Arm: Making Progress vs Strategy

Ian Thornton (Head of Investor Relations, Arm Holdings plc)

November 30, 2017

Good morning, everybody. It is really very good to be back here again today. For this presentation, I am very much going to be focusing on, less about who Arm is and how we make money, and much more about some of the progress that we have been making against our strategy in the year that we have been part of SoftBank.

[#: indicates slide number]

#2

What I am going to do today is first of all, I am going to give just a very, very brief refresher on Arm and our business. Then, I am going to give a quick update on our business performance in the first half of this year. Then, I am going to dive into some of these key strategic areas for our business over the next 10 years, which is particularly in the areas of servers, of artificial intelligence and machine learning, and also in autonomous vehicles in self-driving cars.

#3-4

But firstly, a quick reminder, Arm is in the business of semiconductor design and the creation of computer chips. Over the past 50 years or so, chips have become much more complex as the consumer electronic devices that we like to buy have become more advanced. 50 years ago, one person could design a computer chip. Now it takes tens of teams of engineers, each team being hundreds of people strong. That is because chip design has become, as you can see, far more complex.

#5

How does Arm play in this? Well, we provide some of the key design components that go into computer chips. We are most known for the main processor, which runs the software in your consumer electronics product. Every time you use your smartphone, open an app and play a game, or write an email, you are interacting with software, and all of that software runs on an Arm processor somewhere in your phone. However, equally, every time you put the brakes on in your car, the anti-lock braking system in your car also is controlled by software. We do not just do the main processor. We also have a graphics business as well, so we develop graphics processors. We developed accelerators for things like machine learning. We also do radio protocol stacks for things like Bluetooth and Wi-Fi, memory interfaces, and the like. Therefore, we develop many (but not all) of the components that you need to design a computer chip, and we then provide those to our customers.

#6

Our business model is that we develop the design of these components. We license those designs to semiconductor companies. They pay us a license fee and a royalty. One interesting thing to note, though, is that, if you look at our business today, the technology that we are developing today is for products that we will license in the future and we will get royalties for many years to come. The revenues that we get today are entirely derived from technologies that we developed years ago. Therefore, the license revenues are for products we developed in the last three to five years. The royalties, some of the royalties come from products that we designed 25 years ago. Therefore, when looking at our business today, the revenues and the costs are quite different. The revenues are from products developed up to 10 to 20 years ago, and the technologies that we are developing today, so our costs today are the products that will earn revenues for in the future.

#7

However, this has helped us become a very highly profitable and a very cash-generative business. What are we doing with that cash? Well, we estimate that, all things being equal, we will generate about $0.5 billion dollars' worth of cash this year, but SoftBank has asked us to reinvest that back into the business. Therefore, we are at the moment increasing our investments in R&D to accelerate our R&D research, so to bring the future forwards, to develop technologies today that otherwise we would have developed in two to three years' time. Some of that is being reinvested back into the business, and some of it is being used to create a new business.

I have mentioned to you before about mbed Cloud, which is a new technology that we are developing to help secure the internet of things. It is a relatively small business today, but it is growing, and it is where we are investing a lot of money at the moment. I am not going to spend too much time talking about this today. I think the next time I come back maybe in three months' time will be a time when I will spend maybe a longer session talking about this new business. It is not yet ready for scrutiny, but it has been moving.

#8

Where are we investing? Well, broadly in these five areas: artificial intelligence and machine learning are workloads that are going to appear in pretty much every product that we use, be it working on the cloud or even in your phone or in your car. Intelligence and machine learning is going to be everywhere, so it is going to become a very important technology for us in the future. Some of that intelligence is going to go into vehicles and into robots. That is going to help them to become more autonomous and more capable devices. Some of our investment is going into computer vision to help computers to see and to understand their physical context, and that also helps bring about technologies such as virtual reality and augmented reality, which we can then experience. We have been putting for maybe the last 10 years quite a bit of investment in servers and into hyperscale technology. We will talk a little bit about our progress there.

Then, all of these are underpinned by the need for increased security, because I think we have all seen from recent hacks and cyber-attacks that if you attach anything to the internet, someone is going to try and steal it. Or, if they cannot steal it, they will deny you access to it, and so therefore try and extort money from a ransomware attack. Therefore, we need to make sure that all of these technologies are secure against a malicious attack.

#9

Just a quick reminder of our business model, and then I will talk about the H1 result. Arm develops processors. We license those to a semiconductor company and they pay us a one-off upfront fee to gain access to that design. About 30% to 35% of our revenue comes from the license fees. The semiconductor company can then go and put our technology into one of their chips, so they can build lots of chips like this. They can then sell those chips to their customer to go into a phone, or TV, or a car. Then, every single time the OEM buys a chip from a semiconductor company, for every single chip that gets sold with an Arm technology in, we get a royalty. About 60% of our revenues today come from the per-chip royalties. I will talk a little bit about both of these in a moment.

#10-11

First of all, licensing, in Q2 we signed 26 processor licenses. What I have done here is I have laid out for the years 2013 through to 2017 the per quarter number of licenses that we signed. The lowest we have signed is about 20. The maximum we have signed is about 56 I think, but you can see that the typical range is a roundabout 25 to 40. That is where I would expect. In most quarters, that is the number of licenses that we signed. In Q2 it was 26. It was very much at the bottom end of the range, which you have to then ask about, "Well, should I be worried? Is that a concern?"

If we dig into the numbers a little bit more, what would we normally expect to sign in an average quarter? Well, maybe we would license one classic processor, maybe an Arm 7 or an Arm 9. We did not sign any in this quarter, but that is nothing to be worried about. It is only one difference. Normally, we would sign about eight Cortex-A processor licenses. We are one off there, so is there anything to be worried about? Probably not. We normally signed roundabout four Cortex-R processor licenses, and we signed four. We have signed about four or five Mali licenses, so that is normal. However, normally, we would sign about 20 Cortex-M processor licenses, and we only signed 11. Is there something to be worried about? Certainly, this needed more investigation.

#12

But if I cast our minds back just over 90 days, at Q1 results, you may recall that Mr. Son put up this slide when talking about SoftBank's Q1 results. Just to remind you, at this presentation he was talking about how Arm is talking its two most popular processor products, Cortex-M0 and Cortex-M3 and was going to make them available for free, or at least for no upfront fee to anybody who wanted it. The reason for doing this was to help accelerate companies who want to design chips going into the internet of things, and, of course, a lot of SoftBank's future business is believed to be derived from IoT, so this will help accelerate IoT.

#13

How have we done? Well, there are 92 days in Q2, and we signed 92 of these new licenses. Even my math tells me that is about one per day on average. As part of this, we also have an evaluation program where you can gain access to these Cortex-M processor licenses on an evaluation basis. You cannot manufacture with these, but these are processors that are optimized to go into prototyping platforms such as FPGA. Again, there are no fees associated with this. It is all free, and from the end of June, we have seen people download over 1700 Cortex-M processors. Now, bear in your mind, before this, this was not available, so it has gone from 0 to 1700 in just over four months.

#14

Therefore, if we have actually signed 92 processor licenses, what would that have done to our Q2? Well, if I included the 92 processor licenses with no upfront fees, it would have been by far the best ever quarter we had have ever had. Q2 was down here. Now it is actually off the top of the chart with 118 licenses signed. Therefore, instead of having 11 Cortex-M processors, we signed 103. That is quite a substantial extra set of licenses, but if we had not made this for free, would they all have come to us and paid us for the Cortex-M processor license? Probably not, but we probably did lose around 10 paid-for licenses in Q2. Now, these are relatively speaking quite low-cost licenses, but we have probably lost $1 to $2 million worth of license revenue in the second quarter because we have this new program where we are giving the processors away for no upfront fee, which sounds like maybe that is a bad thing to do, but we did not lose just 10 to 20 licenses. We also gained a total of 92 companies using our technology.

The question I ask is, what would have happened had we not made Cortex-M0 and Cortex-M3 available for free? What would these companies have done? Well, there have been lots of free processors on the internet for some time now. You can download an 8- or 16-bit processor and use that in your chip design. That is what many companies have done for years. Often, we talk about Arm having about a 35% market share, and the other 65% (a lot of those being very dumb processors without much intelligence), well some of those have been downloaded for free.

However, in things like the internet of things, actually the 8- and 16-bit microprocessors are not good enough. They are not smart enough. For the internet of things, you need a 32-bit microprocessor. Therefore, our concern was that, if we do nothing, then these companies needing a 32-bit microprocessor would be creating enough demand that somebody would start to fill that demand by creating a new 32-bit microprocessor to be downloaded for free, and effectively create a competitor to Arm. Therefore, by making Arm processors available for free, we have now, hopefully, fulfilled that demand.

Now, although these are license-fee free, they do have a royalty associated with them. There is about 4% royalty per chip from anybody who actually ends up manufacturing and selling a chip. Actually, the first 1000 chips are free, but the 1001st is when you start paying a royalty. We do hope that in some point in the future we will start to see royalty revenue from some of these. I did have a debate to that whether to actually have this as my headline number for licensing. 118 licenses would have quite a nice number to have had.

#15-16

Okay, so let us have a look at the royalty side of the business, then. Now, when looking at Arm's royalties, we need to first look at how the industry as a whole has been doing. If you look at the overall industry, then this is for fiscal H1 2017. The semiconductor industry as a whole was about $190 billion worth of computer chips, actually showing some very, very strong growth. 21% growth year-on-year, it is very strong. This actually is a very strange number because, for the previous five years, the semiconductor industry has grown on an average about 1%, so 1%, 1%, 1%, 1%, 1%, 21%. That is a very, very unusual growth. This is primarily being driven by a very strong sales of memory chips, particularly into servers, but also into other consumer electronic products like phones and TVs, and increasingly cars as well. Also, there are lots of graphics processors from companies like NVIDIA in here as well, also the chips from Intel.

However, Arm is not exposed to all of those markets. You do not find Arm processors in memory chips, for example. Therefore, if we just look at the market that is applicable to Arm, just chips that could have an Arm processor in them, then the market is a bit smaller. It is just over half the size or it is on the half size. This market grew about 9% year-on-year, which is still pretty good growth. Normally, this market would be growing at about 2% to 5%. That is typically its growth range over the past five years or so. If you look at how Arm did, there is about $450 million worth of royalty revenue, so there is just about half a billion dollars of royalty revenues. That is up 11% year-on-year, so slightly ahead of the industry, which is good.

#17

However, one of the interesting things here is that the value is up 11%, but the volume is up nearly 30%, so almost three times greater volume in numbers than in value. Obviously, that implies that the average royalty that we are getting per chip has declined over the last year, so the question there again is, should we be concerned? Let us have a look at this. If I look back over the last maybe 20 years, you can see that the average royalty revenue per chip actually came down over the first part of this period, but then since then has been pretty flattish. That has been driven by two things, really.

Firstly, over this period, Arm processors have increasingly been deployed in microcontrollers and smart cards. Now, these chips have an average selling price of less than $1, and the royalty that we get per chip is about 1%. Therefore, we have been getting less than a penny per chip from all of those chips. They ship in huge volumes, but at the same time, we have also seen the growth of the smartphone market and also the growth of Armv7 and then Armv8, which have much higher royalties on average. A typical smartphone may have a $10 chip and may have Armv8 plus MALI, so maybe we get 2% or 3% royalties from a smartphone, so maybe 30 cents from each chip going into the smartphone. Therefore, we have seen growth of smartphones with 30 cents royalties and growth of microcontrollers with 1% royalties. What has that done on to the average? It has roughly cancelled out because the growth of microcontrollers and the growth of smartphones have happened at about the same time.

#18

What has driven this line then is basically four things: firstly, the average selling price of the chip, so it is $10 chips versus $1 chips. What is the average royalty rate percentage per chip? That is 2% versus 1%, which customers are winning share and which markets are growing faster. By far, the biggest influence on this graph is the market mix. This is microcontrollers growing faster than smartphones, or smartphones growing faster than microcontrollers has the biggest impact. Then also the royalty rate that we are getting, because every time we design a new processor, every time we license that new design, we try and get a slightly higher royalty percentage per chip.

However, over time, these have all netted out to be about the same, to be flat. You can see this to a certain extent is because Cortex-M is the technology going into most of these microcontrollers and smart cards. You can see that since the last 10 years, we have grown from practically zero to around seven to eight billion of these chips. Cortex-A has grown from practically zero to about half that number, about four billion chips. We tend to get one penny for these. We tend to get tens of cents for these. Therefore, the growth of this versus the growth of this has given this sort of flattish line. Going forwards, we expect that the average royalty revenue per chip probably stays fairly stable because the growth of Cortex-A will continue and the growth of Cortex-M will continue.

If I go even further forward, however, if the internet of things does indeed become a trillion devices, then that does mean a trillion chips at a dollar or less than a dollar, so over a long period of time, I would expect that do decline, but driven by the mix of more low-cost chips. From a profitability point of view of, though, all of these are 100% profit, so it does not really matter whether I get a dollar of royalties from 100 chips or a dollar of royalties from 10 chips. It is still a 100% profit. It is still a 100% margin.

#19-20

On profitability, let us have a look at our investments so far in the past year. Since we have been part of SoftBank, as you know, we have been asked to significantly increase investments. As you can see, we have, in fact, doubled the normal run rate of hiring. Normally in a year we would add somewhere between 500 to 600 new people. We have actually added this time 1100 net new members of staff. The vast majority of these are, as you would expect, are engineers. They are the people who actually make stuff and make the products that we will license in the future. The number of non-engineers has stayed proportionally about the same, maybe a few less.

One of the things that we are focusing most on is, as well as getting lots of people into the company, is making sure we are getting the right people into the company, so focusing on the quality of the people and also then once they are inside, inculcating them with Arm's culture, our values and behaviors, so that they integrate with the rest of the company as quickly as possible. We understand how to do this. In the five years before the acquisition, Arm doubled in size. We have already gone through the process of doubling in size and realizing that over half of the company has joined in the last two to three years. We understand how difficult that is. It is harder when you are doing it on even larger numbers, but it is something that we have done before. It is something that we are going to have to continue doing for a number of years. We plan to double the headcount from 2016 to 2021, and so we need to sustain this growth rate for a number of years to continue.

#21

Increasing headcount obviously increased our costs. The majority of the growth in our costs is related to the headcount. This £91 million here is entirely related to the addition of those extra people. In fact, if you compare these extra £91 to the £300 million from H1 2016, that is indeed about a 30% increase. Add in a bit of wage inflation, and that pretty much tallies up. There is a couple of items related effectively to the acquisition that you can still see in our costs as well. The reclassification of share-based compensation, which was excluded from this calculation a year ago, is now cash-based compensation, and so it now appears in our costs. Post-acquisition, we are replacing our bonus scheme, so that is really one-off extra cost that we are now lapping, so both of these, when we get to Q3, will go away. We are also lapping in Brexit, which was in June 2016, and so in the last couple of quarters, sequentially, the exchange rate has been more stable, so I am expecting this to disappear as well. Therefore, going forwards, Arm's increasing costs are going to be very closely related to the increase in headcount, which is kind of what you would expect.

#22

Clearly, revenue is growing strongly, up 17% year-on-year, but as we just saw, costs growing even more strongly. This has had an impact on our profits, which are significantly down year-on-year. However, if you remember six months ago, when I presented to you about Arm's plans for 2017, I said that we expected our EBIT margin for 2017 to be 18%, or about half. I think it was at 40%. I said it would be the half. Indeed, it is 18%, so it is pretty much exactly where we thought it was going to be. In fact, in our internal models, we had 18% for 2017, so that is exactly what we thought it was going to be. Of course, we are continuing to invest at a rate higher than our expected revenue growth, so we still expect that our margin this year for 2017 would be lower in 2018, so you should expect that number to continue to decline.

#23-24

Okay, I am now going to run through some of our key strategies and how we have been investing in those. The ones I am mainly going to be touching on here are machine learning, autonomous vehicles, but first of all, I am going to start with talking about Arm in servers.

#25

Now, as many of you have known, Arm has been investing in servers for many years. We are taking the opportunity because the server market is changing. The historical server market has been an enterprise server market where each server ran a very wide range of tasks. Indeed, this is for 2016, and you can see standard enterprise servers dominate the market, but increasingly, we are seeing cloud servers and high-performance computing servers taking over from the standard servers being the main market. It is these markets that we are aiming for, so we are still not targeting the standard server. We are very much targeting cloud and high-performance computing.

#26

The journey that we have been on, as I said, it has been a 10-year journey starting in 2008. In order to start this journey, we first need to develop some new Arm technology, which was Armv8. It took us two to three years to develop the first architecture and first processors for v8. We could then provide those to our customers. Now, our customers, they, of course, have to build a chip that also takes two to three years, but once you have a chip working, that is the first time the software engineers can really start to develop the software, and so we provided chip to some of the software engineers, and then they start to work on getting the operating systems working, and porting some of the applications. That also takes maybe one to two years, or a bit quicker. However, once we have the software, then we have working chips, working software, and it can go to the server users for the first time so that they can start evaluating it. Therefore, really, even though we started in 2008, it was not really until 2014 that companies like Google and Facebook were really able to start using Arm-based servers in their systems and start to do some trials and some proofs of concept.

Now, throughout this process, at each step there is lots of feedback generated. The chip guy is telling us how to do a better architecture. The software guys are asking for more features. The end user is saying, "Hmm, this did not work so good," or, "This worked well.

Can we have more of this, less of this?" so lots of information coming back. That then helps us improve the technology for future generations.

#27

Where we are today, we are just introducing the fourth variant of Armv8, which I will talk a little bit more on the next slide, but we now have multiple production chips from multiple different semiconductor companies targeting the server space. We have all of the main operating systems and many applications now ported, optimized, and commercially available for Arm-based server chips. That, therefore, means that we are only now able to see the first production deployments of Arm-based servers. This is how long it takes to enter a market. I said 10 years. It has indeed been 10 years.

Looking at some of the Arm technologies, if we go back to Armv7, the seventh version of the Arm architecture, which was introduced in 2004, the big new technology that that introduced was multi-core processing for the first time, and also some multimedia acceleration built in to the processor. The v8, as I mentioned, we started working in 2008, but did not really have the first version available until 2011. That introduced 64-bit computing for the first time in an Arm processor, and improved the multi-core, improved the multimedia technology. With v8.2, that had better multimedia, better multi-core, better error management, which is very important in servers, but also introduced some machine learning and virtualization acceleration. This is the technology that is going in to the latest chips that will be going into smartphones that will be available next year. Therefore, this is state-of-the-art today effectively in terms of what is about to become commercially available in smartphones and servers. As I mentioned earlier, we have just introduced the fourth variant of Armv8, and I will go into the future processors that we are working on at the moment and have code names for, but I cannot tell you about the details of those just yet.

#28

We now have seven companies, seven customers that now have commercially-available production silicon, which you can go and buy chip from today.

#29

The most recent company to announce support, announce their chips has been Qualcomm. This is from a presentation they gave about two weeks ago where they sort of dug into how they perform against Intel's incumbent Xeon processor. It is interesting that they used the most advanced Intel Xeon processor in order to benchmark themselves against. They were using the standard benchmarks of SPECint. Actually, for performance, they are pretty good. For some of the SPECint benchmarks, they were able to demonstrate that they can deliver about the same amount of performance as the Intel chips, which is good. On performance per thread, they were again pretty much the same as the highest performance Intel chip, but where they have really shone was in performance per watt. This is a much,

much lower power implementation than what Intel can normally produce. This is demonstrating between 30% and 50% better power performance for the same level of throughput. Because Qualcomm is selling these chips much more cheaply than Intel, where they really shine is in the performance per dollar. Therefore, they are able to demonstrate a two- to four-times improvement or for the same level of performance you have to spend a lot less money. Therefore, the combination, we think, of lower upfront costs and also then lower energy bills, but for the same level of performance, we think it is going to help to give us the opportunity to win some share in this market.

#30

As I mentioned earlier, all of the operating systems are now available. Red Hat and CentOS were the most recent products to have commercial support from. It has been really interesting to see that we have started getting a lot more market pull from companies wanting to deploy Arm-based servers with this technology, and that helps a lot to get them to support it.

#31

We have also started to see more end users of servers announce that they are migrating part of their system over to using Arm-based server chips. Amazon, Microsoft, Alibaba, and Tencent, all now publicly supporting in a small part of their data centers some Arm-based technology. I think this is only the start though, as in I would not want anyone to get overexcited by these announcements. Last year, we had less than a 1% market share of the server market. I think for 2017 maybe we get 1%, but I do think that we are now on target for our goal of a 25% market share in 2022. We are still looking that far out for our 25% market share.

What I think we can expect, though, in the next two years is a very strong response from Intel. Therefore, I am expecting there to be a lot of very strong competition from Intel, particularly in reducing the cost of their products to be more price-competitive with the Arm solutions, providing more lower power technologies to compete with the lower power of these Arm-based chips as well. We will see how that develops. We are expecting, as I said, a very strong response. Even though this is lots of good news, I am expecting there to be some bad days as well as good days over the next few years.

#32-33

Okay, now to talk a little bit about artificial intelligence and machine learning for Arm. Many of you will already have heard a lot about artificial intelligence over the last year or so. Almost all of that will be in the context of cloud computing, big data analytics, possibly with graphics processors either playing computer games against humans, or image recognition, or digital assistants such as Alexa and Siri, and things like that. All of that technology that you will be familiar with is done in the cloud.

At Arm, we are focusing on bringing artificial intelligence and machine learning out of the cloud, out of big data and bring it into the consumer electronics products that we use every day, and having machine learning not just working on big data but also on little data and even personal and private data. We want to see your mobile phone being able to understand your voice, to be able to understand natural language as you speak it without needing to send commands and information up to the cloud, for it to learn how you speak. We want cars to be able to drive themselves without having to wait for information from the cloud about whether they should be slamming the brakes on or not, and for robots to be able to understand their physical context and be able to respond appropriately. Even IoT devices, even tiny electric motors and sensors can learn how to operate more effectively using some machine learning technology.

#34

Now, you all have used, probably even today, some machine learning technology without even being aware that you are doing so. One of the early versions of machine learning that you would have engaged with is if you have a Samsung phone or an Apple phone, you can probably unlock that phone using your fingerprint. Probably when you first got the phone, you trained it to learn your fingerprint by just pressing your finger five or six times on the sensor until it told you, “Oh, yes, yes, I now can recognize your fingerprint.” This is using the same technology that Google uses for ImageNet that looks at thousands and thousands or millions of pictures of cats and dogs to be able to tell the difference between a cat and a dog. Now, this is trying to do something much simpler, is it your fingerprint or is it not, but it is the same algorithm. The first five or six times you pressed your finger on the sensor, that was training the algorithm, and now every time you unlock it, it runs an inference engine that just says, “Is this is the same as the image that I have learned?” It is machine learning.

#35

We think that you will see increasing amounts of machine learning being done at the edge because of things such as security. If you have a security camera, you do not necessarily want every employee that works for you being beamed up to Google’s cloud. You want to be able to do face recognition and local to the camera. For reasons of privacy, I do not want Amazon listening into every conversation I have at home, but I do like the idea of a digital assistant. My fear is that, if I have an argument with my wife and Alexa is listening in, I am going to get an email later in the day, saying, “You might want to buy your wife a present. You were on the wrong end of that argument.” Therefore, privacy is going to become very important as well, and therefore, we want machine learning to be done at the edge.

In a car, latency is very important. If you are driving at 60 miles an hour and somebody steps in front of you, at that speed you need the car to make the decision itself. It cannot wait for the cloud to tell it what to do. At 60 miles an hour, the speed of light is not your friend. For reasons of bandwidth and power and cost, we need machine learning to be done at the edge.

We need to be able to train devices to think for themselves.

#36

Now for machine learning algorithms to be run at the edge, we need to make sure that we are matching the needs of the algorithm to the capability of the hardware that is going to run the software. Therefore, in the case of a security camera, security camera needs to be able to do face detection, face tracking, and face recognition. That is a very different set of algorithms from, say, a mobile phone doing voice recognition and being your digital assistant. Different algorithms require different hardware blocks to be to best run them. Some of those algorithms are best run in a GPU. Some are best run in a DSP. Some may be in a RISC processor if it is decision making, and some can be accelerated using specific hardware accelerators. The needs of machine learning in a camera are very different from machine learning in a smartphone, and therefore, they need different hardware implementations to execute them.

#37

Arm's approach to machine learning in edge devices is to understand, what are the algorithms needed for this specific device? What is the available hardware? What CPUs, GPUs, image signal processors, computer vision engines, DSPs and so on are available in this chip? Then, compile the algorithm in order to make best use of the available hardware. This means you end up with an optimized software implementation to match the underlying hardware components. Then, all of this is available from Arm today.

#38

To help accelerate machine learning in today's (or rather into tomorrow's) smartphones, earlier this year, we introduced a new technology, a new microarchitecture for processors called DynamIQ. This accelerates by up to 50-fold artificial intelligence or some AI algorithms for things like TensorFlow and Caffe, which are some of the main libraries used in machine learning algorithms. This technology processes the Arm's licensing last year and this year. I would expect to see this technology appearing in smartphones introduced in about a year's time, so summer next year, you will be able to buy a smartphone with DynamIQ acceleration for machine learning.

#39

Finally, I am just going to run through a little bit around the road of autonomous vehicles. Although I am focusing here on cars, much of the technology that we are describing here also applies to robotics, to drones, to self-flying planes, and self-steering boats as well.

#40

One of my favorite quotes of this year has come from Audi saying that, 90% of their effort at the moment in innovation and developing new technology is going into improving the electronic systems, and in the software that is going into their future cars.

#41

That does not really surprise us when we look at the amount of software that is increasing in order to create a level three autonomous car. Level three, by the way, is a car that can drive itself in constrained circumstances. Very often, if you think of a motorway or a freeway with slow-moving traffic, then that is exactly where a level three autonomous car can take over and just drive you. If you are in a traffic jam, you can get a book out or your laptop out and do some work and the car will drive itself.

However, as you can see, the amount of software is expected to significantly increase. This is ideal for Arm. History has shown that in markets where the amount of software starts to significantly increase, this is the time when OEMs start to require their suppliers to standardize on a single architecture, because, if their investment in software is increasing, they do not want to have to worry about porting their software across multiple different architectures. They want a single architecture so they can write their software once and we use that software across chips across their entire platform. It makes their R&D much more efficient.

#42

This is some work that we have done recently looking at the compute requirements in a level three autonomous car, and comparing it to the compute requirements in a very high-end smartphone today. A high-end smartphone today needs around 40,000 DMIPS. ‘DMIP’ is a measure of compute requirements, and this includes the main processor and also all of the supporting chips in the phone, so the touch screen, the Wi-Fi, the Bluetooth, the modem. All of the technologies in the phone, all of the compute needs in a smartphone adds up to about 40,000 DMIPS. If we apply the same approach to the car, then the car needs about 10 times the processing requirements of the phone, so think of a car whose compute requirements have been 10 absolute top-of-the-range smartphones that you can get today.

#43

From my point of view, this is great, because I am often told that the smartphone market was unique, and they will never be a market as big and as successful as a smartphone market. Well, there may well be a billion of these sold each year, and only a 100 million of these sold each year, but if the compute requirements of this is 10 times this, then that means that the markets are equivalent in terms of compute requirements. Therefore, I am hoping they are going to be equivalent in terms of royalty opportunity as well, and that maybe I can get 10 times the royalties out of each car as we currently get out of each smartphone. Certainly, that is looking like it could be the case.

#44

We have had some recent good news in this area. Next year, Audi is introducing a new version of the A8. This is being marketed as the first commercially available car that supports level three autonomous driving. If you strip back the technology, the main brain is called zFAS. It is the processor or the board in the center of the car that actually sits under the driver seat, and within that there are two main chips, one from NVIDIA, which contains a Cortex-A15, and another from Intel, which contains an Arm Cortex-A9. Therefore, a lot of Arm technology going is to help make this car into a self-driving vehicle. That is coming available next year.

#45

Then, last quarter, DENSO announced that it had taken its first Arm processor license, the Cortex-R52. This is our highest performance real-time processor. It has been designed for safety-critical systems, and so it is ideal for advanced driver assistance systems needed in a self-driving vehicle. This product, the R52, has actually become very popular with the chip companies building chips for cars, and all of them have now licensed the R52 for their next generation of chips for the cars, even though most of them were not using Arm processors at all just one generation ago. Therefore, this is a very important technology that is now beginning to gain share in chip designs.

It is worth mentioning, though, that the car industry is quite slow moving in terms of getting new innovations to market. Therefore, although we have licensed R52 to the chip companies in 2016 and 2017, we do not expect those to really being in many cars until about maybe 2022 and beyond. Therefore, we are still a few years away before we start to see cars out there based on this technology.

#46

That is the last of my presentation, but I believe we have some time still for Q&A.

Q&A

Q1: I have a question about smartphones. Could you just remind us what kind of trends you are seeing for your products in smartphones and what do you think over the next couple of years about that market?

A1: Yes, sure. Thank you. Obviously, I was just putting smartphones in context for the moment. Obviously, the years of smartphones growing double digit year-on-year are long behind us. However, I do remain surprised at the strong growth we have seen this year with the smartphone growth being around about 5%, which is still a lot higher than I was expecting, but very nice to see.

The primary area of competition between the smartphone OEMs seems to be remaining to be, who has got the smartest smartphone? Who has got the smartphone with the best technology that nobody else has? That is helping us immensely because this is driving the technology providers, who provide the chips and other technology into smartphones to continue to invest in new R&D and in new technology areas in order to provide the OEMs with more and more technologies that go into their devices. As a provider of technology, that is only good for us.

In the sort of 2017 smartphone market, I guess the things that we are primarily seeing is virtual reality is becoming a standard requirement for the high-end smartphones. In 2016, it was really only the Galaxy S8 that supported VR capability. That is the ability to take a normal smartphone and attach it to a plastic headset and have a good VR experience whilst wearing a smartphone. That requires a lot of extra graphics support because these are quite high-resolution displays, and you need to update them very, very quickly so you do not get motion blur. Therefore, you need a lot of graphics processing capability.

Virtual reality at the moment is still quite limited, in that, if you have good eyesight, when you have the screen so close to your eyes, you can still see the pixels. If you see the pixels, then when you are looking around, you do not quite get the fully immersive experience that is required. To get to the point where you cannot see the pixels, you need to get to a screen which is 4K per eye, so 4K and 4K, an 8K screen. These are in development, still a few years away, but this has a 16-fold increase to the amounts of pixels, and therefore to the amount of processing that is required compared to today’s smartphone.

It is impossible to do that level of processing in a mobile phone. It is impossible. However, companies are working on new technologies to alleviate the problem. New headsets are going to have additional cameras which look at the eyeball and track where on the screen the eyeball is looking so that, as you look around, only the area we are actually looking at is done in high resolution and everything else can be blurry. This is a technique called “foveating”. I am expecting to see this coming into smartphone headsets in the 2019 and beyond timescale. This, again, adds a lot more processing requirement to a smartphone.

Apple is trying to get ahead of its rivals and is doing more investment in augmented reality. Augmented reality is the ability to be able to overlay graphics on the real world, usually through some kind of Google glass or headsets or just by holding the phone up. One of the key technologies, though, needed for augmented reality is for the phone to be able to understand its physical context, so basically using a combination of cameras and 3D sensors for it to be able to scan the room and understand what it is looking at, to understand this is a big room with lots of people as opposed to a small room with a poster on the wall of lots of people, so being able to understand how things fit together, then on top of that overlaying graphics that can interact with the room.

The technology at the moment is very rudimentary, and there is a lot of R&D being done by Arm and other companies to enable computers to understand their environment much better. This technology can be used in smartphones. It will be used in security cameras. It will be used in self-driving cars because, actually, nowadays, computers need to understand the visual world in many environments, also in things like robotics as well. Therefore, there is a lot of research being done by Arm and other companies in the area of computer vision.

Just going back to 2017 for the moment and the trends we are seeing there, every major OEM this year has announced that their latest handset has machine learning in it. It has some form of machine learning acceleration, some form of neural network accelerator. These are all, again, fairly basic. It is first-generation technology. I think the most important thing though is that all of the handset manufacturers have introduced toolkits to enable developers to gain access to the machine learning accelerators in their phones. The reason why that is important is that 100,000 brains are better than one because there are lot more software developers developing apps for mobile devices than there are, if you like, OEMs building phones.

In the same way that Apple did not invent Angry Birds, but actually Angry Birds became the app that really first made the smartphone work. A lot of the game developers are the people who really push the hardware platforms. Now, we have software developers gaining access to machine learning toolkits. There will be apps that come along in the future that will do things with machine learnings that nobody thought of. That will help to then set the agenda and the pace for the future innovation in machine learning in smartphones.

Therefore, I think it will be really interesting in the next two years to see how machine learning evolves within smartphones, but I am expecting that will be a lot of new technology being demanded from the OEMs to accelerate AI in their devices. That will be the main things.

Q2: To follow up, how about thinking about 5G handsets? Is there anything that is visible yet, or could you think through or the implications might be for Arm?

A2: I guess 5G for Arm is more data. It is more data. It is a bigger data pipe both up- and downlink. What do we do with that data when it arrives? I think inevitably that is going to create more opportunities for more processing. More processing requires more processors or faster processors that just is going to create more demand for Arm technology. I do not think that 5G specifically has any impact on Arm. Maybe it makes the modem a little bit smarter, but broadly I think it is the case of just more data. It is just more processing and therefore that may help us that way.

Q3: My question is concerning DesignStart Pro. The point of the question is license fee is supposed to be free for easier entry. Is it fair for me to understand that you are imposing a bit higher royalty rate? The second part of the question is, if they convert this contract from DesignStart Pro to the usual contract to pay license and the lower royalty fee, does it lead to some lost opportunities for Arm?

A3: Thank you. Just to run through the DesignStart license, it is zero upfront fee. Therefore, we move the barrier for anyone to gain access to the technology. It is also a very lightweight contract, only three pages, but there is, as I mentioned earlier, a 4% royalty per chip, although the first 1000 units are free, so for very low volume manufacturers, they will effectively pay nothing, but 4% on a Cortex-M0 is about four times higher than what most people would be paying. However, if you are not developing a small number of chips, that might be okay. If you start to hit higher volumes, then that is probably going to start becoming quite expensive. The door is always open for someone who wants to come to Arm and then maybe pay $1 million to buy down the royalty rate to a 1% royalty again. This is intended to be great for someone who wanted to try and do some R&D, great for someone doing some development. If you start to get into high-volume manufacturing, you probably do not want to be doing it on this. You will want to buy down a lower cost license.

Q4: So it is not necessary assuming large clients take this license structure, right?

A4: We are expecting that large companies will be taking normal the perpetual licenses. The companies that will be licensing this are going to be small teams within a large company, or a startup, or a group within a research department of the university who may then get spun out into a commercial organization.

Q5: Okay. This does not necessary explain the reason why the 11-license level is lower than the 20?

A5: I think it does because I suspect that 10 companies that would have scraped together some money to spend the $150,000 of their startup money on a Cortex-M processor license now do not have to. If you think of a startup, they have got a few million dollars' worth of VC funding, and they know what their daily cash burn is. $150,000 going to Arm is a big chunk of cash burn. Then they have the date where they have run out of money and they go out of business. If they can get a free version rather than spending $150,000 of their $2 million VC funding, actually, that may keep the lights on for another two weeks. Therefore, I can imagine that some companies, given the choice between a $150,000 license today or free license today, knowing they can always take the $150,000 license in a year's time once they have actually got something that works may well take the free version today. Therefore, I think this does explain why we only had 11 rather than 20 licenses. That is also why I think I would say we lost probably $1 million, $2 million worth of revenue out of the quarter, because those 10 companies that would have otherwise licensed or paid for Cortex-M processor license instead took the free one.

Q6: Thank you. My next question is, you talked about IoT device and microcontrollers, and you said that 32 bits are required for IoT rather than four bits or eight bits. Probably the devices are like lower one. That should be the case, so what I would like to make confirmation is, in term of the processing power, are 32 bits required for the processing power, or Arm is providing a development environment, so that has necessarily become the requirement for the 32 bits of the speed. Which is the case?

A6: Thank you. If you imagine a sensor in a field or a street light controller or something like that, these are running algorithms that are trying to make sense of the data that they are receiving about the world around them. A sensor in the field is taking temperature measurements or moisture measurements or something else. Maybe it is battery-based and it needs a lot of batteries that needs to last for the entire season. The worst thing it can do is turn on its radio because the thing that is going to really burn through that battery is going to be having the radio turned on transmitting data. Therefore, you need an algorithm to run inside that sensor that can do local data analytics, that can understand that the temperature is rising rapidly now, so therefore maybe the sun has come out and it is strong and something needs to happen, so it can send not just information or data, but actually actionable information that the farmer can do something with.

You do not want a sensor that sends four degrees, four degrees, four degrees, four degrees, four degrees, five degrees, five degrees, six. You want it to go, "Something has happened. You need to do something about it." That requires smarts going into the sensor. The more complex the algorithm, the more complex instructions that need to be run, this is what 32-bit, 16-bit, 8-bit represents. It represents the ability to have a complex algorithm to be run. That is why I say the vast majority of IoT devices will need a 32-bit processor to run.

If you look at some of the markets that used to be IoT but have crystallized into real markets. Let us say the drone market used to be an IoT market. Wearables used to be an IoT market. All of those are 32-bit. Arm has over 90% market share of the main controller chips in the drone, the main controller chips in a wearable, so I think that the vast majority of IoT devices will be 32-bit.

Q7: Thank you very much for your very easy-to-understand presentation. Maybe my question is a follow-up. Cortex-M, then for Cortex-M contracts or regular contract will probably the number of licenses or regular licensing will be low from now on. Depending on the first application, whether to use DesignStart or to go with the regular licensing scheme, that is my question. Does it depend on the first application?

A7: I think from a company's point of view, these are often more of the companies that only have one project. The companies that will be using DesignStart will probably only have one project that they are working on, but that could be anything. We see a lot of Cortex-M processors going into cameras right now. We are seeing here a lot of camera chips being developed for everything from industrial manufacturing, cars, security, as well as consumer electronics as well. A lot of Cortex-M class processors are going into the IoT applications, embedded connectivity applications, security applications as well. This license type is suitable for all of them.

Whether you pay for the license, or whether you use the free version, is probably going to be dependent on your confidence in actually having a business selling chips at the end of the day. If you are a large company and you have a new project, you have got a pretty good chance that you know that you are going to end up selling a chip. If you are a small startup with the first VC funding, you may never get to having a chip. Therefore, signing up to a 4% royalty rate is entirely a theoretical problem that you can deal with in a year's time. Actually, knowing that you do not have a pay anything now is going to be the main determiner to whether I pay it for license or a free license.

Q8: Therefore, it is solely on the volume probably that is the determining factor, not the application?

A8: Confidence of achieving volume, yes.

Q9: Okay, thank you. Second question, free cash flow, you talked about mbed Cloud, and you said you would give us the details for mbed Cloud. It is IoT product, right? Then, making investment in two or three years' time, you are expecting to get revenue. That is my understanding. In 10 years' time, when you look at your roadmap revenue in 10 years' time, what is the percentage of free cash flow that you are investing? That is the volume zone of their investment. Has there been any change before and after the acquisition by SoftBank, just briefly, please?

A9: To your last question first, post-acquisition, there has been a lot of increased investments in the area of mbed Cloud than there was previously. As a listed company, Arm very much had to balance where we could invest and where we needed to time that, how we needed to time those investments. Under SoftBank, our instruction from Masa is to be much more aggressive and to take much bigger risks. One of the areas that he was particularly wanting us to increase investment was in the area of mbed Cloud. At the moment, it is relatively small. At the point of acquisition, it was about 200 people working on this area, to date about 400, so doubled in size, but it is still less than 10% of our total engineering team. This will ramp over time.

In terms of the business that it is going after, this is going to be a cloud-based service that is going to help provide secure connectivity to IoT devices in the field. I think it is worth remembering though that the internet of things is still not here today in any defined way. It is not clear how the market will develop. It is not clear where the revenue flows are going to be strongest. It is not clear where the profit pools are going to be deepest. Therefore, we are having to experiment as much on the business model as we are on the technology side of things in this area.

We are partnering with companies like Microsoft, IBM, GE. We have real implementations with these companies. We have around about 300,000 IoT devices under management within mbed Cloud, but we want this to be 300 million in five years' time. We are a tiny way along in our journey. This is probably going to be the area which sees the most incremental investment over the next five years, but it is going to be done so with caution, not recklessness. I suspect that some of that investment is going to come through acquisition. There are companies that are developing similar technologies to ourselves, and rather than compete with them, we would rather have them on our side, so I am sure that we will be looking at some of those in order to see which are the right ones to combine with in order to create the right portfolio of technologies in order to take the best advantage of the internet of things.

This is also an area where we have opportunities to partner with other companies with the SoftBank ecosystem. There have been a lot of collaboration across SoftBank, and this is one of the areas that we have been doing a lot of work, particularly with Sprint. But, again, it is early days.

In terms of what does 10 years' time look like? Potentially, this is a bigger business in 10 years' time as Arm's IP products business is today, so there is a big opportunity here, but it is not entirely clear exactly what that would look like, and we need to be flexible in order to be able to take the opportunities as the IoT market forms. That is basically what we have been given by SoftBank, and it is something that we could not have done had we remained a listed company.

Q10: I have just one question. You talked about machine learning. There was a reference to Caffe and TensorFlow. If the performance has become accelerated by 50-fold, what type of services or the new services do you expect when the accelerated the machine learning is available?

A10: Yes, the primary things that we have been focusing on for machine learning at the edge has been in digital assistance for smartphones. That means the smartphone being able to understand your voice and your instructions. This actually is not as hard as it may sound. It is interesting that obviously Alexa requires huge datacenters to be able to understand your voice, but actually, if it is just your voice and the 500 words that you are likely to say, there is enough processing power in your phone today to do that. It will be done much more efficiently with DynamIQ and with some accelerators which the OEMs are developing, so smartphone digital assistance is one.

We are doing a lot of work around the cameras, around object identification, face detection, which is something slightly different, and face tracking as well. For small numbers of faces, to be able to do face recognition in a security camera as well. If it is large numbers of faces, it then would need to be offloaded on to a cloud device, but you can imagine that, if you are a small company, a security camera over the door to the entrance to the building you might be able to identify your 50 employees as they come in in the morning, so rather than you having to use your pass to get into the building, you just walk towards your building and the camera lets you in, or redirects the visitor to go to the reception. If you have 50,000 employees coming in in the morning, I am afraid the camera does not have the processing power to do face recognition on 50,000 faces, but it can work with a cloud server for something like that.

One of the most interesting ones, though, we are working on is actually is around IoT, so things like putting a vibration sensor on an electric motor so that the motor controller can tune itself to reduce the amount of vibrations, and thereby improve the efficiency of the motor. This is something that, over time, motors change as they wear, and the algorithm can be tweaked as the motor wears in order to remain efficient, quiet, and vibration free. This, of course, is done entirely without any connection to the internet, so it needs to be able to do it itself as the motor ages.

Therefore, really, if you are thinking or if you are looking at the tiniest IoT device all the way through to working with high-end servers in terms of machine learning is a technique, a workload that will go everywhere. Thank you.