, but we actually confirmed those feelings in the process of the years, and you see here a couple of examples from actual fabricated digital synthesizers. And you see going from 90 nanometers in [unintelligible] to 30 nanometers, a dramatic scaling in the area of the digital synthesizer, and also the power dissipation. And we predict that things are going to continue, going to 14 nanometers, for example.

Justin Rattner: Yeah, so this must've been a big boost in confidence for you.

Yorgos Palaskas: Yes.

Justin Rattner: Okay, let's take another block.

Yorgos Palaskas: So, one of the last blocks that came together was the digital phase modulator. This is a key part of our digital transmitter architecture, this different transmitter we developed. So, digital phase modulators, they existed before. In fact, Intel's satellite division . . .

Justin Rattner: Intel Mobile Communication?

Yorgos Palaskas: Intel Mobile Communication, they have the digital phase modulator in a similar part.

Justin Rattner: In a product?

Yorgos Palaskas: They published it recently in a product. But, that works for, it is [targeted] for 3G operation. So, it is smaller, but it's up to five megahertz.

Justin Rattner: [For your] channels.

Yorgos Palaskas: Yeah. So, in our case, we had to deal with a much wider [unintelligible] requirements for Wi-Fi. So, our system was designed up to 40 megahertz Wi-Fi operation. So, we really had to again come up with a lot of . . .

Justin Rattner: So, I recall asking this question in review once. I said, this sounds like it's going to require some new mathematics not found in the textbooks.

Yorgos Palaskas: It requires some very creative mathematical manipulation of [unintelligible]. It was definitely . . .

Justin Rattner: You put it so deftly, mathematical manipulation. All mathematics is manipulation. Okay, so you know, now we've got all the blocks. Can we build a radio?

Yorgos Palaskas: Yes. We actually also have the complete radio.

Justin Rattner: Okay, I can bring it up, yeah.

Yorgos Palaskas: And this is the first complete digital Wi-Fi radio that we have developed. It operates up to the full, the more demanding Wi-Fi requirements, the 40 megahertz bandwidth. It has been implemented in a 32 nanometer Intel process. The power efficiency is on par with the best analog designs that we could find in the literature. And we know that our radio will keep improving because [this all] flip-flops, transistors, switches – while we know that other radios are actually not going to improve, the analog radios. So we expect that this is going to become better and better moving forward. It is already on par.

Justin Rattner: I've waited 10 years for this moment, Yorgos. Can we see the radio?

Yorgos Palaskas: Yes.

Justin Rattner: Okay.

Yorgos Palaskas: This is our transceiver board here that you see. And we're zooming into the actual digital radio here in the middle.

Justin Rattner: All of these are test connectors. That's the chip itself.

Yorgos Palaskas: Yeah, the black one in the middle. So that [unintelligible] digital transmitter comes here to this, through this cable to this transmitter antenna. It's been transmitted to the other side, the receiver antenna, our receiver. This is the same transceiver, digital transceiver operating under safe mode on this side.

Justin Rattner: We can think of this as sort of the access point, and that's the client end at that –

Yorgos Palaskas: Transmitter receiver.

Justin Rattner: Yeah, okay. But same chip.

Yorgos Palaskas: Yeah.

Justin Rattner: Okay.

Yorgos Palaskas: And these are the configurations here.

Justin Rattner: Right. Okay. I know you've got all this test gear here, but it's best in the interest of time that we just look at the application you have running.

Yorgos Palaskas: Yeah. We have –

Justin Rattner: Trust me, this meets all of the specifications that we've been talking about. But let's just look at the app running.

Yorgos Palaskas: We have a video streaming right now, and we'll see good quality of data going through our system over this entire link.

Justin Rattner: Okay. But how do we know you aren't faking this?

Yorgos Palaskas: You can try to block the antenna signal.

Justin Rattner: Okay. Come over here.

Yorgos Palaskas: Now it is going to take a few cycles because there is some buffering and encoding in the system [to actually stop].

Justin Rattner: I have transparent hands.

Yorgos Palaskas: It is very good encoder. It works well. [Laughter]

Justin Rattner: All right. Two hands. [Laughter]

Yorgos Palaskas: Now you have to wait. This is the slow part, Justin.

Justin Rattner: Hey, it stopped. Okay. I wouldn't work so hard if we were faking it. Okay. [Applause] So there you go.

Yorgos Palaskas: [Unintelligible].

Justin Rattner: First all-digital Wi-Fi radio, 10 years from Pat's challenge to today. Now, actually I should check that. Pat didn't just ask for a digital radio; he asked for a digital radio that you could integrate on a chip with other compute elements – blocks, memories, registers, I/O. What about [a poll check]?

Yorgos Palaskas: [You] still remember the – okay. So here it is for you.

Justin Rattner: So this is actually Radio Free Intel. Well, Not-so-free Radio Intel.

Yorgos Palaskas: This is a research test we developed. It's called Rosepoint. It is 1 32-nanometer, the first [SOC] that integrates a Wi-Fi radio together with two Atom cores and all other required digital circuits.

Justin Rattner: Two Atom cores and the radio. [Applause] Okay. Pat, wherever you are, Radio Not-so-free Intel. All right. That's just terrific. Terrific, Yorgos. And really, you kind of get choked up here if I let myself go, it's so great to see it. So what's left to do? What's next on the agenda?

Yorgos Palaskas: Primarily taking this radio to productization. We're developing the technology pieces, so other people have to –

Justin Rattner: While it's nearly completely digital, aren't there one or two analog elements still left?

Yorgos Palaskas: There are a couple of small things here and there. There are a couple of things that will always be analog. For example, we [unintelligible] the power amplifier. This is an analog signal. But there are a couple more things that we can keep improving, for example, we can put a little more data conversion, a few more bits [unintelligible] that can further digitize the receiver.

Justin Rattner: Okay. Sounds good. Well, thanks, Yorgos. It's just been great.

Yorgos Palaskas: Sure.

Justin Rattner: Thanks for being with us today. [Applause] Yorgos Palaskas, everybody. Yeah. It's not every keynote we get the opportunity to share with you something that really hasn't been done before.

So let's continue with our quest for the next-generation of wireless and wireless computing. Here's a very familiar scene, familiar to all of you. I'm sure most of you have a bag, you know, your computer case or a little zipper bag which has all the cables you need. You never know where you are, you know, "Do I need that USB? Do I need that HDMI?" You know, whatever. So it's fair to ask the question, "Here we are in the new century in the second decade. Are we ever going to get rid of those wires?" Let's watch this video.

[Video begins]

Male Voice: WiGig is an integral part of my digital life. It makes everyday tasks much easier. On the way to work, I stopped by the movie kiosk to grab a flick for tonight. I can download the movie right to my phone and take it with me. A few seconds later, I'm on my way.

My tablet is the center of my high-performance workspace. When I get to the office, all of my devices connect automatically and I'm ready to get to work. My external devices become an extension of my tablet.

All over the office, we found a variety of applications for WiGig. It's especially useful for group projects. I can automatically connect to my colleagues' devices through the corporate network, so I can focus on the task at hand, not the technology.

Whether transforming my thin computing device into a work station, simplifying interaction with my coworkers, or enabling group participation, WiGig enhances the productivity of any workplace. And WiGig is great at home, too. Remember that movie

I downloaded this morning? I can play it right from my phone, and the quality is outstanding.

[Video ends]

Justin Rattner: All right. Well, to tell us more about WiGig and all the exciting progress that's being made, I've invited Ali Sadri, who's Chairman and President of the WiGig Alliance, to join me today onstage. Good to see you, Ali. [Applause]

Ali Sadri: Hi, Justin. Good to meet you.

Justin Rattner: You look great. Very San Francisco.

Ali Sadri: Thank you very much. I borrowed this.

Justin Rattner: Yeah, yeah. I like it. I like it. So I think the video did a great job really of kind of explaining all of the wonderful things that you can do with WiGig. But you're running the show, so to speak. So where are we particularly with the development of the standard?

Ali Sadri: Absolutely. Thank you for the opportunity.

Justin Rattner: My pleasure.

Ali Sadri: Thanks for the audience listening to us. The work has been done by multiple member companies. I have the pleasure to be able to help the group to make it happen. Since the inception of the WiGig Alliance back in 2009, we always had the vision that we needed to develop a very high throughput wireless technology that is capable of doing things beyond what Wi-Fi can do, for example, transmitting to the televisions, to the monitors, unleashing all the devices from – cables, for example, [PC Express], US [SDI] –

Justin Rattner: So no more bags. [Laughs]

Ali Sadri: Hopefully. Maybe you can put something else in your bags, but not necessarily [crosstalk].

Justin Rattner: I saw Apple introduce a new connection or a new cable yesterday.

Ali Sadri: Oh yeah. You can put that one in there. That's right.

Justin Rattner: Okay.

Ali Sadri: So we were able in just a very short time, in a couple of years, to bring the concept that you just saw in the video that we produced – before this becomes a reality, to today, that will be very close to a reality.

Justin Rattner: Yeah. And it does a good job of that. Now, you're up in the 60-gigahertz part of the spectrum, millimeter waves. That's kind of the new regime. And WiGig is not the first occupant of that part of the spectrum. So why do we need WiGig? Why can't we go with one of the existing 60-gigahertz [crosstalk]?

Ali Sadri: Very good question. I've been asked that question many, many times, so I'm prepared for that.

Justin Rattner: I'm sure. Okay.

Ali Sadri: As you know, the legacy Wi-Fi devices operating in 2.4 gigahertz or 5 gigahertz. These frequencies by definition are very high-consumed. So the interference is potentially very high. For high-demanding applications such as video, you know, synchronization [has] very high throughput, and multi-user applications in the office environment, we need to move into a frequency where there's a lot less interference. And that's why we chose 60-gigahertz.

Justin Rattner: Okay.

Ali Sadri: But we also looked at a multitude of applications, as I mentioned, for PCCE, handheld, handset, you know, ultrabooks and so forth.

Justin Rattner: Right.

Ali Sadri: So we designed this new standard based on that vision.

Justin Rattner: So that's very reminiscent of Wi-Fi itself.

Ali Sadri: It is. As a matter of fact, if you look at some of the board members of the contributors that we have – if we move on with the slide –

Justin Rattner: Sorry, my fault.

Ali Sadri: – you have the same potential characters and companies involved. And that's why we actually designed our system based on Wi-Fi so we're backward compatible to the Wi-Fi system. It does the applications that Wi-Fi can do today, plus additional applications that we always discussed.

Justin Rattner: Okay. So you've been working since 2009. Now it's 2012. When are we going to see something working?

Ali Sadri: We just did the production, so that's it. Just kidding.

Justin Rattner: You're in production? Oh boy. I'm sorry, not in this keynote.

Ali Sadri: Video production. So if we move into the next slide, we actually turned that vision a couple of years ago into reality. You can actually see it, hopefully [in life] –

Justin Rattner: Can we do that?

Ali Sadri: Absolutely. I'd like to ask my friend and colleague, George, to help us with this demo. So what we see here, there are two monitors that are connected to the docking station.

Justin Rattner: Wow, two monitors.

Ali Sadri: Two monitors.

Justin Rattner: High def?

Ali Sadri: High definition 1080p. And there's an external hard drive and a printer all connected to the docking station device. And also – let's wake this up first. It's not waking up? Oh, this is off. It's supposed to be on.

Justin Rattner: Wake up the Ultrabook.

Ali Sadri: You know what? It was on since last night. I think it ran out of battery. What it's supposed to do –

Justin Rattner: No, really? Well, it's live television, everybody.

Ali Sadri: It's live television. It's not connected? So I can explain. So what it was supposed to do –

Justin Rattner: Oh, he's got an image up now.

Ali Sadri: Yeah, he's waking it up. I think the laptop went to sleep. The device will connect to the docking station. It's not paired yet. When it connects, actually a light comes up and it's paired. And as I said, all the peripherals and the external drive and the monitors are connected to the docking station. The content that we saved earlier was supposed to show – hopefully we can get it done in time – is from the hard drive, the content is going to be rendered into the laptop, and then it will translate back into the monitors.

Justin Rattner: I see. Okay, right.

Ali Sadri: So at the same time – again, we were supposed to have it done but not done here – we have Wi-Fi connections. So not only we have Wi-Fi connections, we also have a complete unleashed environment for the laptop. And hopefully it will come up live. One of the main things to look at is the size of the laptop in the back is actually constrained by the connector.

Justin Rattner: Sure. Yeah, you can see the connector.

Ali Sadri: So even if we were able to remove all the connectors, we can actually shrink that size in the back as well. That is going to hopefully give you a better and sexier look of laptop.

Justin Rattner: Super. So when do you think we're going to see this actually in the market?

Ali Sadri: Well, this was definitely a prototype of the wireless docking station. I assure you it was working yesterday.

Justin Rattner: I did see it, yes.

Ali Sadri: It always happens when you have a live demo. We have several member companies developing the prototype. And we went through two rounds of the test events, one in 2011 and one in 2012. We are set for the next event later on this year and the beginning of next year. Our certification program is going to start mid next year, and hopefully – are we connected? It seems like we're connected.

Justin Rattner: Oh, fantastic.

Ali Sadri: So our certification program is about to launch in mid-2013. So hopefully the certified products will be shipped shortly after.

Justin Rattner: Okay. There they go. Okay.

Ali Sadri: We're connected.

Justin Rattner: I think that deserves a round of applause.

Ali Sadri: Okay, so the hard drive is actually now disconnected. So we're going to turn on the hard drive, and the system is going to detect it. The Ultrabook is going to detect the external hard drive. And it pops a window. And we're going to go and choose the video content. Okay, we'll make it larger. Here we go. So the video is being, again, rendered from the hard drive through this Ultrabook, and goes there. And again, there's no wire whatsoever.

Justin Rattner: And soon no connectors.

Ali Sadri: No connectors. And to make sure that this is the system and we're not putting –

Justin Rattner: No, I don't think anybody believes you're faking it at this point.

Ali Sadri: As soon as I opened it up, it actually disconnected.

Justin Rattner: Okay. Oh, there it goes.

Ali Sadri: There we go.

Justin Rattner: Right. All right.

Ali Sadri: So this is probably the best thing can happen to all of us, that we get rid of all the unwanted cables.

Justin Rattner: So when?

Ali Sadri: Hopefully, as I mentioned, right after certification. Mid next year, we will see more and more devices in the market. And I look for the day that –

Justin Rattner: Now I know why you're in charge of WiGig. Thanks, Ali. Ali Sadri, everyone. Thank you, George. See, I was worried about the Wi-Fi radio. I figured the WiGig, that's standard, so everything like that. So you never really know. Okay, we've got to rock on. We've got a lot of ground to cover and not much time.

When you look at mobility and wireless technology and everything being connected, you start to realize that there are really a large number of tradeoffs. And one of those tradeoffs is this notion of convenient, right? How convenient is the device to use? How does it fit into your daily life? Versus battery life. And it turns out that this tension is actually quite common. You know, people want their devices to appear to be awake, appear to be always available. But as we all know, that's going to reduce the battery life, and in some cases dramatically reduce it, particularly if it involves wireless communication. If you're like me, I'm constantly managing the radios on my phone. I'm on the plane, the radios are off. I land the plane, the radios come back on. I'm not in a Wi-Fi zone. Wi-Fi's off. Oh, I don't need location, I'll turn off the GPS. You know the drill. So all of those contribute to reduces battery life, or some loss of convenience.

So if we look at what's been going on, Intel has already brought some of the technology required to really solve this problem, sort of eliminate this tradeoff between convenience and battery life. And Paul Otellini, our CEO, in his keynote at Consumer Electronics CES I guess in January of this year, showed this Dell notebook. This is an XPS 13 Ultrabook right there, which is equipped with Intel's Smart Connect Technology. And what's really going on here, to make it simple and brief, is that the behavior you associate with smartphones and tablets, which is they always seem to be receiving the email and the tweets and other social media traffic. We've brought that capability to the Ultrabook platform. So when traffic comes in, the platform wakes up and receives that traffic, and then immediately goes back to sleep. But of course, we're on the lab side. This is already in the marketplace. What are we doing to improve over the Smart Connect Technology. And to explain what's up, I've invited Intel Labs Principal Engineer, Charlie Tai, to join me onstage. Come on out, Charlie.

Charlie Tai: Hi Justin.

Justin Rattner: Welcome, welcome, welcome.

Charlie Tai: Thank you.

Justin Rattner: Okay. Thank you. So I was talking about this tradeoff between efficiency and battery life. How are we moving forward? How are we going beyond Smart Connect?

Charlie Tai: Well, as you know, Intel Labs has deep expertise in energy [efficiency], from the transistor circuit level all the way up to the system level. The Intel Smart Connection Technology you mentioned earlier is a good example of that, which we helped develop a couple years back. And now we are expanding our research to not only support standby but also active device use cases.

Justin Rattner: Okay, so you know, what is it that's going on that we really need to address?

Charlie Tai: Well, as you know, we are seeing more and more of internet-driven call-based use cases, however internet traffic tends to be erratic, right? And it's hard to manage, which leads to all kinds of inefficiencies here, and we would like to tackle that problem.

Justin Rattner: Okay, so can you give me an example of what you're doing?

Charlie Tai: Yeah, actually, we're going to like to show you a couple of examples for our new technology code, code-named Spring Meadow. Maybe the best way to explain it is to actually show you the demo. You want to talk a look?

Justin Rattner: Okay, let's go, yeah.

Charlie Tai: Okay, so there are two cases. The first one is for idle, right, where a user kind of knows, sitting in front of an Ultrabook reading an [idea press] articles. But in that scenario, I guess, this is relatively idle, right, because the user's not really doing anything. However, even in that case, there's tons of background traffic which keeps bombarding your system in order to maintain network connectivity, to support service discovery, and so on, right? So, as you can see on the lab screen, right, every time a new pack arrives, and heat assistance, it basically will light up the whole house.

Justin Rattner: Platform powers up.

Charlie Tai: Yeah, and sometimes takes some time to turn off the lights. So, obviously, you cannot, you know, consume a little bit of energy here. And on the right-hand side, with our Spring Meadow technology, and we actually analyze the traffic in order to identify the most important packet.

Justin Rattner: And you're analyzing it out at the [nick], not in the process?

Charlie Tai: Yeah, we're analyzing that in the nick, and we can remove all the unnecessary ones which are either not [interesting] to me or can be deferred for later processing, for example. By doing that, as you can see in the middle, we can dramatically reduce the platform activities. Hence, you know, save a lot of CPU power.

Justin Rattner: And you mentioned a second example of this. Want to show us that?

Charlie Tai: Yes, let me show you the second one. Let me get it started, there we go. So, the second one is for active, right, where a user accessing content in the cloud by downloading, streaming, audio, video, image, or any other type of traffic. So, in that case, our Spring

Meadow technology is actually proactively managing the network flow in order to minimize the interruption to the CPU.

Justin Rattner: So, I mean, give me a sense of how much power the Spring Meadow technology has.

Charlie Tai: So, let me show you. So first of all, as you can see, you know, you can download the same set of files in the same amount of time. So basically, there's no impact on the performance, as you can see from both sides.

Justin Rattner: They look right in-sync to me.

Charlie Tai: And in terms of your very good question about how much power you're going to save, as you can see from the middle screen, actually you can save about half of the CPU power.

Justin Rattner: Wow, half of the CPU power?

Charlie Tai: Yes, yeah.

Justin Rattner: Wow, that's amazing, that is amazing. So you know, it's clear this has obvious benefits for Ultrabooks, but it might even be more beneficial in phones and tablets where you don't have as much battery to play with.

Charlie Tai: Absolutely, you are totally right, yes. Yeah, those devices tend to have a smaller battery, and this technology would be even more crucial for that kind of devices.

Justin Rattner: Okay, so you know, with all this technology, all these advancements, you know, how soon do you think we're going to see this in the market?

Charlie Tai: Well, it's, you know, still a prototype, but it's getting better every day, so stay tuned.

Justin Rattner: Okay, all right, I'll wait. Hey, thanks for showing it to us.

Charlie Tai: Thank you.

Justin Rattner: Charlie Tai and the Spring Meadow technology. I want that, I need that, very important technology.

All right, so let's keep rocking here. Another thing that has become very popular in this age of internet-connected devices and wireless technology and so forth is streaming video over the net. I think the proper term is, over the top. And that creates yet another one of these trade-offs. Do you optimize for more video streams, or do you optimize for higher quality, better video experience? So, we've been looking at this, you know, at this tension between, you know, between these two technologies, and I'd like to – before I get there, actually, let me give you sort of a quantitative assessment of what's going on. You can see on the left-hand side how internet traffic just continues to explode. There's no abatement in here, and a lot of that is being driven by internet, and you know, the amount of consumer internet traffic has really come to dominate that equation. So, this is clearly an area that can use some help, can use some better technology to manage the available bandwidth while the other technologies advance, and we'll talk about those in a little bit, to just increase the total absolute bandwidth.

To give us some insight into this, I've invited Chris Neisinger from Verizon and our own Jeff Foerster from Intel Labs to join me here onstage and talk to us about some very exciting work in Video Aware Wireless Networks. Gentlemen? You snuck up on me, don't do that again. Hi, Chris.

Chris Neisinger: Hi.

Justin Rattner: Thanks for being here, and Jeff, great to see you. Okay, gentlemen, so, we're working on this problem of, is it more channels or is it better channels? Do you want to, Chris, just kind of give us an idea of the Verizon perspective here?

Chris Neisinger: Well, from the perspective of Verizon and other service providers, you know, the fact that our customers are loving video and want to use our systems is very good. We have a huge demand for our products, and so my job is to look at the supply side, as you mentioned, try to find ways to keep the high quality that our customers expect while we're dealing with limited resources, limited spectrum.

Justin Rattner: There's only so much spectrum, only so much bandwidth available, absolutely. And Jeff, how do you see the challenge?

Jeff Foerster: Yeah, so I mean, one of the things that we're trying to do is work with industry partners like Verizon and Cisco on a university research program that we call Video Aware Wireless Networks. And we're funding five leading universities in this space, looking at really fundamentally new ways to manage video within the network. And the research spans a lot of different topics as part of the end to end system. I think the major one is really looking at and understanding video quality from a human visual systems perspective and really looking at, how can we track that video quality within the network, especially on mobile devices, which becomes critical to be able to share that information with the network as well.

Justin Rattner: So Chris, again, your perspective on this. You're part of this big research effort. What are you looking for to come out of it?

Chris Neisinger: Well first, one of the things that's really terrific about this collaboration with Intel, Cisco, and the universities is that we're trying to shorten the time from ideas to getting into the network, and it's a terrific advantage here. So, one of the things that [Von] has done that we're particularly interested in is ways of measuring the video quality. We have very good ways of measuring voice quality and using that in our network to adapt the network to get the maximum capacity out of the network, but in video, we're a little premature on that. We don't have all the tools we need to really maximize the capacity. So, the video quality awareness is very important to us.

Justin Rattner: Well Jeff, I'll toss it back to you. So, can you give us some sense of the ideas that are coming up from this collaboration?

Jeff Foerster: Yeah, so internally, one of the things we're doing is actually partnering with our user experience group within Intel Labs to really understand the fundamentals of the factors that impact video quality. And we're, so they've actually done a number of subjective studies asking real people, you know, how is video quality impacted by a number of factors, like what kind of device you're watching it on, whether it's a phone or tablet or TV, what kind of content you're watching, whether it's a newscast or sports or TV show, what kind of data rates, how does the data rate impact the video quality, and also how do the different ways we can compress the content, how does that impact the video quality. So, there's actually a tech talk later this morning going through the details of that subjective study. But the basic outcome is, you know, for a fixed data rate, the video quality varies significantly depending upon content and the device you're watching. So, you really want to take into account this kind of information when you try to optimize the video over the network.

Justin Rattner: Can you give us an example of it?

Jeff Foerster: Yeah, so we actually put together a demo, a fairly simple scenario where we have 10 devices, four tablets and six phones, all streaming different content over a wireless network, over a constrained, let's say 10 megabit per second pipe. And so, the initial, the current way of managing that video is kind of dividing up that capacity based upon throughput, so giving everybody about a one megabit per second length, and then seeing what the quality results in. And based upon our subjective studies, about half of the users won't be satisfied with that kind of quality of experience, and so I think we can do a lot better than that. And so, what we do is now incorporate some knowledge about the quality of experience, some knowledge about the type of content that we're streaming, the devices that we're streaming to. And when we take that into account in the allocation of the resources for that fixed channel, we can actually satisfy all the users within this network. And we have kind of one example with two tablets, two of the tablets in this particular case. And we can show the videos here, where the tablet on the left is kind of based upon this throughput managed network, whereas the tablet on the right is based upon this quality of experience managed network, where we take a lot more intelligence within the management of that network. And you can clearly see, hopefully it's clear that there's a big quality difference between these two tablets.

Justin Rattner: So Chris, what do you see as the challenges of actually doing this in the network?

Chris Neisinger: Well, there's technical challenges which we're addressing, but also this perception of fairness and openness.

Justin Rattner: Right. That's been a big deal.

Chris Neisinger: Yeah. As we adapt the content, we're really concerned about being fair to maximize quality to everyone.

Justin Rattner: Right. But this gives you a way of sort of getting the quality, and making more efficient use of the bandwidth, right?

Chris Neisinger: Yeah, absolutely. As you said before, the spectrum's a limited resource. There's a high demand for this, and building out the network is important. We invest in our network. But we also have to do it in an efficient way to maximize the –

Justin Rattner: And your thoughts here, Jeff?

Jeff Foerster: Yeah. So I think we need to kind of build the tools that analyze the video content, that extract features and information about that video content to be able to share with network operators, like Verizon, so they can better manage the network. But it will take kind of a collaboration with the other ecosystem partners here to really optimize the video. So I think that's the challenge is to work with the other industry partners to actually realize some of these ideas.

Justin Rattner: Well, it sounds like very promising research. The best to both of you, and of course, the whole research consortium. And we'll look forward to a day when we can all get better video and more of it. Thank you so much, Chris. Chris Neisinger, everybody from Verizon. Jeff Foerster of Intel Labs.

Okay. I said there were a series of tensions. The next one I want to address and I want to talk to is the tension between ease of use and increased security. And I think you're all well aware of this. I mean people have been turning off the antivirus, anti-malware protection for years just because they think it's going to make the system more responsive, run a little bit faster. You know, we believe it doesn't have to be that way, and it is possible to bring some new technology to bear to solve the problem. So my next guest, who's going to talk to us about the efforts we're making in this area in Intel Labs, is Sridhar Iyengar from the Security Research Lab, part of Intel Labs.

Sridhar Iyengar: Hi, Justin. How are you doing?

Justin Rattner: I'll get myself in the right place because everybody keeps sneaking up on me. Thank you. All right. So ease of use, more security. Talk to me about this.

Sridhar Iyengar: Well, let me start by asking you a question, Justin, if I may.

Justin Rattner: No, no, wait a minute. I ask all the questions.

Sridhar Iyengar: Well, bear with me for this time around. How many passwords do you use, Justin?

Justin Rattner: I've lost count. You know, I have my Intel passwords and my personal passwords and my super personal passwords.

Sridhar Iyengar: And how many of them are the same?

Justin Rattner: Well, they're variations on a theme.

Sridhar Iyengar: All right. Well, see, that's the problem with passwords. We use too many of them, the rules are complex, and they differ for differing Web sites. So what's a user to do? They either use simple passwords or reuse them, right? So they sacrifice security for ease of use.

Justin Rattner: Okay. Is there some way around this predicament?

Sridhar Iyengar: Well, essentially, in the case of wireless devices, it gets worse. But there is a way around it, and biometrics is one option. It's a right option because replacing what you know – passwords – with what you are – fingerprints – it's an ease of use issue, it's harder to spoof, Justin, and you're not likely to forget your fingerprints anytime soon.

Justin Rattner: But, you know, I have a fingerprint reader. I don't really use it. It's kind of unreliable and everything. So is there something better than that?

Sridhar Iyengar: So from a security spectrum, because you leave fingerprints on your screen, and it's like writing your password and sticking on a device. So yes, there is something better than that. And in our research, what we have is something called client based authentication. And the way it works is the user walks to the device and locally authenticates to the device using whatever factors are available, including biometrics.

Justin Rattner: Okay. So now I'm authenticated, but there's more to it than that.

Sridhar Iyengar: Indeed, absolutely. What the device then does is it securely asserts your presence to the remote party. It does a secure assertion of the user's credential to the service provider. So we don't send their fingerprints over, naturally. We just send a secure assertion to prove that you are who you claim to be.

Justin Rattner: And there's no passwords here. It's not like I've got a password wallet or something.

Sridhar Iyengar: Not at all. It sends an assertion to the service provider that proves you claim who you are.

Justin Rattner: Okay. Can we see it in action?

Sridhar Iyengar: We can certainly see it in action. Come on over.

Justin Rattner: Okay. And Jason over here. Hi there. Okay, you want to take us through it?

Jason: Yeah. So what we have here is a tablet that we've added a biometric sensor to. So when I want to actually –

Justin Rattner: What's this? [Faring]? You've got –

Jason: Yeah, this thing is kind of hanging out, outside here. So when I want to actually access my device, all I have to do is wave my hand in front of it.

Justin Rattner: Oh, it's unlocked. Wow, that's cool. You just kind of – can you play music on it too? Okay, so now we're authenticated to the tablet. Now how do we get connected?

Jason: So what we want is that now that the device knows who you are, it can actually tell all your services your identity on your behalf so you don't have to do any more work. So I'm going to actually go and check my bank account real quick here.

Justin Rattner: Okay. Clicking on Acme Bank up there.

Jason: You see it kind of flashes, and the service provider is asking my device who's there –

Justin Rattner: So essentially you go right past that home page, no username, no password. You're into the account. You can make that quick transaction. You can see how this is going to be so important on mobile devices, right? You just want to get in, get out, and so forth. In fact, getting out is an interesting problem. Have you addressed the getting out problem as well? You know, you're at Starbucks, and a buddy walks in, or your coffee order is ready. And you just want to put them device down and walk over, but somebody could just grab it, and then you're in a big trouble. So have you thought about the backend of the problem?

Jason: Absolutely. So one of the other sensors that this device has on it is an accelerometer. And we can use that device actually to tell whether the device is still being held and used by you. So if you want to go pick up your coffee, you're just going to go ahead and set it down.

Justin Rattner: And it locks immediately.

Jason: It protects your session, takes care of it for you.

Justin Rattner: Okay. So the accelerometer is sensing that it stopped moving. You pick it up again, right? I've got my coffee now. I didn't actually get through that transaction – the banking, the stock, whatever it is. Can you show me how you get back in?

Jason: Yeah. You just say hi again, and you log right back in where you were, and pick up where you left off.

Justin Rattner: Wow. That just takes a few seconds. I think a round of applause is very appropriate. Well, I don't know. I won't bring up the movie. I've got a scene in a movie, but we can talk about that after the keynote this morning. So this really seems like cool technology. What's next? Where we do go from here? And perhaps more to the point, do we have to invent new standards? I mean is this going to be a multiyear effort?

Sridhar Iyengar: No, absolutely not, Justin. I think the good news is the client talks to the server using existing protocols, like [SAMA], Open ID, Kerberos, and things like that. So we plan to work with service providers to take full advantage of this using open standards. So no new standards. That's the good news.

Justin Rattner: And in terms of new silicon –

Sridhar Iyengar: No new silicon.

Justin Rattner: I guess the big thing, you know, in the labs we don't have the kind of budget to do the packaging work, so we've got to get that sensor integrated.

Sridhar Iyengar: Integrated into the device.

Justin Rattner: But you get that, the software standards are in place, and we're good to go. All right. So can we say great ease of use and improved security?

Sridhar Iyengar: I think that's the idea.

Justin Rattner: All right. Thank you, gentlemen. Jason, good job. So you're probably thinking of that movie scene right now where somebody's hand gets removed. Sorry about that. Hey, I haven't mentioned the word condom in my entire talk. Give me a break. Sheesh. Well, I guess now I have. All right. In any case, the last thing we want to talk about this morning is going to take us into the cloud. And what we want to talk about is building the next generation of wireless infrastructure. You know, we've talked about all the other bits and pieces. How do we move video? How do we make things secure? How do we get long battery life? But is there an opportunity to rethink the fundamentals of the radio access network to deliver more bandwidth, better experience, greater ease of use? Let's watch the video.

[Video begins]

Male Voice: As the mobile computing revolution continues to drive the explosion of data creation, consumption, and sharing .Wireless operators worldwide are exploring new ways to approach the build-out of their next-generation broadband network infrastructures.

Historically, those networks have been based on proprietary hardware that is both expensive and difficult to maintain and upgrade, hardware platforms that require so much power and cooling that the cost of electricity can represent over 40 percent of ongoing operating costs.

In the current approach, base stations have no ability to share capability with other sales sites on the network. Each base station must be designed to support maximum loads during peak times for the particular cell area it supports. The result is an infrastructure model in which valuable network intelligence and computing resources are dramatically underutilized.

But now, by combining open architecture servers with fiber optic backhaul, operators can centralize network intelligence in the cloud, utilizing a virtualized environment where resources can be allocated to different base stations on demand when and where they're needed.

In the C-RAN architecture, radio protocols are implemented using software-defined radios in the cloud, rather than hardware at each tower. And research projects have already shown that a complete LTE-based station stack can be implemented in software running on a modern IA processor. This software-defined radio design offers the versatility to add new radio protocols over time as they're introduced with a simple software upgrade.

With software-defined radio running on hundreds or even thousands of IA-based servers, a modern datacenter can replace much of the computing hardware traditionally installed in base station towers scattered across a wide-area cellular network.

The result is a total C-RAN-based solution that not only reduces both capital and operating expenses, but allows operators to bring new services online faster and with greater flexibility than ever before, ensuring the future potential of the mobile computing revolution for generations to come.

[Video ends] [Applause]

Justin Rattner: Yeah, it's a pretty good video. I was so impressed with it myself, just a few pencil sketches and here you go.

Well, if you were here last year and you were paying attention, we actually showed sort of the first piece of this idea of moving the heart of the radio access network to the cloud. We created a complete LTE base station stack that was running on I think a Sandy Bridge second-generation Intel core processor, just a PC. So that demonstrated that it was possible to replace those custom and proprietary mostly hardware radios at the cell sites with software running on standard server hardware in the datacenter.

You may also recall that the work was a collaboration between Intel – in particular, Intel Labs China and Beijing – and China Mobile Research Institute. And I've invited Chih-Lin I from China Mobile Research, where she's Chief Scientist, to join me this morning to talk more about the ongoing effort. Chih-Lin? There she is. [Applause] Let's give her a big welcome. Come on out.

Thank you so much for coming. I've made that trip from Beijing to San Francisco many times. That's a long flight. I sort of stagger off the plane. So grateful that you're here.

Chih-Lin I: Thank you, Justin. And on behalf of China Mobile, the largest wireless service provider in the world, I believe –

Justin Rattner: By a fair measure, I think, yeah.

Chih-Lin I: – I'm really delighted to be here. And it's been a wonderful experience collaborating with your labs over the past few years. And we are very happy with the progress to date.

Justin Rattner: Yeah. I'm stunned, I guess, when I hear the statistics for China Mobile when it comes to towers, right?

Chih-Lin I: The scale.

Justin Rattner: You're 2G, 3G, getting up there towards –?

Chih-Lin I: Yes. We have about close to a million base stations just for 2G and 3G.

Justin Rattner: I just a headache just thinking about it.

Chih-Lin I: It's more than 900,000 base stations.

Justin Rattner: Wow, wow. That's really amazing. So you're obviously part of the research project. Tell us more about your interest in Cloud-RAN.

Chih-Lin I: Sure. The thing is even though we have that many base stations already, some of you may know that our CEO had announced recently the rollout plan for 4G, the LTE.

Justin Rattner: 4G, right, right.

Chih-Lin I: The second is that we actually – even though we have close to a million base stations – we have 680 million subscribers. So that's not quite enough, right? So we are planning to deploy 20,000 LTE base stations by the end of this year, which is a small start, and 200,000 next year, and then 350,000 by the end of 2014. So that, of course, brings the question of how efficient and cost-effective we can make our RAN going forward.

Justin Rattner: Right, yeah. I'm convinced you have a lot of motivation, I mean, if you're planning hundreds of thousands of new LTE base stations in the coming year. What do you see as the benefit of Intel China Mobile Research and Intel working together on this?

Chih-Lin I: Well, I think, of course, we all know that Intel Labs is a premier R&D labs in the world, but more so in the sense that in the challenge we are facing today, we need to make our radio access network much more power-efficient and much more cost-effective. The traditional standalone base stations really cannot meet that need. And with the large number of cell sites, you know, our power consumption last year was like 30 billion kilowatt hours of power was consumed by us, and 70 percent of that was in the base stations. So we need to make this part much more power-efficient and cost-effective.

Plus, our deployment is getting denser and denser, so the interference across cell sites is becoming a big problem. And we need a clever way to turn that interference into use for signal. And also, the fact that we are very happy with the exponential growth of the traffic. We're not very happy with the fact that our revenue is not growing proportionally, so we need to be much more cost-effective. And all of that leads us to the C-RAN. Okay?

And the partnership with Intel Labs actually enables us to work on this concept from day one, when we proposed it in 2009. We actually started a close partnership with Intel Labs as well as a few other industry leaders on this concept.

Justin Rattner: So tell us about the goals here. What do you expect to achieve?

Chih-Lin I: Okay. We see that, in C-RAN, we are combining centralized processing, collaborative radio, and also cloud computing all together in this one new, I would say, revolutional architecture, but it has a smooth evolutional path for our migration.

And that can [reach clean] and the cost-effective both deployment and operations. And by partnering with Intel Labs, who has the expertise in the open common platform, the server platform, we see the opportunity of replacing the proprietary hardware with actually what – we're very happy with what we have shown already the labs did last year in the IDF – I think they all have seen, right?

Justin Rattner: Yep, Sunny was here last year.

Chih-Lin I: Yeah. Last year, we already saw the capability of replacing the proprietary system with an open platform and the software stack of the LTE.

Justin Rattner: Okay. Well, Sunny's not been idle. A year's worth of progress. Why don't we walk over there and see what he's up to? He's always got something good.

Chih-Lin I: Hi, Sunny. Nice to see you again.

Sunny: Hi, Justin. Hi, Dr. Chih.

Justin Rattner: Hi, Sunny.

Sunny: Hi, Justin.

Justin Rattner: Welcome back. Not many people are in an IDF keynote back-to-back years, so you're very special.

Sunny: Yeah. I'm very glad.

Justin Rattner: All right. You've got a server array sitting here. Why don't you tell us what you've done?

Sunny: Yeah. As you mentioned, Justin, last year we demoed a single LTE base station running on servers. So we made a lot of progress since then. So today, I can show you [unintelligible] base station running on servers. And the load can be shared among different servers, as in C-RAN base station [two].

Justin Rattner: So we have a little Cloud-RAN running right here?

Sunny: Yes.

Justin Rattner: Wow.

Sunny: You are right. So here you can see we have two Intel core i7 servers, and they are running base station workload. And we also have emulated mobile devices connected to the base station, and they're streaming video.

Justin Rattner: Okay. So again, you make it go.

Sunny: Yeah. So here, I start with a configuration, and we have only one server is running and processing two base station workloads. Each base station has a few mobile devices. So, as I add more mobile devices into this network, so [unintelligible] the load of this server starts to increase, and this is the type of scenario we see every morning, such as one more user come online in an office area. And then, so one of the base stations can be offloaded to another server, in order to do load balancing and also to support more users in this [recruit].

Justin Rattner: I mean, these days, right, the base stations are typically over-provisioned. They're designed for the expected peak load, right, and so you wind up with a lot of hardware in the field that you can't really bring into play when hotspots develop in the network. You can continually balance the load.

Sunny: Yes.

Justin Rattner: That's fantastic.

Chih-Lin I: This is wonderful, but Sunny, I have a question, because actually there is another challenging situation we'd like to deal with, which is when like, at midnight or off peak hours, when the traffic load actually is not that high, can you somehow help us to conserve the energy in the system when we don't really need that many processors going?

Sunny: Yeah. So, in addition to the load sharing when the traffic load is high, our C-RAN prototype can also consolidate multiple base station workloads into a smaller number of servers, such as one, turn off to mobile devices. And then, you can see one of the base stations can be moved from this serer to another server, and then, so we can turn off this server in order to save energies.

Justin Rattner: Right, and I imagine, as you say, in lightly-loaded situations, that could save a tremendous amount of power, right, because in the current model, those guys are up just all the time waiting for traffic.

Sunny: Yeah.

Justin Rattner: All right, well, that's really great. So, you know, if you're going to come back for a third year . . .

Sunny: I hope so.

Justin Rattner: You've got to do something really cool. No guarantees, I'm not making any promises, we'll delete this part of the video.

Chih-Lin I: No, wait, let's try to make a promise. I think the next, so we really have a few more even higher challenges to deal with, can we put it out here?

Justin Rattner: All right, let's put it out here. What's next?

Chih-Lin I: The next, virtualization, maybe?

Sunny: Yeah, so right now, so we are continuing working on the C-RAN architecture optimization. So, we want to support a different loading levels for different scenarios. So, our research target is to support up to 100 base stations in just one server when the load is light.

Justin Rattner: 100 base stations on one server, wow.

Sunny: So, that's where our research goes.

Justin Rattner: Okay, you do 100 base stations, I'll bring you back next year.

Chih-Lin I: Okay, that's a promise.

Justin Rattner: Fantastic. All right, let's shake on it.

Chih-Lin I: Very good, thank you.

Justin Rattner: It was really great having you here. Thanks so much for coming.

Chih-Lin I: Yes, thanks very much.

Justin Rattner: Thanks so much for coming, and we look forward to a great continuing partner.

Chih-Lin I: Well, I look forward to a continuous and close partnership with Intel Labs and all the industry leaders that wants to join this revolution or evolution technology.

Justin Rattner: All right, thanks again, thanks for coming. Chih-Lin I, China Mobile Research.

All right, time to wrap it up. We really pride ourselves at Intel Labs on working on research that matters, and I hope this morning you've seen us focus on issues that go all the way literally from the transistor level with the Moore's Law radio up through this whole range of issues and sort of the middle of the stack, you know, dealing with the trade-offs between convenience or security and long battery life and so forth. And then, to conclude, as you just saw, with this notion of using software to find radio technology technologies to completely transform the structure of the radio access network. Looks like it's going to happen in China if Sunny keeps producing, his team keeps producing, and China Mobile Research is just, I think, a great partner for doing this, because nowhere else on earth do you find the kind of demands on radio access network technology. Lots of exciting work, and of course more work to do.

So, this is, as I said at the outset, this is my seventh IDF keynote, research and technology keynote, and it really amazes me when I reflect back over the last seven years, how far we've come at Intel, and how far we've come as an industry. If we were giving this keynote seven years ago, people would just be walking away scratching their heads and going, that's just a bunch of pipe dreams. It used to be, we'd come out, all we'd talk about is the silicone. I think that first keynote, I talked about the core micro architecture, introduced core micro architecture, was all about silicone. As we've moved along, we talked more and more about platforms and all the different kind of platforms. And I think today, yesterday, and the day before, you've really come to appreciate the focus we have on creating great experiences across the entire computer continuum.

I won't even try to speculate where we might be seven years from now. I won't even speculate if I'm here seven years from now, but I'm sure that we'll enjoy the ride into the future together. Thanks very much, see you next year.

Male Voice: Attention, ladies and gentlemen.

[End of recorded material]