Justin Rattner - 9-15-10

[Beginning of recorded material.]

Female Voice: Ladies and gentlemen, may we have your attention, please? 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 and our most recent form 10Q or 10K filing available on our website for more information on the risk factors that could cause actual results to differ. Ladies and gentlemen, please welcome Renée James.

Renée James: Good morning and welcome to day three of IDF. At Intel, we have this thing, if you don't do a good enough job, you get a redo. I'm just kidding. I'm here to introduce day three and just recap what we've done in the last couple of days together.

We have talked to you a lot about the continuum of computing, and I just want to just recap a little bit about what we mean and hopefully you've all gone to the sessions and you get it. But basically, we have a point of view that Intel can deliver the best products across the entire range of computing from Xeon servers to our newest Core processors and, of course, you heard from Doug yesterday, the ever power-efficient Atom family of processors. The combination of hardware and software, we think, allows each of you the opportunity to create and deliver something really spectacular.

So hopefully, as you leave today, after of course you listen to Justin and go to all the sessions, you'll see that Intel's working very, very hard to enable you to build fantastic products for whatever part of the market you serve.

Yesterday, Doug showed you the combination of software in the Atom family and all those really wonderful products that we could create, and there were some surprises for all of us, with that wonderful Dell convertible tablet. I talked to you about software and what we're doing for developers, and hopefully you went across the street and saw AppUp. It's still there today. And, of course, we had a tremendous number of technical sessions and our insight session from Geoff Lowney on parallel computing. Hopefully, we were able to give you insight into where we're going and the really difficult problems we're trying to solve.

Today, however, is our chief technology officer Justin Rattner, and for those of you that have been to IDF many times, you know that Justin's keynote is one of the highlights of the entire IDF experience. He'll be sharing with you the future of computing and where Intel's going with contextual computing and some other really, really cool surprises. So we're looking forward to hearing that in a second.

The technology sessions and showcases are open. It's the last day for the tech showcase. If you haven't been through it, take the time. There's a lot of really interesting things down there. And it wouldn't

be Intel if we didn't tell you about the next IDF, because you know we like to plan ahead. So we're already thinking about it, and our first one next year will be in China. It'll be April 12th through 13th, and then we'll be back here in San Francisco in September, the 13th through the 15th. So those of you who like to plan ahead, you can look forward to that, and we look forward to having you back again.

Finally, because, as you know, if you were here yesterday, I love to give away prizes, I'm going to get to give away another one. And I'd like to congratulate – now this is hard name for me – Chaim Gartenberg -- for capturing the datacenter IDF experience, which we're going to see in a second, and then I'll tell you what he won.

So for that wonderful datacenter dude video – it was funny, thanks -- he won an ultimate home system, a 55-inch LCD TV, a Boxee Box, which Doug talked about yesterday, and an Atom-based home server. We want to thank Cisco and D-Link for sponsoring that. And with that, I want to thank you for coming back for day three and on to Justin.

Female Voice: Ladies and gentlemen, please welcome Justin Rattner.

Justin Rattner: Good morning. Thanks for sticking around for day three. Always appreciate that. I think from the video it's pretty clear that even the best relationships have room for improvement. And that's certainly true when we're talking about the relationships we have with our various digital devices. Whether we're talking smart phones, tablets,

netbooks, notebooks, TVs, cars, you name it, these relationships are ready and ripe for improvement.

The question is how can we change the relationship so that we think of these devices not as devices but as assistants? Or even companions? You know, things that are just indispensable in our daily lives. As the opening video suggested, we believe that context-aware computing is poised to fundamentally change the way we interact and relate to the devices that we use today.

Through a combination of sensing and inference – and we'll talk more about what that means – future devices will constantly learn about you. They'll learn your habits, the way you go through your day. They'll really learn about your life. They'll understand your friends and so forth. They'll probably even know how you're feeling. They'll know where you are, and more importantly, they'll know where you're going. They'll anticipate your needs, as we saw in the video.

They'll know your likes and dislikes. As you shop and browse and surf and whatnot, they'll learn where you like to go on the net and what topics interest you. Now that may sound a little crazy, but let me show you in a very real way what context-aware computing is. Please welcome Tim Jarrell, VP and publisher of Fodor's Travel. Come on out, Tim.

Tim Jarrell: Hello, Justin.

Justin Rattner: Thanks for coming.

Tim Jarrell: Thank you.

Justin Rattner: It's great to have you and great to meet you. We just had a few minutes to chat. Tell us about Fodor's interest in context.

Tim Jarrell: Well, Fodor's is a 70-year-old travel content brand. Our mission is to provide up-to-date, accurate, timely, but relevant information to consumers to make trip planning better. Consumers trust us because they want us to provide content that is relevant to them, that matches their style and their values.

Now we're trying to figure out, as a lot of people are, how to take our content and put it into digital phones. So far, we're not so impressed with the technology, and we're not impressed with the applications that are out there, and consumers don't seem to be willing to pay a lot of money for the destination apps that are currently available.

Justin Rattner: Okay. So did Intel and Fodor's come to be working on this problem?

Tim Jarrell: Well, about a year ago, Intel came to us with this concept of context awareness. And we were blown away, because for us, it was exactly what we wanted to do. That we could actually create an application

where we could really be at the consumer's elbow, guiding them, shepherding them through their trip with recommendations – timely recommendations – based on their past preferences, their behavior, our personas, but more importantly the circumstances and the context which they were in.

For us as a consumer brand, that's exactly where we wanted to be. So we worked with Dave Sandage, in Intel, and his team to develop a working prototype.

Justin Rattner: Can we see it?

Tim Jarrell: Absolutely.

Justin Rattner: All right. Great.

Tim Jarrell: Dave?

Dave Sandage: Tim.

Tim Jarrell: For the prototype that we built, we used content for New York. But for today's demonstration, we actually have loaded the device with San Francisco content. Now if I'm a visitor to San Francisco and let's say I don't know a whole lot about San Francisco, but one thing that I've heard about is the Ferry Building. And I understand the Ferry Building might be by the waterfront, so I'd like to go there. So, Dave, how would that work?

Dave Sandage: All right. Well, on this MID, we have an application that's running on top of our context framework, which enables it to be context aware. You can see we have a map that shows your current location, indicated by the red dot. What we'd like to do now is make our device believe that we're actually at the Ferry Building. And you can see that the dot has moved over to show where we are.

Now, while I'm here, I'd like to find a restaurant to go to dinner at tonight. So I can go ahead and press search. I can look for restaurants – in the future, it's dinner tonight -- within a half mile of where I am currently.

Now, at this point, other travel apps would have me actually enter the cuisine types that I'm interested in for tonight. However, since this is context aware, it already knows what kind of cuisine I like, what style of restaurants I like, even how much I am usually willing to spend on a night out when I'm on vacation.

Justin Rattner: And Dave can be a real tightwad.

Dave Sandage: Yeah. I know who signs my expenses. But in this case, let's go ahead and let the device use what it has learned about me to filter my restaurant choices. And it's now looking through the Fodor's database. It's come up with four restaurants that are near the Ferry Building. On these, I could go ahead and plot them on the map so I could see where they are, or I could, in fact – if I'd made a

reservation – add them to my calendar, which the system would then calculate the travel time based on where I am, update it throughout the day, based on traffic, and make sure to remind me so I'm not late for the restaurant.

Tim Jarrell: Now what's important about this selection is it just isn't randomly all the restaurants that are in that area, nor is it just the restaurants that we at Fodor's may recommend. But these are choices tailored for you.

Dave Sandage: That's exactly right.

Tim Jarrell: And that's what makes it, for our standpoint, much more workable. Now, when we developed this prototype, we talked about how travelers may be in different modes as they go through and explore a city. They may be in mission mode. I need to get to a particular restaurant at 7:00, and I don't want any other interference. But they also may be in wander mode. So if I'm in –

Justin Rattner: That's the way I travel. I just kind of wander around.

Tim Jarrell: Absolutely. So if I'm in wander mode, I'm down by the Ferry Building, and I have a couple of hours to kill, I want to explore the neighborhood. Again, based on my preferences, what will the device do?

Dave Sandage: Well, when you're in wander mode, you don't actually have specific things you want to see. You just really want to wander around the neighborhood and maybe have the device make recommendations of places to go, things to see, based on my current circumstances, what I like, and what's around me. So the context engine is, in real time, generating suggestions for what's around you.

And you can see the suggestion button has lit up yellow, which means the system has a new set of suggestions for me. So let's press that and take a look. We're by the Ferry Building, so that's actually the closest thing to me. But I notice down here also the San Francisco Railway Museum is here. Now I can go ahead and plot that on the map, and we can see that, in fact, it is very close. It's within walking distance. And before I got down there with my device and used this, I didn't even know that was there. So I wouldn't have found it without wander mode explicitly giving me those recommendations.

Justin Rattner: Perfect.

Tim Jarrell: Now, Dave, the other thing this device does is that it's collecting information about me all the time. So it knows where I am and how I'm going through the city. Is the device collecting that information in a usable way?

Dave Sandage: Yes. All of the context information that's being created and gathered by the device is being stored and uploaded to a web

service. Now what we can do with that is actually allow the user to automatically generate a blog of where they went, what they saw, what they did, that they can then – optionally – upload to a –

Justin Rattner: Oh, this is perfect for me, Tim. My wife is always bugging me to, you know, get the pictures organized and tell the story.

Dave Sandage: And the beauty of this is it's done completely automatically for you. All the pictures that you took with the blog, all the places you went, have been recorded and generated. Now one of our engineers wandered around San Francisco on Sunday with his device and went to the Palace of Fine Arts, took a picture. This is actually the unedited output of our blog generator.

Went to Fishermen's Wharf. You can tell -- it said when he arrived, when he left. And down at the very end, it even generated a map that showed all the points of interest that he visited, as well as links to the Fodor's data, so that when he shares it with friends and family, they can click on and see what he saw.

Justin Rattner: That's fantastic. Is there also a mode where you don't have to do all that walking around, and it will generate the blog for you?

Dave Sandage: The fake vacation mode. We'll work on that.

Justin Rattner: Perfect.

Tim Jarrell: You put it in your wife's purse and let her wander around.

Justin Rattner: Exactly. All right. Well, thanks, Dave.

Dave Sandage: You're welcome.

Justin Rattner: So, Tim, let's talk about where we go from here. You talked about working on the prototype and basically doing it in New York. Did you actually run field tests with this?

Tim Jarrell: Yes, we did. We conducted a live test with 30 participants last spring in New York City, and we gave these poor research subjects a working prototype. And the working prototype was really clunky. It was big. They had a battery pack. There was a funky little antenna.

Justin Rattner: Not this, like, little MID.

Tim Jarrell: No, not the slick device we have today. And we gave consumers that device, and we let them wander the city, and they loved it.

Justin Rattner: Really?

Tim Jarrell: They absolutely loved the device.

Justin Rattner: Well, I guess the acid test is would they be willing to pay for it?

Tim Jarrell: Well, that's right. And what the participants did is they came back and, on average, they said they'd be willing to pay $20 or more for this particular application. And that compares to the destination apps that are out there at about three or four dollars. But even more importantly, about 10 percent of the subjects gave the device a name.

Justin Rattner: A name?

Tim Jarrell: A name. In one case –

Justin Rattner: Like a pet rock here.

Tim Jarrell: Well, it wasn't -- well, it could be their pet rock. But in one case, it was sort of a beloved relative who was always giving suggestions. And it was this way of honoring that person, as if she were on the trip, herself, and giving suggestions on what they would do. And, in fact, in the middle of the day, some of the subjects would say, I wonder what the name of the device says we should do, and they would go to the device. And it became sort of a ritual in of itself.

Justin Rattner: I mean, it sounds like this was having the exact effect I talked about a few minutes ago, where people began to think of the Vacation Assistant as an assistant.

Tim Jarrell: Absolutely, as an assistant, as a trusted companion. And for us at Fodor's, that's exactly where we want to be. The ability to be with

that consumer, with that traveler, and shepherding them through their itinerary and giving them relevant contextual recommendations as they go through their trip, that, as a consumer brand, is important for us.

Justin Rattner: Excellent. Where does Fodor's go next with this?

Tim Jarrell: Well, we're excited about it. We are continuing work on the Personal Vacation Assistant, based on our research results, and we are optimistic that we will have a product.

Justin Rattner: Outstanding. That's just great, Tim. Thanks for being here.

Tim Jarrell: Thank you. Thank you very much.

Justin Rattner: Tim Jarrell, everyone. And as you can see, we really had a lot of fun working with Fodor's on this project. Well, I wanted to get that out early so you weren't sitting here for the majority of the morning saying, you know, well, what is he really talking about, you know, context and all of that. So I think the Personal Vacation Assistant really epitomizes the kind of technology that we're talking about.

But I think it's important to ask where did all this come from? What are the origins of context awareness? Did this thought just come out of the sky, or have we been thinking about it for a while? And it turns out that there's been more than 20 years of research in context-

aware computing, and it's really been simmering in the research community all that time.

Most famously, and if you've ever touched this technology, I'm sure you've heard of Mark Weiser of Xerox PARC, who articulated the original vision for context-aware computing back in 1991. Yes, almost two decades ago. And I thought it would be valuable this morning to hear from one of Mark's colleagues, to sort of talk about how this idea evolved at PARC. So please welcome Bo Begole, Xerox PARC Ubiquitous Computing research director. Good morning, Bo.

Bo Begole: Good morning, Justin.

Justin Rattner: Come on down.

Bo Begole: Good morning.

Justin Rattner: Well, help me understand something. I mean, you know, PARC is famous for inventing personal computing. We wouldn't have all these people here this morning if you guys hadn't hatched this idea decades ago. But literally, I mean before people were going out and buying PCs so they could connect to the Internet, you were talking about making PCs disappear. What on earth were you thinking?

Bo Begole: Right. Well, the irony is that even though it was invented at PARC in the early '90s, people at PARC had been living with personal computers for a while. And they were starting to think –

Justin Rattner: Altos and all of that stuff.

Bo Begole: Right. And they could see the future, with Moore's Law, things were miniaturizing and becoming faster, and it wouldn't be long before there'd be more than one computer, there'd be multiple computers. And another aspect was that a lot of the social scientists at PARC were pointing out that, you know, this personal computer still required you to do everything manually. You had to drive it in every possible way, and it didn't seem to be very intelligent.

Justin Rattner: Right.

Bo Begole: So one of the big ideas that came out of the research there in Ubiquitous Computing was this notion that when our environment had a proliferation of devices that had some ability to sense the physical context, then they could do things more appropriate to the situation by detecting what needed to be done.

Justin Rattner: So, I mean, the essence of context awareness.

Bo Begole: Exactly right. And one of the things they created were services that could do things like automatically connect your handheld computer to displays and projectors in the room, for example.

Justin Rattner: Well, I'm still waiting for that. I carry around my VGA cable. In fact, now I have to start carrying around an HDMI cable. Maybe the Intel WiDi technology will put us out of our misery. Let's hope that's the case.

Bo Begole: Yeah, we're all looking forward to that. And then they also created services that would find your colleagues, where they're located, and whether it was a good time to try to reach them to discuss some issue.

Justin Rattner: Sure.

Bo Begole: Finding equipment in your building and in hospitals and other things like that. I want to point out, though, that even though a lot of these services were based on location detection, it's not the case that context is solely about location.

Justin Rattner: Yeah, a lot of people do get those confused.

Bo Begole: Absolutely. It also includes things like who you're with, what you're trying to do, and how you're going about it.

Justin Rattner: Well, we're going to talk a lot more about it. But we're talking about thoughts that the folks at PARC were having almost 20 years ago, and, you know, here we are today, and we haven't seen a lot of commercial success for context-aware computing. What's going on?

Bo Begole: Yeah. Well, the sad reality is that researchers are always overly optimistic about how soon their ideas will achieve commercial success.

Justin Rattner: Tell me about it. I have to deal with about 1,000 of them every day.

Bo Begole: Yeah, yeah. But what we've seen recently, though, is with the advent of high-powered handheld computers and ubiquitous wideband networking, wireless networking, and GPS-enabled smartphones, that we now have the infrastructure that's needed, the critical mass needed, to create these services.

Justin Rattner: Yeah, so that's really the big change, that all of the glue, if you will, is now in place, and so the opportunity is out there. Hey, thanks, Bo.

Bo Begole: It's been a pleasure.

Justin Rattner: Thank you so much. Bo Begole, Palo Alto Research Center. Okay. So, you know, you're developers. At least some of you are developers. And you're probably sitting there asking yourself, am I ever going to get to the meat of the matter and actually talk about how you get the sort of things working that we saw in the Vacation Assistant and the kinds of things that Bo was just talking about.

And that's a perfectly legitimate question. We don't want to disappoint you this morning, so we're going to take off the covers and really tell you what's involved in context-aware computing. And to do that, I thought it would be great if we could get one of the goddesses of context-aware computing at Intel to join us onstage, and that's Lama Nachman. Come on out, Lama. Good morning.

Lama Nachman: Hi, Justin.

Justin Rattner: Okay, well you're going to be our guide and our navigator and kind of take us through all the ins and outs of context-aware computing. And I really can't think of anyone better prepared to do this. It's like you've been working towards this your whole life. And so now you've got a rapt audience, and they're ready to learn the ins and outs of context awareness. So are you ready to do that?

Lama Nachman: Absolutely.

Justin Rattner: Okay.

Lama Nachman: So we've talked about context. We've talked that context is really everything from what you're doing, who you're with –

Justin Rattner: Oh, I like that. Context if everything. Keep going. Keep going.

Lama Nachman: Absolutely. It's like, it really needs to understand where you are and what the environment around you is, the weather, your calendar, all

of that information. But really, if you think about context, let's start with sensing, since it's really at the core of what context awareness is all about.

Justin Rattner: Right.

Lama Nachman: So –

Justin Rattner: You want to drive?

Lama Nachman: Absolutely.

Justin Rattner: Okay.

Lama Nachman: Okay, so devices today, basically, we're seeing a lot of sensing getting integrated into these devices, and it's really simple. Sensing. You can understand that the phone is next to your ear or you're holding it this way or where you are, and so on, but even though these are really simple things, it generated quite a bit of excitement in the market, right? And then once these APIs were opened up to developers, many new applications emerged, and these actually become even more exciting.

So, you know, this clearly is just basically skimming the surface. So let's move beyond the notion of device context and start seeing what . . .

Justin Rattner: Okay, getting beyond the basics, great.

Lama Nachman: Exactly. So the first example that we have here is really about using the accelerometer to infer what the personal context is.

Justin Rattner: Okay, so not just what the device is doing --

Lama Nachman: Not just how you're holding it and what the device is doing.

Justin Rattner: Right. But something about the person who's actually interacting with the device.

Lama Nachman: Exactly. Exactly. So the first example we're going to start with is essentially really trying to talk to the problem of falls with adults over 65.

Justin Rattner: Right, which is a big problem.

Lama Nachman: It's a big problem. More than one-third of adults over the age of 65 actually fall and end up, in many cases, in hospitals because of that.

Justin Rattner: Yeah, so can we see it?

Lama Nachman: Absolutely.

Justin Rattner: You going to show it to us?

Lama Nachman: Let's go over.

Justin Rattner: All right.

Lama Nachman: So, Jigna is going to walk us through this example.

Justin Rattner: Hi, Jigna.

Jigna Kapadia: So, like Lama said, falls is the number one reason why elderly over 65 go to hospital emergencies and have trauma. What I have here today is our Shimmer sensors. These are wearable, and I believe you're wearing a couple today.

Justin Rattner: Yes, okay, we've got a camera on this. Here's my Shimmer, and there's one there, and there's one over here. The latest in wearable computing fashion, for you fashionistas out there.

Jigna Kapadia: So as you can see, you're going about uninterrupted through the keynote wearing those sensors. What these sensors do is essentially measure your stride time and your swing time. The stride time and swing time are the two most important variables in human locomotion, and these are the variables that are used to measure your walking speed and your running speed.

Justin Rattner: So this is me walking around on stage?

Jigna Kapadia: Right. So I took this data while you were walking around on stage before.

Justin Rattner: Okay.

Lama Nachman: I'm not sure what we infer from that right now.

Justin Rattner: I don't know. But at least I didn't fall down yet. Okay. So that's good.

Jigna Kapadia: Right, so medical researchers have found that irregularities in your gait or poor movements are excellent to identify any risk of falling, and if identified, we can predict falls.

Lama Nachman: And the cool thing is that as we start looking at this over time, then we can start seeing these things emerging and your increased risk of that and then it provides an early warning.

Justin Rattner: So if you're able to collect a lot of data over a period of time, you really get a lot more insight?

Lama Nachman: Exactly.

Justin Rattner: Yeah, you can't do that in the doctor's office. A few feet across the exam room is not going to do it.

Lama Nachman: Absolutely.

Justin Rattner: Okay. Super.

Lama Nachman: Thanks, Jigna.

Justin Rattner: Thanks, Jigna.

Lama Nachman: Okay, so basically, over the last multiple days at IDF, we've seen a lot of examples of personalized TV experience, understanding what the user cares to watch, and trying to tailor the content accordingly.

Justin Rattner: There were a lot of people at the TV exhibit.

Lama Nachman: Exactly. So one of the next examples of context awareness that I want to show you has to do with really trying to make that happen. So to make that happen, you really need to understand, you know, who's actually watching the TV, so that we can start understanding what specific shows they're actually going through and then, based on that, create this --

Justin Rattner: You're not going to put some big camera on top of the TV?

Lama Nachman: Ah-ha, so let's see that. So actually, we're going to walk over here to Geoff, who's in the living room.

Justin Rattner: Should I sit down?

Lama Nachman: Yes, please, go ahead.

Justin Rattner: Hi, Geoff.

Lama Nachman: And Geoff will show us his remote control.

Geoff: So what we have here is a remote control that enhances the smart TV experience. It allows us to distinguish who is wielding the remote control based on the way they hold it and lets you provide a personalized experience for the television.

Justin Rattner: Okay.

Geoff: So let's give it a try here. So what we have here is a small sensor pack that we've augmented on the remote control that --

Justin Rattner: Kind of like the Shimmers.

Geoff: Kind of like the Shimmers that you have on your sock there, that lets us understand how you push the buttons and how you wield the remote control. And then as soon as you pick up the remote and start surfing the television, you can see that, you know, this is the data that we trained on before, so you can see that it's recognized me as person one and that I am the one using the remote, and it's provided me some little personalized recommendations there.

Justin Rattner: Where did you get all this video?

Geoff: This is from our friends at CBS. They've provided us some great new content here of their upcoming season here.

Justin Rattner: Okay. All right. We thank them for that.

Geoff: So you can give it a try here.

Justin Rattner: Okay.

Lama Nachman: Let's see. All right.

Justin Rattner: Oh, okay, so it's recognized me. Now, how's it doing that?

Geoff: Well, it recognized based on the motions of your hand and the way that you do the remote control. It actually uses something called unsupervised learning, so it's learning all the time in the background.

Justin Rattner: So it's learning right now?

Geoff: Yep. You bet.

Justin Rattner: Okay, I don't feel anything.

Geoff: Well, it's working in the background. It's watching just the same way as when you were walking before with the Shimmer. It's looking at your behavior and understanding it.

Lama Nachman: Thank you so much, Geoff.

Geoff: Yeah, no problem.

Justin Rattner: Fantastic. When do I get one of these? You don't have to answer that question.

Geoff: Thanks.

Justin Rattner: Thanks.

Lama Nachman: Okay, so we've talked about multiple examples, and really what I want to walk you through now is what's really under the hood, how we actually do this, right? So we talked about this notion of starting with simple data, right, that comes straight from the sensors. And then, you know, we do some inference on this data, and what we're trying to understand is what classes we want to get out of that. So let's give you an example.

Justin Rattner: So this is, you know, when people talk about sensing and sense-making, we're going to put that together?

Lama Nachman: Exactly.

Justin Rattner: Okay.

Lama Nachman: So let's walk through the example of an accelerometer. So you might have an accelerometer data, you get that through the system.

Justin Rattner: I happen to have two.

Lama Nachman: Exactly. You extract some features out of the accelerometer data. So let's say, for example, the minimum or the maximum of the signal or something like that. Then we feed that into an inference algorithm. And this inference algorithm, like in the case of the TV remote, was like a support vector machine. That inference algorithm then would essentially generate the class in the specific case, which is who's the person who is actually holding the remote. The cool thing about this is that kind of inference pipeline extends across multiple types of sensors and multiple types of classes.

Justin Rattner: So that's a model, that's a paradigm for doing this sensing and sense-making.

Lama Nachman: Exactly. So you can think of many other examples, like, for example, inferring emotional state from your GSR and heart rate, and that's actually pretty cool.

Justin Rattner: Okay, getting into interesting territory.

Lama Nachman: It's actually kind of along the lines of a lie detector, but you can think of --

Justin Rattner: Oh, great. Now you've made me feel a heck of a lot better.

Lama Nachman: And then, you know, the same thing kind of applies to audio, right? You can think of what's playing in the background, is it music or is there somebody talking in the background, and so on, or even camera, which ultimately is a great source of rich context, where you can start seeing who's, for example, in front of the camera and help you understand who's coming towards you and stuff like that.

Justin Rattner: Okay, so we've got this model.

Lama Nachman: Yeah, and so basically, really, the next step, you know, hard sensing, which is everything that we've talked about so far, is real sensor data that actually we infer some context out of it. But really that's not really everything. There's a lot of other soft sensing information that's very relevant.

Justin Rattner: Right. In fact, we saw some of that in the Vacation Assistant.

Lama Nachman: Yeah, exactly. So, for example, you can get information from your calendar. I get a note from your calendar that you're at IDF, doing something. I can understand what meetings are coming up, I can understand what's going on in the world.

Justin Rattner: Where I've been, where I'm going.

Lama Nachman: Exactly. So there's a lot of stuff that's coming through, applications that you're running, what you're browsing, et cetera, that we use. So what we actually end up doing here is essentially fusing information from that hard sensor, plus all of these soft sensing, and then what happens is we get a richer context and improve the accuracy, because what you really don't want to do is be very dependent on a single source of information.

Justin Rattner: And those soft sensors, they can include things like social networks and preferences and all these things.

Lama Nachman: Exactly. Exactly. Exactly.

Justin Rattner: Okay. Actually, there are probably more soft sensors then there are hard sensors.

Lama Nachman: Absolutely.

Justin Rattner: Okay, good.

Lama Nachman: So let me take you over there to a demo, right. And basically what we're showing over here is, in this demo, we've collected data from different sensors, both soft and hard, and we basically collected those in different scenarios that we want to actually go and detect.

So something like you're presenting in a meeting, right, or you're shopping or sleeping or whatever.

Justin Rattner: I'm just doing an IDF keynote.

Lama Nachman: Yeah, exactly.

Justin Rattner: Not a very busy morning.

Lama Nachman: Not really. So we collect all of this data and then basically what we do is we run this data through our system to essentially infer what the person is doing. And what you can see over here is . . .

Justin Rattner: So these are all the . . .

Lama Nachman: These are all the sensors.

Justin Rattner: Right.

Lama Nachman: So from the accelerometer, what we have is, you know, we're trying to infer the physical activity, so if somebody's walking, sitting, et cetera. From the microphone, we want to infer what's playing in the background. And actually, in this specific case, we have part of the real audio clip. We can listen to that.

Justin Rattner: Oh, it sounds like a –

Lama Nachman: A typical meeting.

Justin Rattner: It's a typical Intel meeting.

Lama Nachman: Exactly.

Justin Rattner: Yes, risk/reward ratios and all that. Yeah.

Lama Nachman: Exactly. And then from the soft sensing you were getting –

Justin Rattner: Oh, this is the soft guys.

Lama Nachman: These are all the soft guys. This calendar. You're probably presenting in a meeting through PowerPoint.

Justin Rattner: Right.

Lama Nachman: There is all of this other information that is coming in.

Justin Rattner: Actually, I just looked at my PowerPoint. I think you pretty much have the day.

Lama Nachman: Exactly. And then from location, we can tell where you are.

Justin Rattner: Okay.

Lama Nachman: So now the system will go recognize these, and you can see that now the system just recognized that you're sedentary, so you happen to be sitting at that meeting.

Justin Rattner: That's me, yeah.

Lama Nachman: And then there is human speech in the background. It makes sense, right?

Justin Rattner: Right.

Lama Nachman: There is PowerPoint that you're actually presenting.

Justin Rattner: PowerPoint.

Lama Nachman: And you're actually in some meeting room at Intel, right? So then all of this information, which is that first level of inference, gets fused into this high-level algorithm, which is a base n graph, in this case, and it detects that you are presenting in a meeting. So that's what you can see and hear.

Justin Rattner: So this is like looking – you know, how confident, right? That sort of confidence –

Lama Nachman: Yeah. That's the confidence level that you're actually presenting.

Justin Rattner: Okay. Well, I'm pretty confident that you've got it right here. I don't even see a bar coming up in any of the other categories.

Lama Nachman: Exactly.

Justin Rattner: Okay. So that's how we bring it all together, but are there ways that don't require so much direct interaction?

Lama Nachman: Yeah. Basically, what's happening here -- so let me show you an example. Presenting in a meeting, for example, you could use that to infer things like that your wife, for example, is in a meeting and you don't want to call her or something like that.

Justin Rattner: Oh, yeah. So that happens all the time.

Lama Nachman: Yeah. So, here, we have actually this -- take a look at this.

Justin Rattner: Okay.

Lama Nachman: This is actually a sense system that we've built and, basically, in here, you can imagine using this information that we talked about, the activity that you are doing, to actually animate an avatar, right?

Justin Rattner: Okay. Oh, yeah. All these guys are moving.

Lama Nachman: Yeah, so all these guys are moving. These are avatars of people going about their normal lives. And if they choose to share that

information with you, so, for example, your wife, you could see that, for example, she was in a meeting.

Justin Rattner: So I know she's in a meeting.

Lama Nachman: Right. So let's see, for example, what Wendy is up to. All right?

Justin Rattner: Okay.

Lama Nachman: So here – oh, I'm sorry. That's actually here.

Justin Rattner: That's Kieran.

Lama Nachman: Yeah.

Justin Rattner: Okay.

Lama Nachman: That's kind of an interesting thing here. So Kieran is actually sitting at a café, you know, drinking coffee, and he gets a phone call.

Justin Rattner: Right. Oh, okay.

Lama Nachman: So now we can see him actually on the phone, and then he stands up, walks.

Justin Rattner: Walks there.

Lama Nachman: He's going out.

Justin Rattner: Late for the keynote.

Lama Nachman: Hopefully, he actually paid for his coffee. We don't know that part of the context.

Justin Rattner: Oh, I trust Kieran.

Lama Nachman: I'm assuming that.

Justin Rattner: He's very responsible.

Lama Nachman: Exactly. So that's one example.

Justin Rattner: Okay.

Lama Nachman: So basically we talked about this notion of seeing this in the moment kind of representation of –

Justin Rattner: So, yeah, at a quick glance I can tell exactly what's going on –

Lama Nachman: Exactly.

Justin Rattner: With all my friends or family or whatever it is.

Lama Nachman: Exactly. So another kind of compelling concept of this is if you think about aggregating this context over an extended period of time. Think about, for example, knowing your activities over an extended period of time. That's actually quite compelling.

Justin Rattner: Right. We saw that in the gait demo.

Lama Nachman: Right.

Justin Rattner: Much more interesting.

Lama Nachman: Or of, for example, your higher-level activity, right?

Justin Rattner: Sure. So this is all building up the context, I guess, if I can put it that way, right? So you're context is growing and growing.

Lama Nachman: Exactly. And you might ask why, for example, it's important to know what activity you're doing over an extended period of time, right?

Justin Rattner: Uh-huh.

Lama Nachman: But I would argue, for example, the same thing applies to finances, right?

Justin Rattner: Sure, right.

Lama Nachman: So even though you might know every specific transaction that you make, understanding how that impacts your whole financial picture really leads you to look at data over an extended period of time and then these interesting things start to –

Justin Rattner: Okay. So, besides time, the other dimension is?

Lama Nachman: So there are the other devices. You don't just interact with one device. You're using your laptop. You're using your phone. You go to the TV. So being able to aggregate context across all these different devices is quite –

Justin Rattner: So context is even a bigger idea than any particular device.

Lama Nachman: Absolutely.

Justin Rattner: And I want that context to follow me wherever I go. So I get the feeling we've seen all the pieces now. Can you put it together for us? How would you actually build an application like the Vacation Assistant?

Lama Nachman: Absolutely. Basically, the diagram that you're seeing here is really this kind of architectural framework of how we can build these context-aware applications. So really at the heart of this, there is this context framework, which you've actually seen in the Fodor's demo.

Justin Rattner: Right.

Lama Nachman: So there is the context engine piece here, which is really kind of a platform service that's running on the platform. And what it's really doing is, at one level, it's providing this consistent API to be able to bring in this rich data of context –

Justin Rattner: So kind of a middleware layer, if you will.

Lama Nachman: Exactly. So basically, you can get data from sensors, you can get data from applications, from web services. It can go from all of these different sources, and you have this consistent API to flow that data in.

Justin Rattner: Okay.

Lama Nachman: At the same time, we talked about all these inference algorithms that we need to run to infer this data, and there is this notion of an analyzer where it's really extensible. So you can build all of these different algorithms within that analyzer and really leverage across different types of contexts, as well.

Justin Rattner: Okay. So we can have a whole collection of inferencing algorithms, if you will, to help us do this.

Lama Nachman: Absolutely.

Justin Rattner: Okay.

Lama Nachman: And then for the application writers, there is this API that provides a way for the applications to extract the context. So you don't necessarily have to know how the context is being inferred --

Justin Rattner: Oh, okay. So I don't have to be an expert in analysis or inferencing.

Lama Nachman: Exactly. So you could say, for example, something like either, "Let me know what the person's activity is," or you can say something like, "If the person's activity changes, let me know."

Justin Rattner: Sure, sure.

Lama Nachman: So all of this is part of that context, the engine. And then there is this notion of sort of a context proxy service that's running in the cloud that actually is responsible for sharing context across the different devices.

Justin Rattner: Oh, okay. So that's the aggregation you were talking about.

Lama Nachman: Exactly.

Justin Rattner: Both over time and over devices.

Lama Nachman: Over devices and possibly over different people.

Justin Rattner: So I don't have to be walking around with this complete context that might actually be quite sizable over time.

Lama Nachman: Absolutely. And then, furthermore, you can take that contextual information that's coming from the proxy service and then feed it back into your context engine, to actually use that to infer more stuff.

Justin Rattner: Terrific. Now, it's pretty clear that as we aggregate this context over time and over devices that the information is getting pretty valuable. I'm not sure I want the world to know my television preferences, for example. Does the framework do anything to protect the context information?

Lama Nachman: Absolutely. So, again, you can see here in this red release policy, right, so the idea here is that the user is in complete control. So they can control what context gets shared or released, to whom that context gets released, and up to what time you want to actually even enable that release to happen and when do you actually want to expire the data.

Justin Rattner: So digital rights management for context.

Lama Nachman: Exactly.

Justin Rattner: Okay. Well, that's fantastic, and I really want to thank you.

Lama Nachman: My pleasure.

Justin Rattner: I really appreciate you coming out here and joining us and giving us sort of this under-the-covers look at context-aware computing.

Lama Nachman: Any time. My pleasure.

Justin Rattner: Thanks again.

Lama Nachman: Thanks.

Justin Rattner: Bye now. Lama Nachman everyone. Well, I couldn't get away at an IDF keynote without saying something about the platform implications of this topic. I think Lama did give us a pretty good understanding of what's involved. But it prompts the question; what does this mean to platform architecture? What kinds of capabilities do we need to be thinking about and designing in for the future?

And I think right at the top of the list is this notion of always-on sensing and inferencing. You know, we can't just let the devices go to sleep and wait an indefinite amount of time. They really have to be at least in a state where they can sense what's going on around them and be able to detect, most importantly, the transitions that take place in our immediate context and wake up other parts of the system.

It also means that this whole sensor capability has to be designed for very low power. If we at least need to keep the sensory aspects of the system up and running all the time, we need to do that at minimum power. And, of course, when we think about the processing technology, the computing technology, we're going to have to look at features and capabilities that optimize performance now for the sense-making part of the inference pipeline. You know, the kinds of algorithms that Lama was just talking about that turn the raw sensor data, be it hard or soft sensor data, into meaningful results. You know, these conclusions about our current context.

So lots of implications on the platform, and I think you can expect to see features that support context aware computing starting to appear in Intel products in the not-too-distant future.

Well, that brings us, I think, to a very interesting point in the conversation this morning, and that's really this question: Is the market ready for context? I think Bo did a good job talking about why it's taken so long to get to a point where we can think seriously about deploying context-aware technology. The question is will people want it? Will people buy it? And to answer that question or at least get a start at the answer, we went on the street and asked people what they'd like to see in their devices. Let's watch this video.

Well, to help us understand what people want, there is no better person than Intel fellow and director of interaction and experience

research at Intel Labs, Ms. Genevieve Bell. Come on out, Genevieve. Here she is.

Genevieve Bell: Hey, Justin.

Justin Rattner: Good morning.

Genevieve Bell: Hello, sir. Nice to see you again.

Justin Rattner: Good to see you. Yeah. It's perfect. So we've heard a lot about the technology.

Genevieve Bell: Appropriately.

Justin Rattner: As a social scientist, you live with this every day, of course, at Intel. There are all these technologists scurrying around. And we've talked about possible applications. But how are we going to find out whether people want all this context awareness?

Genevieve Bell: Well, you can see just from that little, tiny video shot on one day in a street in one town in America that people have incredibly ambivalent and complicated relationships about what it is they want their devices to know about them. I think working through that problem and making sure that we're actually applying context in the right ways means actually working out what it is that people are going to care about and what they love.

Justin Rattner: Okay. What they love? Yeah. Tell me more.

Genevieve Bell: I could tell you all kinds of things about that.

Justin Rattner: Okay. After they keynote, please.

Genevieve Bell: Okay. In a different context?

Justin Rattner: In another context.

Genevieve Bell: Exactly. So I think you know -- as you know – because you're now my boss, and you've been in charge of these things for Intel for a long time – Intel has always had a kind of interest in and a center of excellence around ethnographic research. So studying what it is that people do, what they care about, what's important to them, what motivates them, what frustrates them, and we actually have a process for doing that that really is, in some ways, driven by getting a sense of what it is that people love. What it is that they care about?

Justin Rattner: Well, you know, love is not something the average hardcore Intel developer deals with on a daily basis. So you're the love doctor. Tell us how we come to know what people will love.

Genevieve Bell: Wow. Is that a promotion?

Justin Rattner: There will be something in your pay envelope.

Genevieve Bell: Excellent. So, I mean, for us, you know, part of the reason why you want to talk about love is it's incredibly important that people have emotional relationships. And if you listen to and think about all of the things you've seen on your stage today, part of what all of that's about is creating things that are going to be in people's lives, that are going to be about monitoring them and helping them in their social relationships – don't interrupt your wife in a meeting, make sure you've found a good restaurant. And as we work through this whole process of spending time with people, what we're really after is making sure that we understand what people care about, what's going to motivate them, and making sure that we can actually then do something for a technical company with all of that work.

Justin Rattner: As I said, we were talking to developers this morning, and I think they really want to know whether it's worth all the bother. What's the path to understanding what people love and, more importantly, what people are willing to pay for?

Genevieve Bell: I think one of the things we've actually worked out how to get right at Intel is how to go from that sense of what it is that people love to actually a set of compelling experiences that we can use to motivate new technology development. We actually have a process for this now that looks remarkably like an engineering process, I'm happy to say. Look at that. Where we actually go from thinking about what we know about people – so what's that going to say about the experiences they're going to care about? How do we create a set of

compelling usages? How do we test that with consumers, both in our labs but also in people's lives? How do we give people stuff to go play with and work through where all the kind of messiness of daily life happens?

Justin Rattner: Okay. So that's great, and I know you kind of reduced this to practice. But how do we answer this question of whether context-aware computing is really going to make a difference in terms of user experience?

Genevieve Bell: Oh, I think part of it is we need to go and actually talk to people more about this stuff. I think it's also about how we go through this process. But I think it's also about a mind shift change, and it's one that – you know, it's a conversation you and I often have, which is the engineers talk to me about context, and I want to say, well, everything you talked about today is context at one level, but at another level it's just what people do in their daily lives.

And so I brought some photos here, because I couldn't take you to all these places, to tell you a little bit about why context isn't just a technical problem. It's also everyday life.

Justin Rattner: Okay.

Genevieve Bell: So the photo in the bottom corner here is in my favorite cricket ground in Australia, a place where one should spend one's summer workdays, as it turns out, on your Bluetooth headset with your beer

and your phone. Which of course creates some interesting problems about context, right? As you're sending photos of the cricket, do you send them to your boss or your friends? What does that all look like? How does your phone predict what you're doing and what you should be doing?

And all of that is sort of in the context of both blending work and life. We know that there are other places in life where multiple devices are starting to cluster, so we have a photo here of someone's living room in Manchester, in England. Where -- I'm sure you're guilty of this too -- a laptop and a television.

Justin Rattner: Right.

Genevieve Bell: And we know there are lots of complicated rules there about who controls what's on the television, what you're doing on the laptop, and starting to think about how all of that is going to play out and what context will mean there.

Justin Rattner: Tell me about the Buddhas.

Genevieve Bell: From a temple in Busan in Korea. Every single one of those little Buddhas is someone's wish or someone's prayer. And when you come to this temple, you leave them there. And thinking about what it would mean to have a temple that was context aware or a Buddha that was context aware is kind of a wonderful notion about what it's going to mean to think about space and place.

And it's actually why I have this other sign here with the line through the mobile phone. It's also from Korea, and the sign says, "It would be a blessing if you turned off your mobile phone." Of course, what they don't say is the church also has a cell site dampener, so it's kind of a zero sum for you turning off your phone.

Justin Rattner: As long as we can leave the sensors on, I think we'll be okay with that.

Genevieve Bell: Absolutely. But I think what that one says to me is that we're going to get to a place where context awareness is also going to be about where things aren't going to happen as much as where they are.

Justin Rattner: Right.

Genevieve Bell: And so creating spaces where you're not being tracked or where you're thinking differently about what you're doing is also going to be a really important part of this. It's not about everything being enabled, but also working through what the social rules are as well as kind of the --

Justin Rattner: Well, knowing when to turn off those annoying instant messages that are always popping up on your computer screen in the middle of a big meeting.

Genevieve Bell: Absolutely. Yeah, and knowing that, you know, the sensor on your ankle should probably stop tracking your gait right about now.

Justin Rattner: Okay. Really, I guess this gets down to the nitty-gritty. I mean, do you really think context awareness is going to lead to a better user experience? You're the expert here.

Genevieve Bell: Well, I think it's absolutely the case. Right? I mean, it is my very strong sense that if we get context right – even a little bit right – what it does is propel an entire new set of experiences. It will be what distinguishes this last generation of smart devices from the next generation that are much smarter. And I think getting what it means to talk about context – devices that know more about us after we've owned them for a year than they did when we first bought them – is all about what it's going to mean to make something that's incredibly compelling for consumers and that delights them.

Justin Rattner: Outstanding. I love it. And I'll get more of this love advice in a little while.

Genevieve Bell: I'll be out back.

Justin Rattner: Thanks. Genevieve Bell, everyone. Okay. I know I should be putting on a black T-shirt, but lacking a black T-shirt, I've got my Intel Labs jacket. Because now we're going to go way out. All right. I got that. I don't need the turned up collar here.

Okay. All right. So if you were here a couple of years ago at the IDF keynote where I took this 40-year look into the future, you may remember that we showed a very interesting technology, which took the form of a piece of headgear by a company called Emotiv, that let you play a simple game just by thinking about your next move in the game. As provocative as that was – and it turned out to be extremely provocative at the time – what I'd like to show you last this morning is the next generation of that idea. Well -- and it's the reason for the lab coat – maybe a few generations beyond that.

The ultimate form of sensing, of human-machine interaction, would be the direct understanding of human thought by machines. And it is possible, and in fact progress continues to be made. We're moving past where the Emotiv headset was, really, to a future where machines understand what we're thinking. Let's watch this video.

Amazing stuff. That's probably IDF, well 2020 seems like a reasonable stop for that, but pretty amazing. And the fact that, you know, we can do those basic experiments -- in fact, we can use the inferencing pipeline that Lama talked to us about -- to understand human thought is truly amazing.

Well, I hope from all of that, you're as excited as we are about context-aware computing. I think we're shown you what it can do to create a whole new generation of applications and experiences.

We've shown you the underlying technology or as much as we could get to in the time we had this morning.

But most importantly, I think we've shown you how it can create devices that people really love, that become indispensable in their daily lives and even to the point where you're giving them affectionate names. And that's really why we think that context-aware computing is the next big thing, and we hope you think so as well. So thanks, and we'll see you next time at the Intel Developer Forum. Thanks, everyone!

Female Voice: Ladies and gentlemen, technical sessions will begin in 20 minutes. Press and analysts --

[End of recorded material.]