Transcript of ARM TechCon 2013 Investor Event
October 30, 2013

Corporate Speakers

• Ian Thorton	ARM	Head - IR
• Simon Segars	ARM	CEO
• Tom Cronk	ARM	General Manager - Processing
• Dipesh Patel	ARM	EVP, General Manager - PIPD
• Charlene Marini	ARM	VP - Embedded Segment
• Lakshmi Mandyam	ARM	Director – Enterprise Segment Marketing

Participants

• Simon Schafer	Goldman Sachs	
• Ian Huh	Samsung	Director Marketing
• Sander Arts	Atmel	VP - Corporate Marketing
• Andrew Feldman	AMD	Corporate VP, General Manager

CONTENTS

Introduction

Simon Segars: Afternoon, everyone. We've run this for a few years now, and it's proven to be a great forum, just to have some interactive discussion about what ARM and its partners actually do together. You saw in our Q3 results, we had a fantastic set of numbers there, revenues etc, and that comes from the work that we do with our partners, the work we do in putting these chips together, and the work we do throughout ecosystem of making it as easy as possible to actually use the semiconductor devices.

So we try and use this event to put some members of the ARM team together with key partners to talk about the actual experiences and work that we do together. So today we've got four presentations. We're going to cover first of some of our processor work, so on my right here is Tom Cronk, he is the general manager of our Processor Division. And with him is Ian Huh from Samsung. They're going to talk about some of the work that's going on in Samsung's implementation particularly around big.LITTLE.

The things that we spend a lot of time thinking about are applications and manufacturing silicon. The people who we license to are in the business of building chips, and that is a nontrivial exercise. And so after the first presentation, we're going to have Dipesh Patel. He runs the Physical IP business out here in the valley, and he's going to give you an overview of what's going on in the manufacturing landscape.

We're going to take a break, then we're going to talk about two ends of the computing spectrum, starting with Embedded. We're going to have Charlene from ARM and Sander from Atmel, talking about what's going on in embedded and what's going on in IOT. If you were in my keynote this morning, hopefully you left with a sense that we think this is quite a big market and a big opportunity and Atmel has done some real pioneering work there.

And then we're going to talk about servers, which is the other area. This is particularly exciting as the ARM world evolves. We have Lakshmi from ARM and Andrew Feldman from AMD going to talk about the progress that's going on there.

So hopefully this is going to give you a good opportunity to cover the – the spectrum of ARM devices and provide plenty of opportunity for – for discussion. I'd like you to come back to the Mission Ballroom; I'm going to be taking part in a fireside chat with Fortune Magazine, and so if you can stay for that, that'd be great, and we'll have a few drinks afterwards as well, so I hope you enjoy the rest of your afternoon.

With that, I'm going to hand it over to Simon Schafer from Goldman Sachs who's going to moderate this afternoon and keep you guys under control. Thanks very much.

Simon Schafer: Thanks, Simon. I'm chairing this event today, and I'm going to moderate some of this, but the idea clearly is to ask lots of questions. But without further ado, thank you very much Ian and Tom for doing this afternoon and I think the two of you together are going to kick us off with a couple of slides, just for introduction, just to set the scene, and we'll take it from there.

I will moderate some of the questions, but the idea really is for all of you guys to stick your hand up whenever you want. And someone will come around with a microphone and allow you to participate in the forum, that's the idea. And yes, we'll take you through the sessions as Simon said, a couple of sessions now, quick break, two more sessions after that and then we can go onto the keynote.

SAMSUNG and big.LITTLE

Tom Cronk: So I've been asked to talk a bit about big.LITTLE, and let me just start by reminding everybody what that actually is. It's an interesting technological innovation that's come from ARM, something that we had in development for probably five or six years. And it's really come to the fore in the market very recently, initially in the form of some Samsung products.

It builds on a basic paradigm which says there are two fundamental drivers for processor design, particularly in a mobile environment. There's a class of applications that require very complex algorithms to be processed, and often very, very quickly. And that tends to drive the need for high-performance cores. High-performance cores are typically quite large, so that's the "big" part.

That's driven by things like video, the initial rendering of very complex web pages, gaming, for example. These might drive the need for high-performance, and, of course, so does benchmarking, somewhat artificially. But if you want to sell your product, you have to run good benchmarks. That will drive the need for a big high-performance core.

At the other end of the spectrum, though, there's a drive for ultra-energy-efficiency, and this comes from tasks that tend to be on all the time: monitoring your Facebook account, monitoring texts, a whole plethora of always-on computing. And this drives the compute design process and microarchitecture design in a different direction. A very simple generalization is that the smaller the core, the more power-efficient it is, and therefore, the longer the battery life of your product.

These two things are clearly diametrically opposed from an engineering perspective. So big.LITTLE, simplistically, is a paradigm that allows you to couple a big core and a little core together in a system, then automatically the software runs either on the big core or the little core, depending on the profile of the software or the task at the time.

Therefore, when you need to have a very fast response or run a very complex algorithm, the software will fire up on the big core, but most of the time, most of the software in these phones sits on the little core. The little core will typically be 70% more energy-efficient. We're seeing energy saving in the order of 50% to 60% in system, using round numbers, through the implementation of big.LITTLE.

And so with the recent advent of the latest generation of the software which we call big.LITTLE MP, or HMP, some of you may be aware of it, the technique and the philosophy's really now coming into life. It's very synergistic with the Android kernel, for example, and we're seeing it deliver real benefit. So that's by way of the scene-setter. We'll hand out to Ian to talk a bit about deployment in the market for real.

Ian Huh: So I'm from Samsung, and obviously we work very closely together with the ARM team. And, I mean, Tom just mentioned big.LITTLE technology, HMP stuff, but I would add that it’s a very advanced technology to implement because we are talking about microseconds to switch tasks from one processor to another based on what you are doing.

So we introduced the initial big.LITTLE product at CES in January, and the technology first came into production in the Exynos 5410. We have just launched another product, the Exynos 5420, which is the best version of big.LITTLE and that product is getting to some products such

as the Galaxy Tab, Note 3, et cetera. And again, as Tom had mentioned, we just completed work for HMP stuff, which is giving us at least 20% to 25% increase in performance, plus a similar amount of power savings, basically.

When big.LITTLE was first launched there were critics saying ‘is it really going to be Octa core because we cannot run eight cores at once?’ But now we can switch cores by any combination, which means we can activate any combination of cores, from one to eight. So that is the current situation right now, and we are really excited and we hope to be a really good partner with ARM in that regard. Thank you.

Simon Schafer: Maybe, I can start when you mentioned that the implementation is very, very difficult. You are the first to try it, but of course, I think ARM has got something like 16 big.LITTLE partners. You're one of the lead ones. But talk a little bit more about some of those challenges that you've seen in just the implementation stage.

And the reason I ask is that, you know, clearly some of your competitors are a lot less forthcoming about the real benefits of this architecture, so maybe just as a reminder to why you think that combination is really what allows you a competitive edge, and why you're better at implementing it. And perhaps in a little bit more detail, just address some of those implementation challenges.

Ian Huh: OK, so I don't really know what other competitors are doing that much. I mean, they probably have their own technology to implement, to solve these type of a performance versus power problems. But obviously, we felt that this solution, I mean this combination of little cores and big cores, is going to be the real solution.

I mean, for example, in every-day work, pretty much 80 to 90% of work has to be done in small cores. Big cores don't have to be awake at that time period. But then you start some applications, or browse web pages, or flip your flipbook, then the big cores are come in and enable a much smoother operation.

So again, the beauty and the challenging part is scheduling. So whenever you have whatever task, we have to calculate the workload very accurately, and have to switch cores in milliseconds – a very, very short time period. And again, everybody has different usage behavior, so calculating and preparing for all those types of usage scenarios has been very difficult for us. But again, I think we are in very stable stages right now, and you're going to see more big.LITTLE products from Samsung from now on.

Simon Schafer: Tom, your perspective on that? I mean, we've had a big endorsement from one of your lead partners, but again, you haven't really been able to convey that wisdom to everybody. So how do you fix that and how do you provide a roadmap that everyone in the landscape actually adopts that point of view?

Tom Cronk: So I think there are a couple of dimensions to it. One of them is timing. Let's start, though, with one of the great powers of the ARM partnership. And the ARM partnership model is people tend to make choices about implementing things and often implement similar things in different ways, and to some extent, that's what we've seen so far with big.LITTLE. So if I use, NVIDIA as an example, they had their 1+4 technology, which they first talked about maybe three years ago.

But the basic paradigm was it makes sense to put a core in the system which is optimized for energy efficiency, and it makes sense for other cores in which were optimized for performance. Now they chose to make the difference in the implementation rather than in the micro architecture, but at the highest level, it's philosophically a very similar approach. And again, with Qualcomm, they do a very similar thing. They don't call it big.LITTLE, but the basic premise, the underlying physics that are driving it are the same, so they've implemented similar things in a different way, called it something different, and that's just fine.

The other thing was the timeline. We've worked through two, three generations of the software on big.LITTLE, and we've now got to the big.LITTLE HMP solution. We've arrived at the place where it's basically invisible in the kernel and eventually becomes part of the OS. I can't talk for what people will do, so I'm not making any assertion as to what other partners will do in the future, but I have to say as an engineer, it's hard to see why you wouldn't adopt it. If you've already made the choice that you're going to make a multicore solution, why wouldn't you implement the software that made it easy? And by the way, there's no physical change required in the hardware underneath in any of the implementations; it's just a software change.

So the original custom migration implementation through the CPU migrations get us to where we are now with big.LITTLE HMP. Samsung have really blazed a trail, and I think that's great. We really appreciate their support and partnership in doing that, and I passionately believe it's the way forward.

Simon Schafer: Any other questions from the audience on – on that big.LITTLE? Go for it.

Unidentified Audience Member: Two questions. Please could you tell us about ARMv8-A big.LITTLE and the timing of the ramp at 16 nanometer. And maybe one for Tom in terms of graphics. big.LITTLE, you're kind of running two streams separately to graphics at the moment, and I just wonder whether you'll see something big.LITTLE-like for graphics in due course? Thank you.

Ian Huh: So the first question is about Cortex-A57, Cortex-A53, and whether we are going to use big.LITTLE for that architecture, and the second question was about 16nm, was it? Okay, so Cortex-A53, Cortex-A57 is our next generation as far as I know.

After the Cortex-A15, Cortex-A7, I think we are looking very closely to implement this, we need to produce product, using Cortex-A57, Cortex-A53, but I don't think I can say about the future of product roadmap precisely at this moment, but I think we are pretty much positively looking at it to use it, because again, the fact that end users are using most of the small cores most of the time and big cores are needed for some performance peak, that fact is going to hold for the future. So as far as that holds true, we are going to use the big.LITTLE technology stuff.

Regarding 16nm, we don't really don't do 16nm. We are doing 14nm, I think the expected date, for mass production is going to be end of the 2014 or early 2015 for 14nm process technologies.

Simon Schafer: There's a question on – on eventual big.LITTLE configuration for GPU.

Tom Cronk: Yes, and a good question, very insightful. It's something that we're obviously looking at. I'd go so far as to say we're experimenting with it. It is a logical progression, and as people find more and more things to run on their GPUs, and then the notion of GPU Compute, I can see it makes some sense, yes.

Unidentified Audience Member: The mix of Exynos in the Galaxy S4 seems like a little lower than you might have liked. Can you talk about why that might have been, and how you expect to be relative to the S5 or will there be things that you'll be ahead on so they are more in your favor?

Ian Huh: So obviously we have our own assessment on that regard, but that decision was made by our customer, which is Samsung Mobile. And as you know, we are operating very separately right now, so unfortunate I cannot answer how they made these decisions. I mean, we have our own assessment, but you should actually ask that question to Samsung Mobile, I guess. Sorry about that.

Unidentified Audience Member: From the earlier question about combining the GPUs and CPUs to create a heterogeneous architecture: when you're looking at the memory architecture, would that be something on its own or would the big.LITTLE be in combination with a GPU, would that shared memory architecture be in one bucket? Any thoughts there?

Tom Cronk: Yes, typically the GPU does sit with the same memory hierarchy as the processors, and the interconnect that we provide along with the processors actually supports that. We have deliberately set up some of the memory, things like the way page tables work and stuff like that, so there's a synergy between the processors and the GPUs, so I think it would just be a continuation of that.

Simon Schafer: There's another question over here.

Unidentified Audience Member: Thank you. Just thinking about the different segments of smartphones, I mean obviously it makes complete sense in the premium segment. How about like your middle and very low-end smartphones, using big.LITTLE, and how much incremental cost would it add, for example?

Ian Huh: Of course this is very different case by case but big.LITTLE, I mean, the cost of the big.LITTLE is obviously you need a little bit more silicon space, and silicon space means there will be more cost from our side; so far we have just been focused on the premium segment which has room to absorb that cost increase.

But now we are seeing more and more pressure, even in mid-range segment for saving power, and basically the rule in this industry is this year's premium goes to the next year's mid-end. That sort of the thing is happening every year.

So since we started this type of premium technology this year, I think next year there will be more of this technology in high segment, probably not for low-segment, but phones, for example, like $300+, you might want to see this type of technology for there.

Tom Cronk: Yes, we definitely see it coming down into the mid.

Unidentified Audience Member: Is there low-end opportunity in the future, or is it something that's just going to be exclusively that domain?

Tom Cronk: I certainly haven't considered it coming up all the way down to the low-end. But I'm going to consider it. I mean, we're about to release the Cortex-A12 next generation which will support big.LITTLE, so it is conceivable that you could do a Cortex-A12, Cortex-A7 LITTLE.

Unidentified Audience Member: But just on the same thought, how about 64-bit, because the power savings, even running legacy apps, is better using the v8 core is what I've been reading so far. More power savings because it's a better architecture? Does that make sense on the low-end or does it just make sense to stay with 32-bit?

Tom Cronk: There's no doubt that Cortex-A53 is coming off the train in the low-end. Yes, I mean Cortex-A53 – as with Cortex-A57 – was designed to be fundamentally a very good 32-bit core and a natural progression from a Cortex-A7. And then of course, with the recent heightened interest in 64-bit it happens to be in the right place at the right time. So I – I suspect the marketing on it's '64-bitness' will probably accelerate, but Cortex-A53 into the low-end is a done deal, if I may say that.

Simon Schafer: I'd love to make you stay on the 64-bit topic. We talked about big.LITTLE quite a lot, but I think even ARM was surprised that 64-bit happened as quickly as it did. But I think on the Samsung conference call last week, you actually became more public about your own intentions on 64-bit.

So – so maybe talk a little bit about the rationale and how you're thinking about the adoption curve for that. Is it just a marketing gimmick or do you genuinely see such a performance advantage that it is actually a necessary step to make. There's obviously some very mixed emotions about that architecture genuinely bringing any advantage whatsoever at this point, just given the lack of OS and so on and so forth.

Ian Huh: So here's our assessment, basically. In SoC hardware level, there should be some benefit, but I think the most benefit is coming out of the application developers. It actually provides the tool to develop an application more efficiently, for example, security password encryption, for example. They can program more efficient graphics by using better, bigger and a floating-point calculation. It really depends on how they utilize the application, the coding, that 64-bit datapath. And we are just providing the infrastructure and baseline to make that happen.

So the real answer is going to be actually depends on the applications developer and how operating system makers such as Android or Tizen or those guys are doing. So it's really ecosystem, not really the implementation by humble SoC guys like me, I guess.

Simon Schafer: And just as a reminder, and I think some of the commentary on the Samsung conference call was just sort of a headline, hey, we're working on 64-bit, but is there a more concrete time line that you as a company have actually outlined as it relates to adoption?

Ian Huh: We don't have any concrete time line. I mean, this is relatively new and we sensed that the customers that demanding more features we probably going to accelerate the development speed, but there's no clear definitive time line at this moment.

Simon Schafer: Any question from the audience on 64-bit?

Unidentified Audience Member: 64-bit will kind of combine different parts of Samsung, just because you guys are in a unique position where you've offered a tablet using both ARM and Intel processors, and --

Ian Huh: Not me!

Unidentified Audience Member: Just framing that with – with the big.LITTLE, and 64-bit, I mean, how – how does it stack up? How does this stack up against Intel's offering for tablets? I mean, I see your phones, but on the tablet side, how does that look? The benefits?

Ian Huh: So again, it's not really me that uses that Intel processors in a tablet; it's one of our customers. I don't really know, and can't comment on, why they choose Intel over ours in some of the models.

But we believe Intel has a long way to go to compete in this mobile space, because it's not just about the product and price; there's an ecosystem behind that. So in the tablet and smartphone space, thousands and thousands of applications have been written to optimize their performance on ARM-based processors. And for Intel to catch up… of course, I really respect the company. That company has a lot of money, people, software engineers, all kinds of stuff. But it's not really like they develop better products, and would be the best on price. Catching up the market is not going to happen, you know, that's my assessment.

Simon Schafer: Yes, please?

Unidentified Audience Member: As the usage of big.LITTLE moves to mid-priced products, do you see price pressure on Cortex-A12 products? I mean, will price come down as you move to more mass-market products in the future?

Tom Cronk: Well there is always price pressure, clearly, it goes without saying. But the way we've set the business up with respect to the Cortex-A12 products, mobile is just one of the application domains, which is a big driver following up from Cortex-A9, but we've licensed Cortex-A9 north of 60 times, something like that, into many, many domains, so there's a lot of places that Cortex-A12 will go with different markets, but we set the business up such that we came in slightly lower with our pricing expectation on Owl at the start, so I don't see any negative impact, if that's where you're driving on the question.

Unidentified Audience Member: Thank you. I had a question for the gentleman from Samsung. I was wondering if you could share your perspective on what Intel tell us is when they get to 14nm, they're going to shrink the time between when they implement that process node for their core CPUs and the SoCs, so I just wanted to get your perspective on, we've heard ARM talk about all the innovation that ARM is doing from their side to counter that. But what's what your perspective? And when you look at that, say, OK, competitively is that changing something from the process side that you will have to counter for your business? Thank you.

Ian Huh: So, your question is, are you going to just to be investing on the implementation, things like that? I mean, could you just ...

Unidentified Audience Member: No, I just wanted your perspective on – just share your thinking on how true is that contention from Intel that at 14nm, with their lead, it will get harder for the ARM camp to compete.

Ian Huh: So, I mean, that's not really my specialty, talking about the – the process node, but again, I think we in general as a company, we believe in the process node, having all the best process node: we are doing 28nm right now, and 20nm, probably end of this month, end of next

year, probably around 14nm. That is very much part of the overall competitiveness. So now, regardless of what Intel is doing, we take that effort very, very seriously.

Because in the memory space, ultimately the node is pretty much about shrink and cost, but in the SoC space, ultimately performance means more for us, less power, so in order to stay ahead of our competition, I believe helping that ultimate performance and if you ask me how much you are better or worse than Intel, I don't really have that assessment at this moment.

Simon Schafer: Can – can I ask that Dipesh maybe comment to that? Dipesh runs the Physical IP division in ARM as you probably know.

Dipesh Patel: Yes, I was going to say, maybe we could take that question when I talk next, because I can answer that from my Physical IP perspective. There's a lot of work we do in the division, working with Samsung as well as our processor team to address that exact point, which is making sure that compared to Intel, we are not at any disadvantage.

Unidentified Audience Member: Thank you.

Unidentified Audience Member: Thank you. Two questions. First, I wanted to find out, in the first implementation of big.LITTLE, what went wrong? What did you have to resolve before the second version that you announced recently? Second question, earlier you mentioned that software changes were key to big.LITTLE. Which – can you detail what software needs to be changed? I imagine it's at all levels, but maybe you can comment on which part of the software stack needs the most changes? Are Android changes needed? You know, Samsung, ARM, so on and so forth?

Tom Cronk: So, nothing actually went wrong, but what's happened is that it has evolved and improved very rapidly as we've worked in partnership lessons from the deployment some obvious optimizations came to light. So the fundamental difference is the first generation of the software was set up to do what we call cluster migration, so you either light up, so in a LITTLE core cluster you have four cores in the extreme, same on the big, and in the first implementation, the paradigm is simply either those four are up, the LITTLE four are up, or the big four are up.

So that was designed to actually make it very easy to implement, because all you had to do was go underneath the operating system, not touch it at all, and – and go into the code that controls the power management, so it was a very simple software implementation, but it was a very coarse solution: all up or all down. And it works, and it shows benefit, and we can show you energy savings and that's fine.

Coming to where we've ended up, what you actually do now is you go into the operating system kernel and you effectively get a very effective HMP implementation, so it's much finer granularity, much more synergistic with the OS. You're able to draw out intelligence that the OS has already got, and you end up with a more flexible and refined solution.

As I said at the start, HMP, or MP, big.LITTLE MP will run on the original cluster migration hardware. There's no hardware change. So it's just – it's a way the software has been implemented. It's simply slightly different, and more enhanced, and the result is we now have a system that will deliver higher performance, because you can use all of the cores, so four plus four, as opposed to just four, so you get a maximum peak performance benefit.

And you get energy saving because of the granularity and the fast context which it's just evolved. We're pleased it's evolved very quickly.

Unidentified Audience Member: I just wanted to kind of follow up on the last question here about some of the same issues as at the high end of the market here. You're moving 64-bit into big.LITTLE sort of at the same time with different partners, et cetera. The first big mover to 64-bit controls their own vertical ecosystem completely, and I imagine that takes a considerable number of variables out of the migration for – for them.

Maybe you could comment on, I guess, the 64-bit migration into Android from – from either of your perspectives, given the fact that there are other chip suppliers in the ecosystem that are not going to be doing big.LITTLE at the same time and have to leverage some of the same kernels, the maturity of Linux for big.LITTLE and 64-bit, et cetera. It seems like in Android, and then not to mention Windows, et cetera, the migration gets a little bit more complicated.

Tom Cronk: I mean, yes, it's true, in the sense that it's a much more open environment, so there are more dimensions to deal with, but with respect to the schedule, I would really prefer to defer that to Google, because it's down to them to decide how quickly, and how quickly they want to move, and then which segment.

Ian Huh: So I just want to add my perspective over there. If you see the past PC industry migration from 32-bit to 64-bit takes probably five to six years, I guess, for us to tape out, the operating system is ready, and then development capability is ready, and people are finally start coding with 64-bit took like five years, something like that, I guess.

Since Android is open there are a lot of players involved. It's probably going to be harder than somebody who can actually control everything, but I think it's going to be much, much faster than what happens in PC space because we can see a lot of demand for customer side anyway. So, everybody's going to probably move very fast.

Unidentified Audience Member: One follow-up question on the benchmarking we talked about, the atom earlier, the question was asked, if you look at the Atom versus the Silvermont versus the current generation, how would you say the benchmarking earlier was brought up for its single-threaded application would compare?

Tom Cronk: The benchmarks that we're seeing are measured off real products in the market, so if you're looking at the performance dimension on synthetic benchmarks, the performance benefit is anywhere between one-and-a-half to three times the Intel solution, so greater performance.

If you flip it the other way to energy efficiency, then it's a similar spread. And of course, it's worth noting that that's synthetic benchmarks, so that's basically as good as it gets on an Intel solution. In the real world, most of the time they're having to run emulation, and that's just another layer of software that sits in between everything else and consumes power, so – and it's a long way ahead of where they are at the moment.

Simon Schafer: I'd love to steer the debate a little bit, on the cost of all of this, and I'm sure – and Dipesh touched some that will come back in the context of manufacturing. That's obviously a big part of it. But one of the things that we have, of course, seen is frankly it's getting more expensive for the customer.

So I know, you know, we get the ROI, but at what point do you start to think, "Hold on a second. I'm getting some performance benefits, but perhaps the – the price that I'm being charged is no longer the same that Robin Saxby first offered to my company in the late-1990s, because we've moved on a long way from then in terms of the pricing. So I'd love it if you could talk to that a little bit.

Ian Huh: So we talked a little bit about the pricing pressures and as everybody knows, ASPs of mobiles, mobile devices are obviously dropping right now and people are talking about a stagnation of the premium smartphone. So does that mean more pressure to our AP business? More pressures obviously, but again, we believe, it really depends on end-user scenarios, and innovation.

Not really in my – my area of – again if you see the material cost of the typical smartphone, I think the AP price is somewhere between, I would say, 10 to 20% I guess, in between, right? And there's a memory cost and others. So the challenge is of course ASP is going to come down, as price is coming down, but how much cost benefit we can actually provide to other components such as memory, power chips, et cetera.

So how do we increase the percentage of the AP price, without costing us our jobs. We believe with all types of these power saving technologies, the high-end gaming, 64-bit all kind of stuff, I think there is a fair chance that we could actually improve that percentage within the overall cost.

Unidentified Audience Member: I guess maybe I – I didn't explain the question very well. My – my question was really more around the discussion you would, the two of you would have now that you're being offered, you know, big.LITTLE configuration and 64-bit, that stuff, you being charged more by Tom next to you than you were in previous generations. So I'm thinking more about that aspect of price inflation as opposed to price deflation in your AP, so --

Ian Huh: You are the customer. I really hope you can get it for free, but that's not the case.

Tom Cronk: Whenever I ask I would always turn it back to the business model, you know, the business model is fundamentally very powerful in that really, most of the price cost gets transferred when the customer's getting value from shipping product and having a successful business in the form of a royalty. So if the customer's having a successful business, and they're paying us a royalty, then we're both winning. And that business model works well.

Simon Schafer: There's a question here.

Unidentified Audience Member: We've heard, prior to this, that there is an Intel tax of some nature, possibly that prohibits its use in mass markets, it centers potentially around decode, where 30% of the power budget might be in decode. Perhaps, this is not the case, but I was wondering, is there any merit to this argument? Is there something Intel can do around decode?

Tom Cronk: I mean, I've never heard it expressed in – in the context of decode, but their architecture carries a lot of sort of high-end desktop legacy, whereas ours is architected originally for deployment in very low-power environments, so I think in terms of power efficiency, we gain some benefit from that right at the top, and that obviously flows all the way through.

You look at what they've recently announced with their Quark solution, which is supposed to be a small core, little core. That's still a long, long way from where we are with, say a Cortex-A53. In ARM-speak, that would be a big core. So there are clearly technical challenges there for them.

Unidentified Audience Member: Can you just give us some sense of the size of the market? Say, how big is big.LITTLE and then how do we think about the three to five-year trajectory, and what are some of the things that would really drive it?

Tom Cronk: So there are about 18 licensees with big.LITTLE capability. I'm not prepared to say how many are in track and when they're all going to pop out, but at the moment there are about two or three public.

Unidentified Audience Member: All right, let me be more explicit. For this to be successful, do you need the other two leading AP designers to implement it?

Tom Cronk: I don't believe so. As I said, I fundamentally believe that this is where the physics take you. The benefits are obvious, so obviously the more they're adopted, the quicker, the easier life is, but for me it's just a matter of time.

Simon Schafer: Yes, there's one more question is back here.

Simon Schafer: One question in the back, please?

Unidentified Audience Member: Yes, on the Exynos 5420, one major change you had was implementing Mali to the expensive Imagination. Can you kind of talk about the reasons for that decision? And then going forward, where does Imagination play in your roadmap for SoCs, or is this mostly going to be Mali now going forward?

Ian Huh: Yes, so we switched the GPU from Imagination to Mali and 5420. We analyze and evaluate each offerings based on the performance and more importantly, power. So basically we just evaluate the available solutions at that time and we found that our ARM stuff is probably better at the time frame and our ARM + ARM GPU ecosystem has worked greatly together, very, very collaboratively with us, so at this moment we are very happy to use the ARM Mali GPU.

Simon Schafer: Just over there, please.

Unidentified Audience Member: This question is for ARM. Regarding the evolution of big.LITTLE, it just seems it's really evolved to a concept of HMP. If that's the case, doesn't even the term big.LITTLE confine you to a concept where you have an equal number of small cores and little cores, small cores and big cores, or rather, as now it seems like you evolved, to a true HMP, that you would have multiple types of cores, maybe not the same number, and the concept of big.LITTLE, at least as you originally thought is, is no longer applicable, that it's really – it really about HMP going forward. Is that – is that a correct observation?

Tom Cronk: Yes I think that's right, I think if you're making a high-end tablet, you'd probably lean towards more big and less little, and if you're making a premium phone you'd probably lean towards more little and more big. And that's the other benefit with HMP: it doesn't have to be symmetrical any more. Four little, two big would be a good solution in some of the phone segments, and it's another dimension that partners can choose to exploit, to differentiate their product.

Simon Schafer: Good, so I think we'll have time for one more question, and then we'll move onto the next panel.

Unidentified Audience Member: I guess I had a follow-up from the same question before on deciding between the GPU of Imagination versus ARM's. Now did you do that (inaudible) on a standalone GPU capability, or was there anything to do with the GPU CPU compatibility. Does that any way feature in your decision of going for a – a Mali processor? And then I have a follow-up for Tom.

Ian Huh: Answer is no, it is purely based on GPU performance, and ecosystem stuff, as you know, the GPU ecosystem is very important: how many game developers are using it, it's more important, actually, so it's based on GPU performance, power, and the perspective from building our ecosystem, and not really with the CPU.

Unidentified Audience Member: And just on one quick question on the big.LITTLE and markets. I mean, I know another question was asked about where all big.LITTLE can be used, and I'm just wondering in the internal things, can you think of any end-market where big.LITTLE solutions would be used, maybe both of you can answer that.

Ian Huh: So I'm not really product planning side, but there are a lot of applications in the future, I'd say. For example, think about wearable devices. It's not going to be like the scale of the Exynos 5420 big.LITTLE, but if wearable devices are doing nothing for 23 hours and then do something for five minutes, they need to have very power-efficient processors, but very quick response, so that type of big.LITTLE concept still holds true for very low-end wearable devices. So I think that perspective is still true and we can see a lot in the future for big.LITTLE.

Tom Cronk: I've certainly talked to partners recently about alternative deployment; wearables is another area.

Simon Schafer: Well, thanks, everyone, for their questions, and Tom, Ian, thank you very much for offering your perspectives.

Simon Schafer: Thanks for doing this. And then we'll get onto the next panel. Thank you.

ARM and Physical IP

Simon Schafer: Thank you very much. Dipesh is the EVP of the PIPD division, and he's been with ARM since 1997.

Some of you already asked on some of the more manufacturing-related questions. Unlike the previous session, where we actually only had one slide, Dipesh is going to take us through a couple, I think that just act as a reminder as to how Physical IP all fits into the equation. It's obviously become a significant part of the business since the Artisan acquisition. Dipesh, thank you very much for joining us today.

Dipesh Patel: Thank you, thank you, Simon. So I've probably got about 10 slides, but I thought I would take you through what Physical IP does, how we see the market.

I'll try and address some of the questions that came up with respect to competitiveness against Intel, whether the ARM ecosystem is really going to fall behind or not, and I hope to show you what we are doing, both inside Physical IP but also working closely with our CPU team and the GPU team to make sure that the ARM ecosystem has the capability to compete with whatever Intel is doing.

Intel's big claim is that they've got end-to-end solutions from manufacturing all the way to software, and I believe through a combination of work we are doing inside ARM with PIPD through the manufacturing community, whether it's TSMC , Samsung, and other players, we have the solutions that enable the best SoCs to be built, that go into all the applications that you are seeing.

So first off, just a quick reminder with what we do in Physical IP. We develop basic building blocks that enable our partners to translate the SoC designs that are expressed at a higher level of abstraction, into chips. We develop standard cell libraries which are basic functions, e.g. AND gates, OR gates, flip-flops, the things that you need to store and process information.

We also provide memory compilers. Memory compilers are basically memory macros that are present in SoCs. Nowadays, a lot of the SoC content tends to be memory, because you need different type of memory functions for frame buffers, for caches, for L2 systems, and we develop compilers that our partners can use for those functions.

We have a particular expertise within the division to develop general purpose IOs, so these are high-quality, high-density IOs that get used for general purpose interfaces that the SoC usually needs.

And then we have a product category for CPU which is PoP IP for the Cortex processors. Now this PoP IP is one of the components which enables us to ensure that the best CPU implementations are possible. So when our partners take the CPU from us, take the PoP IP from us, they can create the best power performance implementation of that CPU.

And then similarly, more recently we launched a product line to do the same thing for Mali cores as well, and over the last 12 months we have created PoP IP for Mali, supporting the T628 and

the T678 set of processors. And just yesterday we launched the next generation processors, the T700 series, and we'll create PoP IP for that as well.

Now going forwards, if you look at what the industry is doing, for a very high-level overview of what's happening in the industry, as the previous panel mentioned, Samsung is looking at developing SoCs that will go into mass production on what they call the 14 nanometer node. So that should happen, from what we heard earlier, at the tail end of next year.

But in the meantime, this year, we've been developing Physical IP platforms that they can use to develop those SoCs that will create engineering samples through all of next year. These engineering samples are an important step along the way to enable the software to be debugged, the software to be developed, systems to be designed.

So that's – that's 14 nanometer, as – as Samsung calls it, and similarly, the TSMC ecosystem calls the same sort of thing, 16 nanometer. And there's a whole set of partners that are targeting that node as well.

In terms of the status of each, it's very similar. from my perspective; for Samsung, as far as IP development is concerned we are creating those platforms. And I expect to be engaging many customers for TSMC to meet very similar time lines to what you heard earlier.

We expect, from our silicon partners perspective, mass production to happen in 2015, very consistent with what you heard from Samsung. I expect the same to happen not only with Samsung, but also other silicon partners who are targeting advanced nodes. And I'll show you later what we have done in Physical IP to reduce the risk for partners who want to manufacture Cortex processors.

From my perspective, there are very few partners that are targeting 20nm, so the early partners who made a decision two years ago to go to 20 are sticking with that decision, but actually, a lot of the partners who are coming into the market now, and are just finishing 28 nanometer are really going to skip 20, and they're going to go straight into 16nm.

At Twenty-eight nanometer, on the other hand, I think we are seeing a resurgence, and there's a lot of different applications that are going to stay at 28 nanometer. And with that in mind, we are expanding our support for 28 nanometer, and in the past 12 months, we've doubled the number of processes that we're supporting at 28 nanometer.

From a mobile tablet market perspective, that still continues to be the driver in terms of design starts, as you've seen both at the high end as well as at the medium end and the low end. A lot of these devices run on Android and a lot of those devices are using 28 nanometer or 14 nanometer technologies at this point in time. And then as you go into 2015, we do expect some of these partners will then move into the advanced ARMv8 architecture.

At the other extreme, everyone's talking about Internet of Things, and what's happening in the IOT space. And a lot of these devices are actually at older nodes. And some of these older nodes, are now having to be refreshed, because the IP that is developed at the older nodes was developed five, ten years ago, and some of the modern-day power management techniques do not exist at those older nodes.

So we've been doing a lot of the refreshing in the Physical IP team to enable IOT applications where power consumption is a key factor. From a summary perspective, I would say at the various nodes, it doesn't matter where you want to manufacture, other than Intel, we'll be there for you. We have IP available across all the four providers that are offering advanced processes. TSMC, Samsung, Globalfoundries and IBM.

I talked about using risk and readiness in terms of manufacturing. So, one of the things we've done in the physical IP team is make sure that our processors, combined with the physical IP combined with the process, is ready for manufacturing.

So, we do a lot of early R&D experiments and here's one example where we worked closely with Samsung on their 14 nanometer process tape-out. Actually, this chip contains two different processors. It contains a Cortex-A7 and a Cortex-M3. And both of these processors were present on the left inside of that plot that you can see.

But we also used these R&D vehicles to perform other experiments that give us confidence that the IP that we are designing will be suitable and ready for that application. And therefore, when the silicon partner comes along and wants to manufacture at that node, they will have a low risk to production.

We also worked with TSMC, as well, and early this year, we taped-out a Cortex-A57. Now, this is a pre-release Cortex-A57. It's a beta quality but the idea here, again, was the same, which was to show that the processor and the process and the physical IP are going to work together. So, when partners go into mass production next year and the year after, the risk for them is reduced.

And if you look at that over time, this is something that we've been doing forever. You know, ever since we started on 28 nanometer, 32 nanometer about five or six years ago, we've been adopting the same type of approach -- do some R&D work, develop platforms, qualify the platforms and then the partners go into mass production over time. And 32, 28 nanometer for us has been in mass production for over two to three years. Now, that's earlier than when Intel was producing 32 nanometer devices for the mobile market.

So, this gives me confidence that we can keep up, if not exceed what Intel is able to do. For 20 nanometer, we have two more customers using our physical IP and entering production this year. For 14 nanometer of FinFETs, we have two other lead customers already engaged with us using our IP I expect to get engineering samples next year leading to mass production the year after.

And then we are already starting to work on 10 nanometer. And this goes back to the question that was asked earlier -- you know, Intel will claim that they're the only company that can keep up. But as you can see here, the ARM ecosystem is already at 10 nanometer which is probably similar to what Intel's doing right now. But we are there right now and doing the R&D.

And then POP -- I'm going to actually not go too much on this, but POP is one of those products that I talked about earlier which offers our processors and multi-processors the best implementations. And we support that technology across different foundries, as you can see on the right hand side. We support TSMC, Globalfoundries and UMC, but also across different process nodes. And at the bottom, you can see the process nodes that we are supporting.

And it's a complex three dimensional matrix which I have tried to fit in this picture, into a 2-D picture. And as you can see from this, we have early access quality available for ARMv8 cores

at 28 nanometer today and we're already working on FinFET IP for those same cores which will come up next year.

And then, at the other extreme, the IOT extreme, there are many different processors that can be used, depending on the application that you're going after. Now, IOT means many things to many people. Wearables is one example, and a lot of the wearable devices have Cortex-M3 type systems. But at the other extreme, you might have smart grid meters. Again, intelligence that exists in the fabric of what you're doing day to day.

And some of these older processes are commonly used for the IOT chips and we have support across the board. We support our 16 foundries at the 90 nanometer to 250 nanometer space, and many, many different variants of processes across those nodes.

So, since the last time we had this event, which was about 12 months ago, there's been significant process on the physical IP side, both in terms of POP IP -- we have many new products available on the top left, supporting 28 nanometer processes, as I said earlier, that the number of processes being supported has doubled and we now have five foundries adopting our physical IP technology for 28 nanometer.

And at the other extreme, on the FinFET side, we are doing multiple test chip tape-outs. I showed you two, but we've done another two or three test chip tape-outs since then to demonstrate that the processors can be very compatible with the IP that our silicon partners can use to build their SOCs.

So that, in summary, gives you an overview of what we do in the physical IP division to make sure that ARM based SOCs can deliver the best performance and compete against Intel.

Simon Schafer: And actually, just before I start, does that answer the question about implementation, versus Intel?

Unidentified Speaker: Yes, thank you.

Simon Schafer: Now that you've actually begun to see some tape-outs for the ramp for next year -- 14 or 14 equivalent FinFET, what are the early results and what are the challenges that you see and what are the type of things that are really causing potential difficulties, if you think about the goal of having mass production by 2015?

Dipesh Patel: So, in terms of challenges -- early challenges that we saw when -- this prior today but one of the biggest challenges is making sure that the EDA tools that need to be used to implement the SOCs are ready and that it works with FinFET technology. One of the new things that's come out over the last two years is something called double patterning, and that includes this additional complexity when you're designing the SOCs.

And so, we need to make sure that the tools that work with our physical IP also work with the RTL from the CPU. And so, the projects that I talked about earlier -- both of those projects, I forgot to mention this, but they had an EDA partner in them. Cadence was that EDA partner and so that's one of the big challenges, is making sure that we can actually build the SOCs at a physical implementation level.

The second challenge comes from the manufacturing side, and this is why we have the R&D vehicles. The 14 nanometer chip that we did -- we actually taped that out 10 times -- the same design; because every time we taped it out, we got some feedback from the foundry saying, "Oh, we need to change these rules to honor the tape-out," and we did that over successive weeks -- almost on a weekly basis, until it got to the point where the foundry was confident, at the end of last year, that they could manufacture it. So, manufacturing is another big challenge.

And actually, now we have both of those pieces of silicon -- they are working. We have processor code running on the chips and that gives me confidence that this functionality is good. Manufacturing is good. The EDA tools and the IP can work together so somebody can build an SOC.

And this removes a lot of risk. So when the engineering samples will be built next year, sort of pre-mass production, I have confidence that they will come back and the yields will be good and it'll be able to go into production.

Simon Schafer: Actually, maybe I could stay on some of this -- some of these challenges in manufacturing. And you mentioned double patterning. You know, and this is not necessarily going to be the big challenge of '14, but of course, one of the things that many people in this room have heard about is, you know, the challenges in the EUV.

So, more broadly, in an environment where frankly, you know, multiple patterning type steps are going to be a reality of life, it sounds like, for a number of nodes perhaps -- for longer than people, you know, most people, certainly ASML, were expecting -- what are the dynamics for your business?

Is that a good thing from an implementation point of view or is it more of a challenge? I'm just trying to gauge as to whether that makes a difference to the way you can charge your customers and how would it impact revenue potential for the PIPD division.

Dipesh Patel: Yes, I mean certainly, at one level, if I ignore the manufacturing complexity of EUV, which is actually very, very hard and I think it's probably harder than doing double or triple patterning, and that's one of the reasons why I think there's a further delay to EUV.

But what EUV does is it enables simpler design rules. So, these are rules that we would use to draw the shapes that the manufacturing needs. So, since we don't have that, it means that we have to do double patterning or triple patterning in the future, which basically introduces an order of magnitude of complexity when we are doing the design, which is a challenge, because time to market timelines are not changing. And so, from our silicon partners' perspective, they have a window in which they need to go from 20 nanometer or 28 to 20 and 20 to 14 -- that window doesn't really change.

And therefore, we are left with the same window to deal with more much complexity. And so, we have to innovate on our side and find a way of dealing with the extra complexity in the same timeframe. At the same time, the complexity is going up, but I'm sure the silicon partners, you know, price pressure is there on everyone. And they're not necessarily prepared to pay exponentially increasing rates -- and nobody is going to do that.

And so, what we do is we are able to drive prices up but at the same time, we have to drive our efficiency up. And so, a huge focus in the division is on improving efficiency -- efficiency for developing the IP

Simon Schafer: And -- and just to put into perspective, one of the things that has surprised me in the way that I look at business, your revenue growth actually hasn't outpaced that of foundry. In theory, you would have -- you'd actually think that the revenue growth would be faster than foundry because of course they're going down the node curve, but actually, certainly in the last 12 months, that wasn't the case. I think growth was in the low 20s, both for yourself and for the foundries.

But just to the point about that process complexity, in a world of multiple patterning, does that actually allow you to outgrow the industry or, you don't -- it doesn't sound like you think it would actually accelerate your ability to have a premium versus what foundry revenue is doing.

Dipesh Patel: Yes, I'm not sure that it enables us to exactly outgrow the foundry revenue as there's always price pressure. There's always -- you know, I get the analogy from myself -- my customers, who say, "Five years ago, when you bought a TV it was $3,000 with so many features. Now, the TV is still $3,000 -- you just get a lot more for the $3,000."

We have the same challenges. And so, we have to deal with it. And yes, we're able to drive some of the prices up, but it's not going to be -- I don't think it'll be able to exceed what the foundry revenues are going to be, in terms of percentage growth.

QUESTION AND ANSWER SESSION

Unidentified Participant: Thank you. Talking about double patterning, what are the implications in terms of design constraint for the companies moving to 16 and 14 FinFET node?

Dipesh Patel: With respect to double patterning?

Unidentified Participant: Yes.

Dipesh Patel: From an SOC design perspective, double patterning is hidden away from them. Because a lot of the complexity around double patterning is taking care of inside what we do in the physical IP layer. So, we'll extract all of that away from the designer. The designer doesn't need to worry about double patterning.

So, if you're implementing -- physically implementing an SOC, it's mostly transparent to you. Because a combination of what we do in the physical IP and what the EDA tools do inside the tools hides all of that away from you. What I mean by that is, for example, I don't know how much detail I should go into, but if you're building an SOC and you're laying down logic in a grid, you need to connect the logic together. That logic -- when you connect it together, has to obey certain rules, which will make sure that you can manufacture it.

Now, double patterning at a simplistic level -- what it does is it says if there's adjacent wires next to each other, each wire is a different color. Think of it that way. So, you can have a blue, a red, a blue, a red, a blue, a red. And all the blues are manufactured together, all the reds are

manufactured together. So, from a double patterning perspective, that's simplistically what it means.

And the EDA tool takes care of that. See, it takes care of that coloring complexity. We take care of the same sort of complexity inside the physical IP. So, when the SOC gets it, they just need to run the tool, they need to verify it against the text that the foundry gives them and it should pass; they don't need to do extra things to be double patterning aware. Because the tools do flow and what we've done has simplified it for them.

Unidentified Participant: And maybe -- if I look at the transistor cost from the start -- it never will go down with a shrink? Where do you think that we'll go back to the 16/14 FinFET? Because there are some people that think, for the first time, it's going to be flat or even -- so, I show you that we were going to have a real production next year from your partner, and bringing production into 15, but what's going to be the cost associated with that?

Dipesh Patel: That's a good question. And, you know, personally, I don't have the wafer data because each of the silicon partners is able to negotiate a different deal with their semiconductor manufacturing partner. So, from a wafer perspective, I couldn't tell you how much they're charging.

I've heard the same thing. But at the same time, we do have partners who are moving down to FinFETs, in spite of the fear that the cost per transistor is going to go up. And the reason they're doing that is because there is an actual benefit that they're getting out of it.

You've probably seen the data from Intel which says you can get 50% power savings and 50% performance improvements, depending on whether you're doing same power or same performance when you're scanning. And some of those benefits are what's really causing you to reevaluate that equation -- because maybe the manufacturing cost is higher, but you think if you get a longer battery life, the total cost might be lower.

And we're seeing that and all this is being done by our partners. Now, they don't share that with us, in terms of cost of a transistor at that node. But they're doing it because they can see that the added power saving or the added performance improvements are worth having that advantage to go into FinFETs.

Unidentified Participant: A couple of two very different questions. Firstly, I guess we've seen the competitive advantage of Intel that was stated at the time of FinFET announcement from them -- maybe diminished some of what most recently with the Broadwell 14 nanometer push-out by another quarter.

And I just wondered, you know, whether you see Intel having any real competitor advantage, in terms of timing, or indeed, in terms of transistor entity, which I guess they probably do have. And if so, what the cost of everything is putting aside -- you know the good work you're doing in things like big.LITTLE and software -- what like for like to process advantage Intel does have over the ARM camp and partners? I don't know if you want the other one first before you answer this one...

Dipesh Patel: Yes, because there's a couple of minor, subtle things in there. So, what I -- the basic question is does Intel have a cost advantage because they -- if I'm paraphrasing, they have control over everything, right?

Unidentified Participant: Yes, cost or time, so --

Dipesh Patel: Cost or time...

Unidentified Participant: Time is money, so --

Dipesh Patel: So, I think you mentioned that their Broadwell announcement recently was delayed by another quarter. And one of the challenges Intel faces at 14 nanometer is this is the first time they are doing double patterning. So, if you go back and look at what they did for 22 nanometer, which is their first generation FinFET process, it was basically all single patterning. Double patterning introduces a whole level of complexity for manufacturing that they haven't had to deal with.

The foundry world on the other hand, had already, you know, pipe-cleaned the double patterning flow at 20 nanometer. And for them, the new thing is FinFET at 16 nanometer. So, one of the things that I do think is happening is Intel is struggling with double patterning. They're learning from it and they are definitely delaying their road map appropriately.

In terms of timeline, I don't believe they will have much of an advantage, compared to what the ARM ecosystem is doing for FinFET. If you look at the Intel line for the first mass production devices, we heard late 2014, early 2015. If you then overlay that to where Intel is and then overlay that with what the system level [benefits or not] are going be, I don't think there's a huge disadvantage, if at all, for us. I think we'll remain competitive and I think we'll be able to compete with them.

Unidentified Participant: That's very helpful. Thank you. I sat in a networking presentation earlier and they talked an awful lot on POPs and PIPD and I just noticed, it wasn't one of the applications that you talked about earlier. You know, maybe devices and tape-outs being the obvious one. Can you talk a bit about the enterprise application area -- And the new developments there? Thank you.

Dipesh Patel: The pop IP we developed is a generic pop IP that is applicable across different market segments. I only focused on mobility, but we do have application users that use our IP in the enterprise and networking segment as well. So, there are networking partners who are building chips for networking applications that use ARM physical IP -- both from a platform perspective, which is the basic standard cell libraries but also for POP. And the POP IP that we do is processor specific and node specific.

One of the things we do in POP is we benchmark different configuration options. So, we'll benchmark a high end option, which is sort of the highest performance option -- might be a quad core cluster with a huge L2. If you look at that, that's a very typical configuration for what the enterprise segment's going to use.

But we've also got the bottom end, which is a dual core end, which is all power optimized, which is something that you might use for medium and/or low mobile. And then if you're doing a smart phone, you might want a balance between the two -- you might want to balance between power and performance. And again, that's a configuration point that we do benchmark and make sure that our partners have a starting point, if that's the segment they're going after.

So, what I'm trying to say is we have coverage for every segment that you can think of. Whether it's at the low end or the high end. Does that answer your question?

Unidentified Participant: It does, thank you. And I guess -- sorry, can I do one more? I guess maybe it didn't come as a huge surprise to some, but the Altera Intel foundry announcement; I just want to know what you make of that and, you know, what your thought process, as an ARM person is.

Dipesh Patel: I think it's a good thing. Intel manufacturing ARM processors -- good. I think our ecosystem enables that. It's one of the beauties of the ecosystem.

It's clearly not a surprise for me because the Altera announcement has been out for a while. You know, I think everyone knows they went with Intel for FinFETs. For whatever reasons, they have decided that that's a better solution for them. I don't really know why but, you know, as far as our processors being on FinFET at Intel, great.

Unidentified Participant: Thanks. And you -- I heard you touch on EDA earlier on.

Dipesh Patel: Yes.

Unidentified Participant: And you partnered with Cadence and obviously with relation to design rule constraints that we're seeing, picking up with double patterning, FinFET and possibly with EUV, as well -- it puts you in a very privileged position to -- it looks to me, at least, that you can make up the partner of choice in EDA for some of the foundries. And it's interesting you partnered with Cadence quite a lot for future road maps.

Is there -- is there any reason for that close relationship with Cadence over some of the other guys, and is there a possibility, perhaps in the future, that there's a -- there's a dollar wallet constraint for how much some of the foundries will spend on design so that there may be a push against what ARM gets in PIPD versus what the EDA guys get?

Dipesh Patel: I'll come back to your second question, because I'm not sure I completely understand it yet. But the first question, in terms of our partnership with the EDA community to enable a low-risk SOC design implementation. So, in -- in the project that we did, we partnered with Cadence for the synthesis and placement side of it. And that's really what we were focusing on over there.

But behind the scenes, if you look at the IP development side where we do need to use EDA tools to develop the underlying IP that goes into it, we use Synopsys tools and that's very important. Because we need to have sure that the Spice models that the foundry is producing, coupled with simulators that synopsis is producing work very well together. Because otherwise, my IP development gets stuck.

And then, similarly, on the DRC design tool side that you also mentioned, Mentor is who we use. So, we use the DRC of those decks. And again, the same thing is applicable with that. The decks are produced in partnership between the Mentor community as well as the foundry -- they have to work together because the design community will suffer if they don't. And we are part of that process.

Unidentified Participant: Maybe I'll try and reword this -- the second part of that question.

Dipesh Patel: Sure.

Simon Schafer: So, I guess the way I'm looking at it is we talked about the cost of transistor going up in future nodes. Are a lot of camps anticipated around lithium etch, et cetera when we would do double space technologies? And it seems to me that there might be a limit -- and maybe I'm -- maybe I'm reaching in the wrong place here, but there might be a limit on how much is spent by foundries on design from a dollar perspective and that maybe EDA and PIPD have to, you know, compromise on how much dollar they can get there.

Dipesh Patel: When you say foundries -- how much foundries will spend on design, what type of design do you mean? Do you mean physical IP designs?

Unidentified Participant: Yes -- I guess I'm talking about design as a -- as a whole one-stop solution, really for -- for the -- for the foundries from the R&D process -- for R&D steps right through to, you know, we got signed off and you can...

Dipesh Patel: I understand. I don't think they will stop spending on design so that our share will go down. I just don't think that'll happen, because from a foundry perspective, if they're going to be successful, they need to invest in the ecosystem. And the ecosystem for them is a number of different partners and players. They have the EDA ecosystem because they need to make sure that tools are going to be ready for the SOC design community, so they have to keep doing that. Otherwise, they shouldn't invest in that process in the future.

Similarly, there's an IP ecosystem. You know, there's a -- while we are experts at standard cells, memories and pops, there are other IP providers who are experts at interface IP. And the foundry has to invest so that a complete solution is available.

Unidentified Participant: So, if you look at the different ARM end markets -- networking, servers in the future, right, printing, all -- even all -- some of the legacy ones, which ones are more attractive for you guys in PPID? And you guys have a bigger penetration than others. Obviously mobile's there.

Dipesh Patel: I think anything there that does high volume is attractive. I mean, we are present in every single market that ARM is also present in. Obviously, mobile is a huge driver for us. Our physical IP is used in some of the top smart phones that you can find in the market and that continues to be a driver for all of us.

But our products also used in feature phones and what I would call premium phones, right? So, these are the sort of next level down. And our physical IP gets used in that, as well, in millions of units. But that's the mobile smart phone side.

But we also have usage in enterprise networking that we talked about earlier. On the server side, which is more a feature thing, our physical IP is being used. The volume is tiny right now. In IOT, there's a lot of different growth that's likely to happen in the future. And we want our physical IP to be used there because the volumes are huge. So, we do want to have wider penetration -- we don't want to be focused on just one segment.

Unidentified Participant: So, in the future, I see server and would -- represents as much of a growth for you guys as it does for the regular process.

Dipesh Patel: Yes.

Unidentified Participant: I still want to explore this Intel versus other foundry question. I'll just ask it a different way. And when I think about what Altera talks about, they say that the TSMC 16 nanometers essentially equivalent to 20 nanometer design rules and so you're really comparing 20 versus Intel's 14. The claim is Intel will get a 2x density advantage versus the TSMC 16 nanometer.

The timeframes will be similar, but they're essentially -- the way I understand it, because you're comparing 20 versus Intel's 14 -- you'll be 2x the number of transistors per square millimeter. Can you help us -- I mean, you know the design rules well. Can you help us understand the disconnect?

Dipesh Patel: So, is the 2x -- when they compare -- because I haven't heard this data point. I mean, I want to make sure I understand it. So, the 2x can be multiple ways. If you go from Intel's 22 to 14, you would get a 50% shrink, for example. That's basically 2x area. Or you could say its 2x when you compare TSMC's 16 to Intel's 14. Which one is it?

Unidentified Participant: Well, I was doing 22 to 14 -- is actually 2.5x, when you do the math. So, just trust me on that. And I'm just saying 20 to 14 -- when I spoke with Xilinx, well they said, well, Intel's 20 -- Intel's 14 really isn't quite a 14. And so that was what they said. But I'm just curious -- when you've done -- when you actually looked at an implementation -- because it's what we're all confused about. We hear all these numbers and people say it's 2X difference; other people say there's no difference. There's a big disparity.

Dipesh Patel: Right. I agree. It's a big disparity. I mean, if somebody was to tell me that Intel's 14 is 2x dense than TSMC or Samsung's FinFET -- let's ignore the numbers for a minute, okay -- 14, 16, 14 -- it doesn't really make a difference. I don't believe that.

Right? The physics doesn't support it, right? I just don't believe that because, you know, I just mentioned earlier that Intel at 22 was single patterning, so their back end of line, which is the metal, was 28 nanometer type technology, okay; what the foundries were doing at 28 nanometer. And their front end was not as dense.

Now, there may be other tricks that are being played that I'm not aware of, but I just don't see that, if I just compare the raw information that I have at my hand, side by side.

Unidentified Participant: So, you don't think there's any difference?

Dipesh Patel: I'm sure there's -- I'm sure there's some difference; 2x seems too much.

Unidentified Audience Member: Just two quick questions. One, just looking at your slides, you talked about Internet of Things and you talked about 90 nanometer to 250. I'm just wondering -- I mean, obviously, the previous panel talked about big.LITTLE moving into wearable technology which I consider at Internet of Things, and I know Internet of Things can be many different things.

But I'm just wondering why you referred to 90 nanometer to 250. Do you -- are you talking more about the Cortex-M variations and not really about the Cortex-A variations? Is that what you're thinking around it? And I have another question.

Dipesh Patel: So, yes, that's a good question. I think if you're talking about fancy wearables like high end wearables that have lots of super functions in them, then that's not going to be a Cortex-M application. That's going to be more likely a Cortex-A application.

But if you're then looking at something like what I wear here -- I've got a Fitbit. I don't know how many of you have a Fitbit, but this thing here is a Cortex-M. And I think a lot of applications are going to be like this, right, which from a wearables perspective, a Cortex A-53 would be probably an overkill.

And I hope Tom's not here -- I don't know if you agree or not, Tom, but for me, it's an overkill. But a Cortex-M is perfectly suited for this, right? And it does the job. It does the algorithm processing that you need to do for encryption -- a lot of these things will use hard IP to do the encryption. And they will have built-in radios.

So, I think you're -- you're end node devices are most likely to be Cortex-M based in a lot of cases -- not every case. Your hub -- you know, we talked about the hub were data comes together like in a smart phone -- that's Cortex-A based.

Unidentified Participant: So, I guess on that -- what I was trying to understand was if I consider where, you know, customer base is and this is supposed to be going to a whole plethora of customers -- do you think the 90 nanometer 250 could be revived as a source of revenue generation for you and physical IP? Or do you think all the IP out there for this is already there with your customers, so there's no extra end market for you go after, in this case?

Dipesh Patel: No, I do actually. So, on this slide, you can see there's many different types of spaces that are usable. On the right hand side -- let's just look at the right hand side for a minute -- nano components, RF components -- in some cases, there may be some high voltage needs. We have refreshed processes across many of these nodes -- mostly around 110 nanometer, actually -- 110, 130, over the last 12 to 18 months. So, I do see a refresh happening at these nodes.

On the left, you see the reference to ULL. It's perfectly designed for some of these applications. We've developed and refreshed platforms very recently at that node and I see that happening over time.

Now, some of these -- like, now I'm only showing 90 to 250 on this slide. I fully expect all this to be at 55 nanometer in the next two years, right? And we'll see foundries like TSMC and UMC and SMIC offering 55 nanometer LP with embedded flash capability at these nodes. And the existing platforms that I have there are good, but I'm going to have to go back and update them because the operating point of these newer nodes is going to be quite different.

You know, we're talking about 10- 150 megahertz operation, because all these were designed for 300 plus megahertz operation. So, they will work at that lower frequency point, but they will consume more power. So, if I know I've got a different target point, I can re-optimize it and create a more optimal solution.

Unidentified Participant: And I just have, very quick, had a question. You know, before this, there was a slide where you're talking about your IP on the IO side, as well. And in the last quarter, we heard ARM getting a fair number of licenses from analog customers.

And I'm just wondering that -- is there any area there for you to explore, in terms of development from the IP perspective, whether it's interaction between the analog to additional within the same chip or what you're looking at at this point of time, or is this just screened for you?

Dipesh Patel: Right now, we are focused on making sure we have a core area of expertise. It's around POPs, it's around standard libraries, it's around memory compilers. And then GPIO -- we're not very specific about that, around GPIO, and we're going to stay on those because we can add a lot more value in those areas than we can do in some of the other markets which are becoming a little bit more commodity standard space IP, for example.

So, we're focused very much on this, actually. We're always looking for other opportunities and the sort of criteria that I look at is where can I add value?

Unidentified Participant: Good. I think we probably have time for one last question.

Simon Schafer: I just had a couple of quick follow-ups. In the slide, you showed the 10 nanometer. What's the timing of that? And then going back to an earlier comment this gentleman in front of me made, it seems like if the foundry can't deliver us 14nm by the end of '14, early '15 -- that would indicate that the lead is shrinking. So, what's the 10 nanometer timing?

Dipesh Patel: So, let me come to 14 first and then I'll explain why I'm leading into 10 nanometer from there. So, 14 nanometer, we started engaging on FinFET -- let's just say FinFET -- let's not call it 14 or 16 -- last year. And the status then was early R&D. And we actually started in, I think, Q1, Q2 of last year with FinFET technology. And we are effectively at the same stage with 10 nanometer.

So, early R&D, one test chip has already been taped out. It doesn't contain an ARM CPU at the moment because it's still very early in the process for that, but the next test chip that we do at 10 nanometer, which will be similar to the one that we did on 14 nanometer with Samsung. So, a Cortex A-7 plus an M-3 type of system. That will -- that happened for us in 14 nanometer at the end of last year.

I expect to do the same sort of thing in the first half of next year. So, from a timeline perspective, if you look at it, 10 nanometer is at about 12 to 15 -- maybe 18 months behind where 14 nanometer is.

Simon Schafer: I think we're out of time. Thanks so much for asking questions, participating. Dipesh, thank you so much for doing it and we'll do it again soon.

Atmel and the Internet of Things

Simon Schafer: Great. Thanks, everyone. So, thanks for coming back. This is the embedded panel. I'm sure Internet of Things as a broader topic will come up. It's great to have Sander Arts from Atmel joining us today as well as Charlene Marini from ARM.

Just as a preface -- and some of you who were just talking about it in the break, you may have seen the Atmel just released their earnings and the one thing that we'd love to do in this panel is perhaps avoid any questions surrounding either the quarter or the outlook statement in particular. That's not the forum that we had in mind for this particular discussion. So, if we could maybe try at least to abstain from entering that part of the discussion, that would be -- that would be great.

With that, I think you two collectively want to kick us off in a joint presentation in a few slides and then we go straight into Q&A. Thanks for doing this.

Charlene Marini: Hi. Yes, so just to set kind of the beginning context here, we're looking at broad-based microcontrollers. And so we put a few numbers up here -- kind of comparison between 2012 and 2011, in terms of growth in the market. And so, roughly, in terms of broad-based microcontrollers going from about 9 billion units to 11 billion units over that timeframe to 2017.

And looking across this, one of the drivers we do see is Internet of Things. And so that's what we're going to talk a bit about today. And really, we see it as three components. Microcontrollers are a significant portion of that – they are the main intelligence we see in many of these devices are going to come to market. But other components are sensing and connectivity, with the microcontroller being the core intelligence around those.

And then also, we can see from the bottom here, we were talking about what technologies from ARM are applicable. And, of course, Cortex M, which is a range of processors from Cortex-M0+, all the way through Cortex M-4, as well as the Cortex-R series and their devices in that category of processors.

So, that sets the context of what we're thinking in terms of the volumes in the market and the specific ARM products that we're talking about. Then I'll hand it off to Sander to bring us more into that Atmel context.

Sander Arts: So, my name is Sander Arts. I joined Atmel a year ago as the Head of Marketing on the management team. I report to Steve Laub. We have a lot to do, but we also have an amazing story in IOT. We are known for AVR, right? However, we were one of the first ARM licensees, so there's a long-lasting standing relationship with ARM.

And if you look at the right hand side of this chart, you can see that we have actually stepped up the investments associated with 32-bit ARM this year, as you can see -- a huge push towards the market. And there's going to be more next year. And you heard the story around the acquisition that we did last year of Ozmo, a low-power Wi-Fi supplier out of Palo Alto.

And we're going to go and connect that low-power Wi-Fi piece through to the microprocessor, which we believe ARM does very well in IOT. By the way, I love interruptions. If you have questions, please do ask.

And so we service all the vertical markets. And I think it's sometimes a little bit hard to put IoT into the models. At the same time, there's a lot going on, right. So this is not some chocolate elephant where people go -- we like elephants, we like chocolate. Let's go build a chocolate elephant.

Gartner says that 50% of the companies that are going to make the Internet of Things are companies that don't exist yet. So there's examples related to this in the Maker space, like MakerBot, that emerged from nothing into an acquisition of approximately $600 million a while ago.

With that particular thing in mind, also from a marketing perspective, there's a lot to do, right. So there's links into the Maker community. People are no doubt familiar with the Maker community, Maker Fair in San Mateo, but also the ones in New York. And there's this huge community of people that are connecting to one another, and that's where the next IOT companies are going to arise.

That's why we changed the marketing strategy for Atmel, because we believe we need to go and engage with that particular community. There a lot of creative people, and a lot of things are happening.

The Fitbit is already here. It talks to your phone. It gathers the data. A lot of stuff in the smart-home automation, remote controls. And I go back to my point that I made about the Wi-Fi acquisition that we did, because the connectivity piece is very important.
So what is needed for IOT, and what our technology areas associated with it that we, as a company offer. And so it needs to be low power, because you don't want to continuously charge your device. So luckily, and fortunately, the Atmel MCU family is extremely low power.

Then at the heart of this story, of course, is the interface. And we have a huge touch business, as you know. Also upcoming is our XSense business, with the flexible display technology that we we're going to roll out.

So from a marketing rules of business point of view, how do you differentiate yourself? So you can go and buy an MCU from a few other people as well. And I think ARM has done an amazing job in -- let's say commoditizing the market.

There's a lot of stuff that can provide a different shade of on top of one particular core. Atmel is investing in building a community; making it incredibly easy for people to do business with us; allowing people build designs on-line. We have huge communities of engineer-to-engineer communication going on, a co-development.

And it links into Makers, right. It links into the philosophy of that particular emerging market. So again, I joined Atmel a year ago. We're giving a huge push on the Internet, because this is where the marketing needs to happen. So you see a lot of investment in social content creation, community development, but also tools and software, and everything that an engineer needs to build product.

QUESTION AND ANSWER SESSION

Simon Schafer: Maybe I can sort of dive into this point about commoditization. If I look at your numbers, and sort of touch on it actually, the market share gains that you have achieved looks as if it's been based on AVR, as opposed to ARM-based solutions.

So maybe just remind us in the context of the importance of the ARM IP, in your portfolio, and where they are going forward, that that's going to be a bigger element for you to capture, increased share.

Sander Arts: It's a good thing and a bad thing that we have AVR. The good thing is we have the installed base. So there's significant portion of revenue in AVR. Then I think, to be honest with you, when I joined the company a year ago, there was a little bit of a paradigm shift happening between AVR, which is at the core of the company. And then people go "okay, are we going to make investments in ARM"?

So I think from my perspective, being a little bit more neutral to it, we managed to go and push that hard. So the good thing is, I think the challenge is that we need to go and change the organization towards being focused on ARM and 32-bit, et cetera, et cetera. The good thing is we have the installed base with AVR.

The other good thing is that, in giving people infrastructure and tools, right, we're doing that as well. So it's compatible. And if you want to go and design in 32-bit, you can just go, still go into Atmel on-line. And you can just work in the same environment.

So we're enabling people to go up, or however you want to describe it, and start developing with ARM. But, yes, and the AVR, those are public numbers. We're still predicting growth in the 8-bit market. Of course, the ARM core business is going to be increasingly important for us. But we do both.

It's only positive, as far as I'm concerned. Because there's other creativity, also emotional connotation and connection in the marketplace around AVR. So there's -- I don't know whether you know -- but there's like 250,000 people engaging with AVR technology in a community side called AVR Freaks.

And if you go in there -- so you may want to go do that -- there's a lot of activity and sharing of information actually happening. And there's a lot of creativity there. Does that answer your question?

Simon Schafer: Yes. Maybe -- let's ask Charlene. What's -- you know, your proposition clearly would be that your architecture is obviously the way forward, as you're using something proprietary. You've done that with a number of customers, of course, other than Microchip and Renesas.

But what do you tell Sander? How do you tell him that, actually, maybe they should think about de-emphasizing AVR?

Charlene Marini: Yes. So I think Atmel's a good example of kind of this transition from 8- and 16-bit to 32-bit. And they're enabling their customers and their developers to choose platforms and migrate between them easily.

And I think that's key as we start to look at increasingly-connected, embedded applications. Because now people are going to be developing across different types of platforms, instead of siloed development, where someone might be, you know, looking specifically at, say, engine-management or motor control.

That device is now going to become connected to something else. There's going to be more of a system-level view across the programmers' mind sights. So that really does speak to one 32-bit intelligence, so that you can easily program even the smallest of nodes, but also, being able to scale as a programmer between these very tiny nodes and something that might require a bit more performance, and be dependent on 32-bit capabilities.

Simon Schafer: And in terms of your pitch, and I think Sander's was actually, was one of the fear of commoditization, I think you used that term. So when you sort of try and tackle that conflict, if you will, when you pitch the architecture to the customer base, how do you do that?

Charlene Marini: Yes, so from our point of view, it's about standardization. And if you look at the innovation that we're going to need around IoT, there's going to be a huge proliferation of hardware, different types of hardware, different types of devices.

In order to really be able to garner the innovation, both with the hardware side and the software side, then programming the processor should not be a barrier. Everyone should have the basic understanding and ability to shift from one platform to another, knowing that the processor that they're going to be developing on is the same, and not have to go through that re-spin cycle of learning a new architecture.

If they don't have to readjust to a different architecture every time they move to a different platform, they can innovate on the higher-level aspects of their design, whether it's at the application level, whether it's the power of their system. They're able to innovate in other areas rather than spending their time and resources on adjusting to a new architecture.

Simon Schafer: And we had Simon Segars do to the keynote this morning; the main tenet to his presentation was in fact the size and potential use of processors in different verticals.

So from your point of view, where do you see the earliest adoption to the type of verticals that you could be serving over time? What are the earlier adopters?

Sander Arts: I was actually still thinking about your previous question. I was thinking of it just from an anecdotal point of view, because I agree with Charlene. Your commoditization comment - it's actually a good thing, right? I stood up in front of the board of directors last week, and I showed a big Amazon.com logo. And I said “in a world where you have commoditization, we need to go and be the easiest company to do business with”.

“We need to have the best technical documentation. We need to have the best tools. We need to have the best technical support.” You need to have the best engagement of the market.

And maybe I used the wrong metaphor by putting the particular logo up. But then the tools people are helping us develop Atmel.com into a site that's the go-to place for engineers, that's what I'm working hard at.

And then your comment about which verticals, I had a few verticals on that slide. I mean, some of that stuff is actually happening now. So the home automation, wearables and stuff like that. I think it's still sort of early, but increasingly is going to be just more connected, and there's going to be more data coming out of all those connected devices we're going to go play with.

I don't know whether you read the article about the heart game. I don't know whether you saw that. There's something to look forward to when you retire, or when you become a little bit older. There's actually smart games being developed, where elderly people can measure their heart rate, but also distance, maybe even GPS type of stuff.

That's where some of those devices are going to go, right. So I guess the sky's the limit when everybody talks about those billions, and billions, and billions of devices, going to be connected. That's probably going to be one of them. And the data, the big data play associated with that..

Simon Schafer: Because you -- I think roughly you've designed, or you've announced something like 150 ARM-based designs by now. But the ARM phase is still very early. So over time -- and I'm not talking 12 months, I'm talking three to five years -- relative to your average growth opportunity, how much share potential do you see?

Sander Arts: You know, this company has been on this journey of changing the portfolio, right. If you look at the Atmel results in the last few years on how Steve Laub and others have changed the portfolio and focus, I think we're extremely well positioned to ride that particular wave.

And especially -- I talked a little bit about this low-power Wi-Fi opportunity -- that connected with the micro, the installed base with AVR, us making the investment associated with giving people an easy migration path, right, into a different architecture, I think we're actually very well positioned.

Simon Schafer: Questions in the audience?

Unidentified Audience Member: It's not a very technical one. You've mentioned the website, and you're bringing engineers to -- you're making your company the easiest to deal with.

I just wondered what the other challenges are of having an increasingly fragmented customer base. That seems to be a sort of complete selection of niches going forward, some of which could be quite small. I just wondered what the challenges to you as a company are in terms of dealing with that fragmentation. Does it mean that you have to hire more engineers to help support these customers? What are the kind of other challenges?

Sander Arts: So I think it's -- it's a very complicated thing to look at, right. So I think an increased, and an enhanced relationship with distribution partners is important, because they give you more feet on the street, and more access to customers.

I think technology's going to help us tremendously, right. So you mentioned website in your question. So we're building a go-to place for people that we won't be able to talk to through traditional sales, right. So how do you go and target people in the long-tail? So we're doing that.

Then I'm investing heavily in a content play. So how do I engage people, and attract people to Atmel.com, when they have all sorts of alternatives to go to?

One of the biggest reasons that I joined Atmel is Arduino. Arduino is an open source, tiny computer, and it has an 8-bit micro in there. And it's used by hundreds of thousands of people. Hundreds of thousands of people use it to do prototyping. So there are a lot of people in IoT just doing stuff in their garage.

Some people think that those are hobbyists, in most cases. In some cases, that's the case. So they build a Halloween costume that gives a light show, or something like that, based on Arduino.

That's not necessarily going to be the next killer app in IOT. However, once people start to go and scale, they still use an Atmel MCU. So if you go back to that quote which says 50% of the 50% of the companies are going to make up the IoT industry are new – relating to your question about fragmentation how do you deal with that?

We are at the heart of the Maker community, we have that particular AVR chip on the Arduino. So when people go scale, they stick with us. We help them move up to a higher-performance platform.

So you asked me - what are the challenges? Reach, penetration, feet on the street. But there is technology and communities and smart ways of getting your arms around those individual people everywhere; we need to be local to the locals, especially on building big content plays in China, things like that. Because you have to be as close to people as you can be. It's all about people-to-people marketing.

So I guess other companies -- but you have to ask them -- might have similar challenges. I think we have a few cards in our hands that some other people don't have. Does that answer your question?

Unidentified Audience Member: As Simon mentioned this morning, the term Internet of Things appeared, or came into existence in 1999. But, you know, for more than 10 years, it didn't take off quite significantly. What makes you think that IOT will take off in the near future? And also, why hasn't it taken off in the past 10 years? What stopped it from being a big thing in the past? Thanks.

Charlene Marini: I guess I'll address that first, if that's okay. I think at the end of it, IOT is about people, and the ability of technology interact with people with a very low barrier.

So, you know, a good example is Bluetooth Low Energy, or Bluetooth Smart. We've had Bluetooth for a while. But you would have to pair your device. You'd have to worry about which device it is, connect it, then download all this stuff.

Bluetooth, Smart Bluetooth, Low Energy have changed this. Take my Fitbit, I just walk near my device, and the information's downloaded. I still have to charge the thing, which is a bother for me. But I think that's a good example of a use case where, just in the recent past, we've overcome a hurdle to wide adoption.

And so I think it is around these automated features that technology is starting to be able to present itself to the end user. And that's a consumer case. And similar things are going on in industrial applications as well.

And, of course, the cost barrier have come down. If you look at sensor technology, it's been driven by mobile in the past year. And that has significantly reduced the cost, and broadened the availability of sensor technology. So I think we're getting to a point where it makes sense from the user's point of view and economic sense from the producer/OEM/system-integrator's point of view.

Unidentified Audience Member: Sorry. This is a follow-up. How do you see the opportunity shared between 32-bit and 8-bit? Is that the usage model for the Internet of Things, or is it much more on the 32-bit side, initially? I'm just trying to think about the phases of the adoption cycle.

Charlene Marini: Yes. So my point of view is that people will choose 32-bit, and I do think Internet of Things is a driver for 32-bit. Even in the consumer market, we're going to get two phases where people are going to take devices, and they're going to have them for a long time.

You do have examples where they are disposable; some medical applications for example. But in some of the more prevalent-use cases you're seeing today, like wearables, I don't think people are going to buy a new smart watch every year.

You're going to have sensors that are on bridges for multiple years. They're not going to be replaced every 6 months. And so you need to be able to upgrade the features and algorithms that are sitting out in systems. So I think that upgradeability is a big driver for 32-bit.

So if we look at sensor fusion, right, it's not just one sensor in many cases. It's going to be out in the field. It's going to be multiple sensors. And there's a heavy mathematical algorithm component in sensor fusion which makes that information usable to the end application. So I think it's about intelligence. I think it's also about designing devices that can last into the future.

Sander Arts: There's two answers. The first one is, it's great to be in semiconductors, because you're always ahead. So I marketed NFC when I was at NXP Semiconductors. We did the first pilots around near-field communication in 1999, and it took a while for it to take off.

For IOT, I think the big hits will be in areas where there are huge savings - energy savings, water consumption, for example, - getting savings out of having sensors, measurement and data all talking to one another.

Think about irrigation. It would be nice to be notified that the soil is dry before your plant dies. It makes perfect sense. I guess at a certain point in time there's going to be adoption in big, markets, where there are incredible use-cases to do all of that.

The fish feeder that talks to your mobile phone is really cool, but that's never going to go move the needle. There will be areas, I believe, where there will be big wins, where the business case is so clear that you can't avoid it anymore.

You said, "why hasn't it taken off"? Well, you know, depends on how you look at it. If you're a half-full person, you can claim that it's happening today, right. So the examples that I had on my slide actually is more futuristic stuff.

And I guess also the fabric is being built, but it still has some voice, right. So once the fabric gets a little bit more closely woven, and there's data, and there's more [stem] that's arising, and then you get the segments or the areas where the business case is obvious, I think it's going to go and explode. And sensor fields, in general, have already gone off. Somebody's buying them, right. Does that answer your question?

Unidentified Audience Member: I'd just like to understand how you see you're going to differentiate your processors compared to TI, for example. Because on the face of it, it doesn't seem to me that the manufacturing is going to be a differentiating factor. No graphics, no customized ARM architecture. So I'd just like to understand how you differentiate yourself versus competitors?

Sander Arts: So that's been a journey. When I started marketing our micros at NXP, the ARM thing was a different shade, right. The NXP model needed an ARM core, and it was sort of, you could sell over the channel, so have an ARM thing. That's sort of gone away in the last few years.

I think I touched upon it, but maybe not clearly enough, how we think we can deploy different shades. And it may sound very simple, but being a company that's very easy to do business with, and where people can go and make their designs, talk to peers, talk to like-minded people: that's, in my opinion, very clearly a different shade.

Fortunately, unfortunately, semiconductors is not the most advanced when it comes to websites. You probably do a lot of buying yourself online. But if you go from a typical semiconductor site to a normal consumer type of site, Amazon.com, or wherever you shop, that experience is completely different, right.

The semiconductor site is still a collection of thousand-page documents. And there is huge upside there, that most people haven't really invested in. I guess at a growth rate of, whatever it was a few years ago, there was no real need to go and do marketing.

So the reason that I'm attracted to semiconductors in marketing, and they, in that particular combination, is the fact that you can go and build a customer experience that already exists, but doesn’t necessarily exist in the semiconductor space. And so it may sound a little simple, but that's what I think.

There's another element. I recently did a few focus groups with 32-bit users. And if you ask them why it is that they choose all of the vendors that you just mentioned, they go -- well, it has an ARM core, right.

Nobody, in my opinion, nobody has yet really taken the leadership position in “we have a 32-bit ARM core MCU, and you should come to me”.

So everybody has an offering, but the real differentiator, nobody has found yet, in my opinion. You may differ. And I would love to hear from you on that. So that particular spot, I think we can go take. And there's probably a few other people that think that as well. But again, I think we're pretty well positioned to be able to do that, with all the ingredients that I have already mentioned. Does that help?

Unidentified Participant: I'm sort of surprised the no one in the audience has actually asked about the three other proprietary architectures. Well, the two that we all know about, Microchip and Renesas, of course, haven't -- or at least not at scale, gone down the ARM route. But Quark, of course, is sort of the new elephant in the room, to use another elephant analogy. But talk about Quark a little bit, and how that changes the dynamics. And, you know, they're talking about product next year. So how does it -- how does it change life for you?

Charlene Marini: I guess I'd start out by saying, you know, Quark is a completely different solution to what we're talking about here. It's not at the level of the Cortex-M end where a device can be made that's less than 2 millimeter squared. And we're talking about power in milliWatts, not Watts. And so I think that's an important distinction.

And when you look at ARM, we're really talking about enabling intelligence out to the tiniest of nodes. And that's what's enabled by solutions like Atmel's MCUs.

I think the other thing on the previous slide, there's an example of how many ARM-based devices Atmel has launched in the past few years. And Sander knows this number better than I do, but it's, you know, upwards of 60 to 90, depending on the year. And it's very much that one size doesn't fit all, right.

When you're talking about these extremely low-cost and low-power points, you do not want anything extraneous in your system. The way your IO bus is designed, and how it interacts with the sensor chip can determine whether that end device can live for 6 months on a battery, or can live, you know, 2 to 3 weeks on a battery. So every component in your system matters.

And I think that's a fundamental difference about what ARM is enabling with our partners, than kind of other architectures, and specifically, that Quark example.

Simon Schafer: So you don't see them at all? Which part of the segment will you see them?

Charlene Marini: So where Quark overlaps with kind of the ARM roadmap and where our partnerships is around the Cortex A7 and Cortex A5. And so you might see that in things like gateway devices with an ARM. So you might have seen that Freescale announced a dual-cortex A7 platform 2 weeks ago that is targeted at sensor-aggregator gateways. So that's kind of a similar, if you're looking, processor to processor.

Unidentified Audience Member: Actually, just following up on Simon's question regarding the Quark family, isn't your competition over there with the Quark family more on the big.LITTLE and 64-bit, essentially the more application processor coming down into wearable technology, rather than micro-controller going up and doing more functionality? Isn't that the way to look at the competition against Intel?

Charlene Marini: Yes. That's another way to look at it. It’s an application processor, like Cortex-A7 is. But again, in this space you have to look at the size, and cost, and power requirements, independent of whether it's a wearable or a sensor that's out in a bridge.

Unidentified Audience Member: So maybe just a follow-up on that, when you look at the wearable technology as not wearable technology, but overall Internet of Things -- do you see a drive for multiple chips, or will it be mostly driven by a single-chip SOC, with multiple applications within that?

Charlene Marini: It's going to vary. If you just take a end-node, it could be a very intelligent node that will use a Cortex-A, because it's in a robot. And that particular node needs to do a lot of processing. Or it could be that bridge sensor example, where all it's doing is sensing vibration, and so you need only the tiniest little microcontroller,Cortex-M0+ to do that.

And in that case, it might be monolithic. You might see the integration of the sensor, and the connectivity, and the microcontroller, all to a very small, small piece, whether it's, you know, multi-chip modules, or die-stacking.

Whereas, in the robot example, because the volumes are lower, you might find that the sensors are separate. There might be multiple sensors that are in a multi-chip package with die-stacking. But then the sensor fusion is separate from that, and the application processor is separate, because the volumes for that particular device just don't justify having a fully-integrated solution.

And there's also a barrier when we're talking about sensors and MEMs. You know, that's a different process than the digital. It's analog, right, you're looking at IDMs that have their own fabs. It's a more costly process around manufacturing. So that's going to be a separate die no matter what.

Unidentified Audience Member: Thank you. This is for ARM. When I think about technology disruptions, usually they come from beneath. And clearly, the Cortex-M0+ is a great effort to push downwards. But when I look at a recent startup from China that's recently got some headlines, it looks like their roadmap maybe overlaps with the latest IP, overlaps with M0+ at the high-end.

Does this -- are you concerned that in terms of the IP, that you're designing specifically around the requirement's IOT, that you may need a new set of IP blocks that really address some of these low-end, emerging applications that take advantage of leading-edge process, or somewhat more leading-edge process that are targeted at some of these performance, yet very power-oriented -- low-power-oriented applications?

And also, just along with that, is there something in wireless? Because there's so much fragmentation in the wireless technologies, and that being a key enabler for IOT, or something more you need to do in wireless that you think you may be able to do to jumpstart that? Thank you.

Charlene Marini: So the comment on kind of the low-end, I mean, we're always looking at the roadmap, and the requirements of end-markets, and how do we build our roadmap? And of course, we're looking at IoT.

And as I think you've seen other people talk about the trends that are going to emerge more than they have today around cost, and power, and security, all those things are going to continue to drive IOT.

And the other point is around sensors, and the things needed to enable better integration in sensors. So all of those are things that people consider and are looking at. But in terms of the processors themselves, the one thing that really leads back to ARM, is around the programmability.

So in most cases with IoT, these are going to be open systems. People are going to want to leverage that standardization we talked about earlier to be able to build their systems in a very quick and efficient manner.

And going back, and following what Sander said, I think there's a fundamental shift that people need to realize that's happening with IoT in the types of developers, right. In, you know, even 5 years ago, but certainly 10 years ago, there was a classic embedded developer. And that was assembly-based languages, moving to C-based languages.

You know, folks working within an OEM space, they had their defined project. They would go to an Atmel website. They would search through, saying: "I need this type of memory, this much memory, these types of IOs".

And they would probably look at competitors' websites as well, and they'd find their microcontroller. Today, and increasingly with IoT, you have to think about those application developers who create that huge ecosystem around mobile. They're now going to be turning their attention to IoT.

And these devices all around us are going to be extensions of that mobile platform. And they come from an application viewpoint of: "I want to create this type of application in this system that's going to be connected to the Internet, and this is what it's going to do".

And I think that goes back to Sander's strategy at Atmel, is how do we best enable those people? And that does involve things like mbed, like Raspberry Pi, that enable people to get their hands on the hardware technology quite easily, and supported with great software infrastructure that allows them to look at that higher level of software programming.

So I think that all speaks to the strength that ARM's going to have, even down to these very low nodes; especially at these very low nodes, I would say.

On the wireless point of view, I think we're getting there. And I think you'll find that there'll be holdouts in places like manufacturing, and industrial, where they've had some other standards for a while. But increasingly, as the prices of other technologies go down, as we are ensuring kind of the security they need in their end systems, I think you find that they'll move to the more standardized wireless protocols that are being used in the consumer space.

And right now, what's happening is, you know, gateway providers are just ensuring that they have plug-and-play capability in the types of communication frameworks that they can support. So I think that's certainly a challenge, but I think it's one that's evolving with solutions. And I don't think it'll be too big of a challenge as we go forward.

Sander Arts: So I just wanted to add to the previous point. Just a few data points, right. 50% of the world's population grew up with Google, right. And the average age of an engineer in the United States is like 55. In China, it's 26, 28. So there's something happening in the world where we need to go, and from a marketing perspective, those are business points of view.

And how do you do all of the questions, I guess, that we've had so far, how do you go and engage with an audience that's so diverse? And how do you go and embrace that audience in an industry that is relatively -- I wouldn't say old school -- but sort of old school when it comes to going to market.

So I guess the excitement that I want to radiate from my side about IoT, and who you're targeting, and how, and when, and with what kind of tools, and how you're going enable it to happen is actually a very intellectual challenge that I like to be part of.

Simon Schafer: The wise wisdom of that. So thank you very much for joining us. Charlene, Sander, thanks for doing this. Thank you.

AMD and Servers

Simon Schafer: Okay. Last, but not least on the panels, just to finish off the panel presentations on the Server side. Really great to have Lakshmi Mandyam from ARM and Andrew Feldman from AMD. Many of you guys will know Andrew from his prior life at SeaMicro, but then of course, more recently at AMD.

But what we're going to do is we'll have a brief introduction, and then I think, Andrew, you will take us through a few slides as well. And then we'll go straight into Q&A. Thank you.

Lakshmi Mandyam: So in terms of bringing server platforms out to the market we've had a great amount of a momentum in the last year or so. And I think if you attended the keynote yesterday, you would have seen HP talk about the benefits that they see of adopting the ARM architecture, some of the systems that they have coming, and of course, the workloads.

We view ARM as having a broad opportunity with servers starting next year when we have partners shipping. And Andrew will talk a little bit more about AMD's plans. But we have partners like AppliedMicro, Marvell, TI, who are already sampling. And then we're excited about AMD and the plans that you'll hear about.

You're starting to see a real choice emerge in the market, and I think that that's going to be a game-changer. And it's interesting that when I talk to end users, it's not only about the low-power aspects of ARM servers, but now they're getting excited about workload optimization and choice.

And 2014 will be the year that we see more and more conversations about ARM servers in deployments around the world. So with that, I'm going to hand over to Andrew.

Andrew Feldman: I'm going to explain the changes in the server market by explaining the changes on the client side, and how those drive a set of changes. And I think the beginning of this is to ask each of you to hold up your cell phone. Just hold up your cell phone for a second here.

And what I'd like you to do is turn off the radio, right. And this is an interesting exercise, because I just turned that device that makes phone calls into an Angry Birds player. It's basically a useless device without that radio.

What I did -- and it's probably true for your laptops and your pads, too -- what I just did is take away the data center from you. I took away your access to the world's compute. And I want you to think about that as I walk through this deck real.

So here are two photos, eight years apart, each taken at the Papal inauguration. What you see here on the left side is the way we experienced the world in 2005 – you see one small phone, here. And what you see here in 2013 [*a sea of tablets and smartphones*] is that we are experiencing the world in a fundamentally different way.

And the way we experience that world has to do with the data center. Every single one of those images was stored, was uploaded, was tweeted, was sent to a brother, or a father, or a cousin. Every one of them was backed up in the cloud.

And this photo set for me, sort of galvanized in my thinking, what has changed, and what is driving the change in the data center?

What's interesting about these client-side devices I asked you to hold up -- and I see many of you are using them -- is that they don't have enough horsepower to do real work. What they are, are unbelievably good display devices. The work you want done is actually done someplace else.

You want to go to Google on your phone, you want to go to Maps, you want to go to Salesforce.com on your phone. Where is that work being done? The work, the computational work, is being done in the data center. And your demand for client-side devices that do cool things is driving a change in the data center.

And this change is monstrous. This change has affected every single aspect of the data center. My question for you as we look forward is: when we look out eight years from today, what will the client side be doing, and how will that impact the data center? How will it impact demand for compute? And this is what we're thinking about every day at AMD.

Here are some things we know, all right. We know the world will contain more devices. It will contain more users, more apps. What we know, for example, is only a third of the world's population has Internet access today. What we know, for example, is there are six billion cell phone licenses out there, and only one billion of them are smartphones. So we know that we're going to see more devices, more users, and more apps.

And what I ask you to think about is whether we can envision a world, with less compute, with less data in the cloud, with smaller and less parallelised workloads, with graphics being less important, and with power and space being less important?

We think not. And I encourage you to ask yourselves these questions, and disagree with me if you'd like. Come up with your own vision for what the future will hold in terms of demand for compute in the data center. So what do we think this means for compute?

In my view: one-size-fits-all computing is dead; that ARM and x86 will coexist in the server space; that by 2019 ARM will command about a quarter of the server market by revenue; that custom ARM CPUs will be the norm for mega-data centers. And by custom, what I mean is that embedded in that CPU will be IP developed, and designed, and owned by that mega-data center.

That advantage is that individual mega-data center owner's software. I believe that smaller and more efficient x86 CPUs will dominate. I believe that every major server player will have a fabric-based offering. I believe that AMD will be the leader in ARM CPUs, and will be the leader in ARM servers as well.

Why do I believe this? Why do I believe that ARM will move from basically nothing to 25% by revenue share in a period of five, five-and-a-half years? I think we have an unbeaten track record. And I know you guys – bottom small-print of all your financial marketing material talks about the past not predicting the future. But the truth is, it mostly does. It mostly does.

And what we know is that in the last sixty years of compute – the entire history of compute – low-cost, high-volume has always won. It’s won in every single market. Ironically, it’s how x86 came to dominate. It was smaller. It was lower-cost. It had higher volume.

In 2012, more than 8 billion ARM parts were shipped. ARM amortized their R&D costs over hundreds of partners. That's a force, I think, which is very, very difficult to beat.

So what are we doing at AMD? What are my teams doing? We announced earlier a 64-bit ARM server part, and we are the only people who've ever built a server part to have one. It's code-named Seattle.

It's based on the Cortex-A57. It's running at greater than 2 GHz, has extensive offload engines, all the server characteristics you need. It includes the fabric that my former team designed at SeaMicro. And we're sampling in Q1.

I think now we'll turn it over to some questions. And if there are no questions at all, we'll just all go to the bar. This is an English company, right?

QUESTION AND ANSWER SESSION

Simon Schafer: So Andrew, you laid out a vision of why AMD can dominate the ARM-based CPU landscape in servers. But there's obviously a number of companies who've been relatively aggressive, a number of them startups. And that's in an environment where volumes are still very low.

So ultimately, once volumes start to grow, what gives AMD the real edge in terms of capturing that upside, of the 25% of the market that you think ARM-based solutions can capture?

Andrew Feldman: Sure. First, I think that historically, the CPU market has been a very bad place to be a startup. It's been an unfortunate place for startups. It has a set of characteristics that are bad for the venture community. It requires a great deal of capital. You have to continue to invest. You have an 18-month pattern of having to deliver in silicon.

And as a result, one of the real weaknesses in the x86 history is that for the last ten or fifteen years, we haven't had a meaningful x86 startup. And that has left the ecosystem barren. It's left two players only, and that's a problem.

I think one of the great advantages of ARM is that it scales down the process of building a CPU. It makes it a tractable problem. In the end, we're going to talk about the advantages of power. We're going to talk about the advantages of space.

But at the end of the day, the fundamental advantage that ARM brings is this: if I build an x86 CPU today, it takes me 4 years and $400 million, whereas we can do an ARM server CPU in 18 months for $40 million. And that's a tough force to beat.

And why us? We have 10 years, or 12 years of experience, mistakes, things we've done well, learning that we bring to bear on every part we build. We have core IP blocks that are proven and battle-tested, and deployed millions of times: RAS features, server security, a memory controller that runs 128 gig of DRAM, et cetera. We have vast expertise in building client-side parts.

Simon Schafer: Right. Because the way you speak, I mean, the 25% figure is sort of a small amount in the context to how you can convey it. So (inaudible) be interesting, your view as to

why 25%? Why not more? Why not less? What are the challenges of adoption, and replacing that part of the ecosystem that, Andrew, you mentioned?

Lakshmi Mandyam: Andrew, in the past, has affectionately called me a sandbagger; have you not, Andrew? At the ARM analyst day in May, I talked about the opportunity for ARM servers. We looked at 2017 and said, the addressable workloads for ARM make up around 30% of the market.

And we said we would take 5%-10% of that 30%. That was based on the assumption that it takes a little bit of time to get the 64-bit ecosystem vibrant. And, at the time, Andrew hadn't announced that AMD were going to do server parts.

If we look at how ARM is progressing its roadmap, it's not just Cortex-A57 – there are architectural partners that you've all heard about. I think a couple of weeks ago, Broadcom talked about, you know, going after a 3-gigahertz part. Andrew's talking about scaling eight cores and beyond that in the future.

So I think there's going to be a wider selection of parts that are going to address a broader set of application points. And so I think by 2020, it is conceivable that we could reach a broader set of workloads. But for now, I think -- I mean, I don't think it's out of the realm of possibility. Did I answer your question?

Simon Schafer: Yes. Andrew, maybe I can follow up. I hear you about the opportunity set for ARM-based CPU. But then when I look at AMD's announcement with Verizon, which is a great win for the company, it actually turns out that you're pushing x86 to a large extent in that particular product. And the ARM-content is still relatively low in that particular launch.

Andrew Feldman: For sure. I mean, that's a big win, no doubt. And that's one of the largest public clouds. But we didn't have a part. And I think the market begins in 2014. That's when it begins. And in 2014, it's a year where many will kick the tires. So that means hundreds of thousands of parts sold, if not more.

But what you will have is you'll have next year, in my view, at least one major mega-data center, announce a full port to ARM. You'll see a software company that dynamically does translation from software written for x86 instruction sets – you just use one of their VMs and it does a realtime translation, obliterating any challenges of running x86 software on an ARM CPU.

I think starting in 2015, you'll see the first deployment at a mega-data center, and you'll begin to see a set of rollouts there, particularly in storage. You already saw Baidu use the Marvell 32-bit part in storage. You'll continue to see that happen in storage.

Anybody who's got an infrastructure where the software base is deployed over many, many servers, anybody who's got a scale-out infrastructure today, is interested in ARM. We never go to a call and bring up ARM, and the client says: "sorry, not interested".

It just doesn't happen. The question is: "when can I get boards? When can I try?" And I've never seen anything like it.

Unidentified Audience Member: According to some of the white papers out there on storage, right, like the digital universe, 70% of all content is supposed to be pictures, video, going

forward through X amount of years. And to me, that would be something where low cost, low power efficiency, would make a lot more sense, versus having an x86. So why only 25%? That's my question.

Andrew Feldman: Only 25%?

Unidentified Audience Member: Yes. Why couldn't --

Andrew Feldman: From zero?

Unidentified Audience Member: -- that entire 70% eventually, maybe not in '19, but see if we completely replaced --

Andrew Feldman: I think -- I think it's a mistake to think that ARM is limited to a subset of the market. I think ARM will eventually have offerings in every segment of the server market; from the small core, for the Web scale, to the largest core for super-compute. They will have a play everywhere.

I think you're right. I think cold storage is a phenomenal -- you know, most things we store are what I call 'one-three, never'. I mean, when's the last time you looked at your storage? Your iCloud? And it's written once, and with any luck, you never look at it.

Cold storage is a crazy big business right now, because we don't know what to store. And with our data analytics, we can now do things like a Cassandra analysis, and we don't even know which data is pertinent. So we store it all. And there's no question that that type of, we call it storage server, is a great fit for ARM in the beginning, without question.

That's why Baidu did it. Because what's really happening is that you're using the CPU like a storage controller, over a ton of different disks. And instead of paying that app five times what you pay for a regular disk, you just buy the disks. And, you know, then you've got a cold storage device.

Unidentified Audience Member: I just had a follow-up question on the x86 server versus the ARM-based one. What kind of pushback do you get today if you were to sell it? I understand that everybody's interested in when. But is there no pushback whatsoever if you try to sell them an ARM-based server versus an x86? And has that changed in the last year so to speak?

Andrew Feldman: I think the pushback right now is: "show me the boards. I want to run software. I want to run tests. I want to understand."

We are interested. And we're sampling in Q1. So I think that's the case for just about all the guys who are either sampling now, or planning to sample. We had a given customer come to us. I said, you know, we're sampling if you want. They said: "I want 1,000 boards".

Unidentified Audience Member: And are they worried about security?

Andrew Feldman: No worry on security? No.

Unidentified Audience Member: And very quick follow-up on that also, big.LITTLE, of course, is something which you are considering more and more in this space. I'm just wondering, can you

give us an example of a specific calculation where big.LITTLE would be more useful than just having a full-blown 8-core chip?

Andrew Feldman: Our 8-core is big-big.

Unidentified Audience Member: That's what I'm wondering. Would big.LITTLE be a viable thing in server space at all for you guys? For you guys as customer?

Andrew Feldman: I think we've looked at, as an industry in servers, we've looked at asymmetric size cores for a long time. And there are particular workloads where it makes sense. I think for the first selection between 2014 and 2016, I think you're going to see all big as the opening volley. I might be wrong there, but from us, you're going to see all big.

And I think over time, as individual segments open up as we have more experience, and we see particular opportunities to take down the core size to save on the power, to therefore, pack more in a rack, you're going to see people look to make a tradeoff to say -- Look, I don't need all those cores as big. What I really need is some of them as big, and we'll take some small.

Lakshmi Mandyam: What we're seeing in SDN and NFE and some of these networking applications, which are in some cases converging closer and closer to the server space, is that they are looking at big.LITTLE, but not in the sense that you would in the mobile device. What they're saying is: "I'm going to use a bunch of big -- let's say – Cortex-A57 cores to do control-plane, or higher performance requirements, on a per-thread basis. And then I'm going to use a sea of little cores to do some of the data-plane applications.

Having these on a single SOC enables you to do a lot of creative things and partition the problems in a more creative way.

Andrew Feldman: That's actually a really good point. If I were building a packet processor, for example, I'd go big.LITTLE. And I'd go asymmetrically big.LITTLE as well, because I'd want my control plane and my data path on different-size engines.

And I think you will see some of those merge as you get more detail on what Broadcom's doing, as you get more detail on what Cavium's announced, as you get more detail on the guys who are coming up out of packet-control manipulation and editing, the guys who are trying to build something for network function virtualization, some of the other advanced networking features. That would be a good place for it.

Lakshmi Mandyam: You see it today with the AppliedMicro Solution, where they have quad-core Cortex-A5, I think it is, doing some of the SDN functions. And I think the Calxeda solution also has a couple of smaller cores on it as well. Maybe you should talk to them about that.

Unidentified Audience Member: On the 25%: Andrew, we know there's many people participating in many different applications. And your statement of being a leader, you're carving out a big portion of that 25% for yourself. Could you tell me what application do you see yourself taking there? What workload is it that would drive the majority of that 25% for you?

Andrew Feldman: In compute, in the data center, we know the adoption pattern. We know that what the mega-data center guys do rolls out to the financial service companies 3 or 4 years later.

And the financial service players have been a leader in the enterprise. And what they do rolls back into mid-America, and the rest of the enterprise in 2 or 3 years. So we know exactly how this works.

This is how Hadoop unfolded, right. It was at Google, and at Yahoo, and then it rolled back some 3 years later, and has now rolled down across the enterprise. This is now a well-worn path. It's the same for Memcache. It's the same for NoSQL Database.

This is a well-worn pattern. So I think what you're going to see is adoption and wins in mega-data center, in Web, and top-tier businesses. You're going to see some deployments in this time frame at large financial-services companies.

But I think by 2019, I don't think you're going to get 'RR Donnelley in Cleveland', right. I mean, I don't think they're going to be the target. I think people will adopt first when compute manufacturers compete for their business.

That is almost always how you tell who's going to adopt technology first: who, when they use this technology, will see positive changes to their bottom line. You know, using an Exchange server doesn't change anybody's bottom line. But if you're in the business that Facebook's in – the business of giving away compute for free in return for eyeballs – your primary cost is compute.

And everyday you need to be looking and thinking about how to drive down the cost of that compute. And the same for the 30 other guys of scale in that same category. And so I think you look at us as winning in the mega-data centers, it's rolling back into the financials in that time frame.

Unidentified Audience Member: I sat in a networking presentation earlier. And CoreLink struck me as something very similar to what some of the new OEMs, or new processor vendors, like Calxeda, are doing. And I'm wondering if they're using CoreLink, or if they're doing something proprietary?

Lakshmi Mandyam: So if you go to the show floor, there is an enterprise pavilion with six or seven partners there. There's a 16-core Cortex-A15 from LSI on the show floor. And that uses the CoreLink 504.

And we announced the Corelink 508 at the Lindley Processor Conference, which allows for 32 cores. So people are taking advantage of the ARM SOC fabric as well.

Unidentified Audience Member: That's fine. My follow-up question then is a lot of the differentiation for, let's call it a data center SOC, is in the timing, the IO. It's in the scheduler. Have you gotten any pushback saying, hey, you're doing a lot of the work that we're supposed to do?

Lakshmi Mandyam: To add to what Andrew said earlier, I think the line that defines, to quote Facebook, "gratuitous differentiation versus innovation" keeps shifting. As an IP company, we try to stay ahead of where that line is.

And so three to four years ago, it would have been blasphemous to think about integrating an SOC fabric. But now if you look at where the innovation is shifting, it's more: "I'm looking at an

SOC where I need to add more value to data plane acceleration capability", taking LSI as an example.

Or in the case of Andrew's company, it's Freedom Fabric. The line of gratuitous differentiation versus innovation keeps shifting. As an IP company, we try to stay ahead of that line. So you could look at it as a trend-predictor in terms of where that line is moving.

Unidentified Audience Member: I think Andrew gave us his vision of where we are in 2019, that this is going to be a game played by the larger companies. In most early-stage industries, you generally see a lot of new players come in, some get to decent profitability. They become the stalking horses for the bigger guys to get in, introduce benefits of scale, and ultimately take the market.

Is there something specific about the dynamic of this market that means that's not going to happen this time? Maybe it's just the fact the big data centers are so big, and need such long roadmap, or whatever it might be. And maybe ARM's perspective on whether some of these small players struggle before they've even started.

Lakshmi Mandyam: I think at ARM we're like Switzerland. We love everyone.

Andrew Feldman: I don't.

Lakshmi Mandyam: The new entrants getting market traction today because they have brought some innovation to a new area where there was none. In the case of a Calxeda, no one had really thought about SOCs in the server realm. They certainly had in networking – that's the way of that world. But in servers, that was not. So I think about the fact that they brought some innovation, and that they really had a direction shift.

In the case of AppliedMicro, they brought 64-bit ARM to the market first, right. So, [to Andrew] don't shoot me for saying that.

But the end users will pick who they're going to go with. So I can't give you a direct answer of: "I think there's going to be one person standing". If that was the case, I think we would have lost the value proposition of ARM and choice.

Andrew Feldman: I would say this. How many people here went to business school? And they forced you to read Geoffrey Moore's book? Right. In our industry, it's just crap. Doesn't hold at all, right. We were taught that there were this group of early adopters, and there was this chasm over which you had to cross. It was very difficult.

In this industry now, it's the exact opposite. The early adopters are now huge. And we have a full bimodal distribution, Okay. And so the early adopters, the guys who used to be the tinkerers and the players around, with whom you could never make a business, are now called Facebook, right?

And to give you an idea: JPMorgan Chase is a 120-year-old institution; in their 98th year of life they became one of the Top 10 consumers of compute in the world.

Just after its fourth year of life, Facebook entered that list. OK, that's what we're talking about. That's why this area of early adopters is no longer a little part of the market in which you can't make money. It's now a monster part of the market in which you can make money.

And that's why we have a different dynamic right now in the data center around network function virtualization, around SDN, around a whole set around storage, because the early adopter is a place that money can be made.

And to the extent that our firms like Calxeda or AppliedMicro can carve out niches where they're delivering to guys with specific needs who, in the past, were given a general solution, they can survive, and they'll be just fine.

Unidentified Audience Member: Thank you. The 25% share position for ARM assumes that there's a share loser, clearly Intel in this case. I know why you feel you can win., but why is it that they're going to be so easy to roll over and lose?

Why wouldn't an alternative scenario be that ARM only gains 5% or 10% of the share, but is used as a point of leverage for the mega-data centers to get Intel's pricing down? I realize there are performance/power benefits. But isn't this really about economics in the end?

Andrew Feldman: I think if we run the world forward 150 times, that could happen. It's not what I think will happen. It's not what I think will happen, because we have a long history in this industry of unfolding in a certain way. And we have seen Intel make the exact same mistakes repeatedly.

They made the mistake in 64-bit with x86. Customers wanted it. They wanted to keep it in Itanium. They kept it away. They made the same problem when they told Apple you have to use Atom, which was a terrible part. And Apple said: “We don't want it".

And then they gave birth to the tablet market. And Intel owned zero, right. Intel made the same mistake in the handset market. They thought they could have it. If you went to IDC 4, 5 years ago, what you saw was a huge array of handsets. They won zero.

I mean, there's a long history of this company making this exact type of mistake. And it's not because they're bad. They're some of the smartest people around. It's because everybody in power came up with a particular view, is a classic example of innovator’s dilemma.

Unidentified Audience Member: Hi. I had a quick question on technology. You know, there's a lot of fabrics and interconnects out there. SeaMicro did a great job with Freedom Fabric. So it's at AMD now.

I'm just kind of curious how you think that that landscape is playing out, right? Because everyone seems to have their own fabric. Everyone's got their own interconnect. What are your customers telling you? Are the Tier-1 OEMs happy with kind of adopting proprietary fabrics? Are they looking to do something different? What do you think?

Andrew Feldman: Sure. I think the truth of the market is that what's been in the box has always been proprietary. For example, a Cisco switch: if it's in the box, it's proprietary. If it faces out of the box, right, then it's standards space.

But switch fabrics were always proprietary. The technology that links an Intel CPU to chipset, the front-side bus technology, is absolutely proprietary.

The stuff in the box, because nobody plugs into it, nobody cares about whether there's proprietary stuff in there. What you have to do, is you have to be sure the standard Ethernet gets in, or InfiniBand, or Fibre Channel, or one of the other ways that you can get into the box. You have to be sure that the disks are standard, and that the memory that's plugged in is standard.

But how you connect those CPUs inside the box to make a cluster, we just don't have people who ask us and worry about the fact that the Freedom Fabric is ours. We just don't see it.

To answer your broader question, the reason people want fabrics is very simple, and it's really important we get that right. One of the CTOs at Dell said this best. And he said, you know, you can pull a cart with 2 oxen or 10 dogs. Right, you can pull a cart either way.

Now, if you guys hold that thought in your mind, for oxen, we use a big heavy wood yoke. And if you've ever run with dogs, you use a super lightweight harness. And that's true, too, when we try and build clusters. We don't want heavyweight yokes on large arrays of processors. It's too heavy. It burdens them with power and cost.

What we want is a super lightweight harness that ties everything together and allows a little bit of freedom. And that's why you're seeing people move away from traditional Ethernet as a fabric, and to something that's lighter weight inside the box. And then when you want to connect boxes together, you go with the Ethernet or one of the standard-based technologies.

Unidentified Audience Member: Yes. Andrew, I guess ARM introduced the 64-bit architecture 2 years ago. And then you guys are obviously taking one implementation of that for your initial design.

So early feedback from you as an early adopter – in the big architecture that they've announced, or in the specific architecture that you're taking, what are you missing? What could you have more of, or what features are you missing that would allow you to address bigger parts of the market that you might not have in that 25% number that you laid out?

Andrew Feldman: Sure. I think why we need to go to 64-bit is very simple: we needed more DRAM, right. The reason that the server guys need 64-bit is these servers now are delivery mechanisms for compute cycles and DRAM. That's what we do. And 32-bit, even with some fancy tweaking, can't get you the amount of DRAM support. And so that's why we went to 64-bit.

What we don't have, you know, I think there are -- there are a collection of capabilities. I'd like to be able to clock the part up higher with a Cortex-A57 part.

And the guys at -- at Broadcom announced that they were going to do that with an architecture license. And the guys at AppliedMicro announced they're going to do that with an architecture license. So clearly, the architecture can go there.

I think ARM, in its -- you know, beyond Cortex-A57, will do a design for everyone that allows you to crank up the clock a little bit. That would be nice. I don't think you need to get to 3.5. I think 2.5 or 2.7, 2.9 is fine. I think that's, you know, one of the more important ones for the future. But I think what we'll have is a combination.

If you sit back and say: “how did the server processor world work eight years ago?” We had one design, Intel had one design, and we binned parts. We pretended they were different SKUs, right. This is the way the world works. And if a part came off the bin at the high end, we sold it at 3.5, or the highest clock rate. And we sold for the most amount of money.

And if the part was less within range, slightly out of spec, and only worked at 2.0, 2.5, we called it our low end product. That's not the way the world is going to work going forward.

All right, Calxeda's going to build a part for somebody with a particular segment in mind. And AppliedMicro's going to go after a different segment. And we're going to build a part that tries to cover six or eight segments with different characteristics. And then somebody else is going to try and build a different part.

And it's not one part being packaged on various arrays of whether it's within spec or out of spec. It's individual parts built for separate needs, and separate workloads. And that's what I mean when I say the death of one size fits all.

And so I think that more range on the clock helps everybody. And that's why I brought that one to mind. But we have a bunch of ideas that I'm not going to share today, but that could give a part a unique advantage.

Lakshmi Mandyam: When we introduce CPUs, they start at a certain performance point because of the selection of process technology. And early adopters have a certain design point that they're going after.

But if we look at over the next year or so after that, when people make a process change, when we do more optimizations of the libraries and things that the Physical IP team works on, and the processor optimization stuff that we do, we tend to see that people actually are able to achieve higher frequencies and performance points than the cores were originally designed for.

And I think the other benefit to ‘one size does not fit all’ is it changes the economics. Before, you had to overdesign and live with what you got. Now, when you're designing to a given performance point, your yield point actually goes up significantly higher. So it kind of changes the economics of the whole thing, I think. And that's an important point to remember.

Unidentified Audience Member: Thank you. On the -- on the first round of servers you introduced were Atom-based, right?

Andrew Feldman: They were.

Unidentified Audience Member: The second round, you're doing -- as mentioned by moderator, based on the Opteron to Verizon design. There have been ARM-based devices out there, like Advanced Micro, Calxeda, that you could have implemented. So back with that 25%, in your majority of market share statement, and you said they're going to have parts for other applications.

So let me ask this question. What were the workloads you were addressing that were not being taken care of by the current devices? Did you have to make a device for it yourself? Did you believe that it would address a big chunk of that 25%?

Andrew Feldman: So we did begin with Atom. And it was an underpowered part. It was a 32-bit part. And we learned immediately that that was the wrong place to be. We convinced Intel, and to my regret right now, we convinced them to embark on a roadmap that produced a 64-bit Atom part.

But because, in answer to the question we had earlier, they chose to limit the memory controller. They only allowed it to support 4-gig of DRAM, limiting its applicability.

By this point, we'd made a decision that we couldn't live under those scenarios. The workloads required more DRAM. And so it wasn't even a performance issue – we needed more DRAM.

And so we went to the next class up, which was single-socket Opteron and single-socket Xeon parts. Now, we would be very happy dropping down a little bit into a slightly smaller part if we could continue to get the vast amounts of DRAM that our customers are demanding.

And, you know, when we built the system, and the system we have right now, there were no ARM options, right. And so running both the server business and the system business at AMD, we looked long and hard, and thought this was a market in which we had unique differentiation. We had a real opportunity based on the fact that the core is right now only a modest portion of the value of a CPU. The IO fabric, the memory controller, the security, the RAS features, reliability, these are enormously valuable in a server processor.

And we came to believe that this market was going to explode, and that we had a set of capabilities, a set of IP blocks that would position us well. Which specific workloads? Web Tier, Apache, PHP, Java, right. If you pop up a level, Duke, Cassandra, True Web Tier, whether it's in Genx or any number of the others. These are the workloads that we think the part we're building is particularly well suited for.

Simon Schafer: Fantastic. Well, we learned a lot. So thank you, Andrew and Lakshmi for this panel. And I want to thank the audience as well for participating on all panels.