Arm Business Briefing Q&A

March 21, 2024

Presenters:
Ian Thornton, Head of Investor Relations, Arm Holdings plc

Q&A

Speaker 1
Q1:
Let me ask you two questions, one about the next 10 years and the other about the current situation. First, in 10 years, you mentioned that you are targeting a non-GAAP operating income margin of 60%, so over 10 years, how much do you envision the top line and sales growing at a CAGR? What would be the drivers of that growth by sector?

A1:
(Thornton) Thank you for your question. Yes, we have obviously recently IPO'd. During the IPO process, you are not allowed to give too much guidance. We are still very close to the IPO, so we are not giving detailed guidance at the moment. Because if you give guidance during the IPO and you fail to meet your targets, then you could be in trouble. If you give guidance shortly after the IPO, the SEC will say, you should have provided that information during the IPO. We are not currently giving more detailed guidance than what I just mentioned. The only thing I would say in general, in terms of revenue trajectory, we have indicated that we think that royalty revenue can grow at high teens or 20%. And that licensed revenue should grow at mid-single-digit growth. However, in our most recent period, license revenue has grown more strongly than we had anticipated, growing mid-teens. At the moment, we are actually seeing stronger growth than we had expected. But we are not giving any more update to longer-term guidance because we are not certain how sustainable the near-term higher growth in license revenue will be.

Q2:
When you say that royalty revenue growth is in the high 10% to 20% range per year and license revenue growth is in the mid-single digit percent, is it correct to say that we can expect those numbers for about ten years?

A2:
(Thornton) Yes.

Q3:
Second, I believe that the sales guidance for the fourth quarter is about 40% annualized growth,

but I would like to know why it is so strong only in the fourth quarter, and what we should think about the sustainability of this. Could you give us information on this as far as it is publicly available.

A3:

(Thornton) When looking at a year-on-year growth number, you have to look at both the quarter that we are looking ahead to and also back to the quarter one year ago. The semiconductor industry is cyclical and goes through periods of rapid growth, very often followed by a period of contraction because the companies that buy chips buy too many chips. Then they have them on the shelf, and then they burn down their inventory. This period is known as an inventory correction. It happens in normal times every 18 months to two years. During the pandemic, there was a shortage of many chips. You may recall that particularly in the automotive market, some cars could not be sold because of some basic chips being unavailable. After the pandemic, we saw many companies, many OEMs significantly increase their inventory levels in order to protect themselves against a future pandemic or a return of the pandemic. The phrase that has been used was that companies had moved from just-in-time purchasing to just-in-case and therefore, had more chips than maybe they would normally have had. During 2022, we saw inventory starting to be consumed, so fewer chips were bought. That meant that semiconductor industry revenues declined. That decline went on until, the lowest point which was February 2023. That was the weakest quarter of the semiconductor industry cycle. Because Arm has approximately a 50% penetration into chips with CPUs sold into the semiconductor industry, we could not dodge that cycle. And so, our royalty revenues particularly were impacted by weaker sales during that period. At the same time, if semiconductor companies are seeing reduced revenues for their chip sales, then they may make decisions about reducing their R&D budgets by maybe starting fewer chip designs or postponing a chip design. That can also impact our licensing revenues as well. We did not see much reduction in license revenues a year ago. But you do not know what you do not know. The deal that you did not sign, you do not know whether it was there or not. During 2023, we have seen a strong recovery in the semiconductor industry. Month-on-month sales have been up every month since February 2023. And we expect further growth for Q4 2024. We are comparing what looked like a strong month with the weakest month in the cycle. Therefore, there is a very strong year-on-year growth. In addition, we are now seeing much stronger license revenue than we had anticipated. Much of this extra license revenue is related to companies being very excited about AI. Although much of the talk around AI has been in the data center, we are seeing companies that are building chips for edge devices, for smartphones, for smart TVs, for washing machines, also now wanting to build chips that are AI capable. But they have a problem. It can take two to three years to build a chip. So, companies are looking at what the market is going to be like in two years' time. In addition, what AI capabilities would be needed by customers in two years' time. But the AI models are changing so rapidly. Every six months, the models are slightly different, they are more capable. If it takes me two years, the models will have changed multiple times during that period. They are trying to hit a moving target. We have seen companies

license our more advanced technology and higher performance technology to try and future-proof the chips that they are starting today, so that when they come out in two years' time, they will be able to run the AI algorithms that will be required by that point. Because they are licensing our most advanced technology, that has a higher price tag associated with it, licensing revenue in the last few quarters has been higher than we had anticipated. It will also mean that the royalty in two to three years' time will be a bit higher, because if you license our most advanced technology today, you not only pay a higher license fee, but you sign up to a higher royalty fee in the future. In terms of sustainability, from a comparison point of view, we are currently comparing with the weakest quarter in the cycle. Next quarter, the quarter after, the quarter after, will be with strengthening quarters. But still, if you think of the downturn and the up-cycle, it is still going to be, the first few quarters will still be against relatively weak quarters. As to the sustainability of the strength in licensing, we do not know yet. We will be reporting full-year results on May 8th. At that time, we will take a view as to guiding for the following periods. We will have to see what our visibility looks like right now. From a personal view, I do not think the excitement around AI is going to come over next week. I think it is going to continue for some time. Hopefully, therefore, that means that we see continued demand from companies wanting to use Arm in their AI applications for many quarters to come.

Q4:

Regarding the sustainability of the Q4 sales you mentioned, the Wall Street consensus for next year is about +15%, which is four times Q4 sales. Is there any problem with this idea of multiplying the Q4 amount by four and using it as a starting point to project next year's figures?

A4:

(Thornton) I think we will be reporting Q4 results on May 8th. Please give us a little time to, as we get closer to May 8th, we may have a slightly better view, and we will guide future quarters when we get to May 8th.

Speaker 2

Q1:

Hi, congratulations. Very good presentation. But more importantly, great to see Arm delivering. In fact, I think more than what you suggested in 2019. If people had listened carefully, I think it would have been a very good result. I have a question on the new AI product line. I think before you explained to us one of the great things about Arm was the ability to work with industrial partners to see the future of technology. My imagination is that we are in uncharted territories in terms of the number of companies who now want to use chips. But I do not know if that is really true. Just one thing is, are we seeing a real broadening of the customer base. Related to that, when would we expect to see these kind of real AI products really coming out from those customers? As you said, it seems to me really in the last 12 months, people are getting excited, and CEOs and CTOs all have to join the boom now. If you do not join, you are left behind. Should

we be thinking two or three years is when you really see the output? The final point on that is the right way of thinking about this is the biggest impact would be more complexity per chip. If you could comment on those areas, thank you.

A1:

(Thornton) In terms of expansion in the customer base, I agree with your analysis that we are seeing more companies wanting to build their own computer chips. I think it is probably surprising to many that companies like Amazon, Meta and Tesla want to build their own chips. For years, I thought that building a chip was getting harder and harder, and requiring more and more resources and therefore would become increasingly consolidated around larger and larger companies. But it seems that there are more companies wanting to build chips today. Although actually many of those companies are extremely well resourced and therefore, can afford the very high cost of chip development. I think what we are seeing here is that software is becoming more of the product that we as consumers are purchasing. When we used to buy a car, it used to be based on the quality of the car, how fast it went, whether the leather seats were very nice or not. But now increasingly what is selling the car is the user interface. How does it look and feel in terms of the screen and how I interact with the data. What software functionality does it provide in terms of self-driving capability or lane warning signs and things like that. Those are all software-controlled functions. Because all of that software runs on a CPU, runs on a chip somewhere in the car, more companies are wanting to take control of the chip because ultimately that is what differentiates how their software works versus the competitor. If they can make that software be smarter, faster, more capable than their competitors, then maybe they will sell more cars or more services. I think that is a trend we have seen across many markets, including things like cloud computing. I think we can expect that to continue, that more non-traditional semiconductor companies will want to build chips. One of the reasons for developing our compute subsystems as I mentioned earlier is because these companies have not been building chips for 30 years. They may want a better starting point than just the individual components, and having those components pre-assembled into a subsystem gives them a better starting point. One of our first licensees of our compute subsystem said that they had gone from delivery to basically taping out their first chip, taping out basically means sending it to manufacture, in just nine months. This was for a complex server chip that would normally take 18 months to two years. The design time was more than halved by using the subsystem. Another customer said that they had saved between $20 million to $30 million worth of engineering effort by using this subsystem. The engineering effort implies to the design effort that they would have to pay to their own engineers without the subsystem. We are definitely seeing that subsystems are good use to these non-traditional semi-conductor companies. Regarding to the question, are we going to start to see real AI technology in the next two to three years, I think one of the most asked questions when I go to an investor conference has been what is the killer app going to be for AI at the edge devices? What is going to make me want to go out and buy a new smartphone because it is an AI smartphone? What is that going to mean to me? The analogy that I have been using is that this feels a bit like the early days

of 4G when the 4G networks were being rolled out and I was being asked what is the killer app for 4G? 4G had some basic functionality, such as streaming a video, or downloading an attachment to an email a bit faster. But what was the killer app for 4G? Looking back, I would suggest that the killer app for 4G was Uber or more generally ride services, because Uber cannot work without a 4G smartphone as part of their infrastructure. Equally, you as a consumer cannot access Uber without a 4G smartphone, so the two are needed together. But Uber is not a 4G app. Uber includes 4G, sure, includes a smartphone, but it also has lots of infrastructure. It has got cars with people to drive them. I do not know how many meetings like this I could have sat in before someone said, I have got it, I now know what the killer app for 4G is, it is taxi services. Sitting here now looking at the smartphone and saying what is the killer app for AI? Yes, I can do live translate. Yes, I can do circle to search, but so what? I think what we are really seeing right now is AI PCs, AI smartphones, AI enabled cameras, AI cars, going into the AI capability being provided to the developers as an empty box as something to say go build your application, go build your product, go build your service of which the smartphone or the smart PC may be just a small component of a much bigger infrastructure. I am afraid if I could invent that I would not be here, I would be with the investment bankers taking on large amounts of debt and building a big business so I do not know but I think maybe in two to three years' time we will start to see some of those new businesses starting to appear and maybe then it will appear obvious to us.

Q2:

Just as a final point, for those new chip customers you mentioned that the number of CPUs is just massively higher so we should expect the complexity per chip from that new customer base to be very positive for your revenues.

A2:

(Thornton) One of the things I would definitely point out is that if we look at AI through the lens of its software, it is just another way of writing software using statistical analysis rather than traditional programming techniques. The AI algorithm is very computational heavy. It needs a large CPU, or it needs lots of CPUs to run. Therefore, you can expect that digital electronics that become AI enabled will need more powerful CPUs, at least initially. It will be interesting to see how the models evolve, ChatGPT-3 had 170 billion parameters ChatGPT-4 has one hundred trillion parameters, so a 600-fold increase. We therefore may expect to see that even AI in a smartphone, maybe will not increase 600 times but as new capability is added it will need more performance. But then we are also seeing models become simplified. As models start to become a bit more fixed, then work is done to optimize them to reduce the number of parameters so they fit into smaller memory and need less performance. I showed earlier a chatbot that is running in a smartphone quite happily. We are working with a company on a text to image generator that runs in a smartphone. It is a bit slow today, but once we have done some more optimization then maybe that can run quite happily as well in a smartphone. We will see both expansion and complexity in compute, and then once the models start to settle down, more optimization.

Speaker 3

Q1:

I would like to ask two things. One is the relationship with SoftBank Group. I believe that SoftBank Group will develop various investment strategies, etc., taking advantage of its position of holding 90% of Arm shares. On the other hand, from Arm's point of view, the IPO had the advantage and benefit of not having to make a profit and being able to concentrate on investment, which has led to the current growth of the company. Now that the SoftBank Group is a major shareholder, is there any good that can be expected from it? Or I would like to know what we can expect from Softbank Group in the future.

A1:

(Thornton) Just looking back over the history, when Arm was acquired in September 2016 I remember the meeting with Masa when he addressed the Arm workforce and his primary message was go, go, go. So, we went, to your point, we significantly increased investment in R&D. That was very necessary. I think the Arm portfolio that we had in 2015 was stretched, and we needed to change things. Those first three years post the acquisition was spent in changing the product portfolio from being just one CPU family into being CPUs designed for our four main markets. We have CPUs for mobile, CPUs for infrastructure, CPUs for automotive, and CPUs for IoT and embedded devices. That was something that we would have struggled to do as a listed company. During that period, we took our operating margin down from about 50% to about 20%. I think that if we had done that as a listed company, the CEO and the CFO would have been fired by the investors. Therefore, we needed the support of SoftBank. Since then, we have been very much focusing on increasing profitability by selling our new products, and then collecting the royalties as they are starting to appear, whilst also then developing our new Armv9 family of processors. At that point, that was when we could start looking ahead to an IPO. Then we had the pandemic, then we had NVIDIA trying to acquire Arm and only then could we actually get on with the IPO itself. The IPO was probably a little bit delayed, but now our focus is very much on getting the balance correct between investing in new technology and allowing the profitability to come through. We see lots of opportunity right now for new technology investments, with all of the new opportunities from AI, from the new opportunities that our compute subsystems bring, in terms of helping companies build chips more quickly. We have plenty to spend our money on. We have hired a thousand engineers last year. We intend to hire another thousand engineers this year, and probably another thousand engineers next year as well. Nevertheless, we still think that revenues can grow faster than our costs because of what we have developed over the last few years.

Q2:

I think it is still too early to pay dividends, and I don't think your company's major shareholder, Softbank Group, is looking for dividends, but with the current momentum, I think you will be very profitable in the future. I would like to know what factors, if any, would change your view on

dividends, and under what circumstances would it be the right time to consider a dividend policy.

A2:

(Thornton) Yes, currently Arm has around $2.5 billion in the bank. Last quarter we generated $250 million of cash. If we keep that up for a year that is $1 billion every year. You are quite correct that we have very little to spend our money on. We could do a buyback, but we have only just IPO'd, so buying back the shares after issuing some does not seem very sensible, and also, I think our investors want more liquidity not less. So, a buyback does not really make sense. Softbank could have done a dividend before the IPO but chose not to. So, I think that right now we are not being asked by a major shareholder for a dividend and I think we would need to wait for their request for a dividend before we would do one. We see no need to do one ourselves without permission from Softbank. They own 90% of our shares, so we need to take their lead. We will probably do some M&A. Arm has historically acquired companies from time to time. Usually, we acquire companies as part of our recruitment strategy. If we have to hire 1,000 engineers, maybe we could acquire a company with 100 already working together, so there are senior managers, line managers, junior engineers, already in place. That can be quite an efficient way of hiring. And there may be some interesting technology areas, clearly with AI being a big focus and system design being a big focus. There may be either companies or teams within a company that might have some technology that we could turn into IP that we could then license to our customers. Not many companies have the sort of technology that turns into generally available IP, so we have to be very careful and very selective, but there are sometimes opportunities to acquire companies and we will save our cash for that sort of opportunity.

Speaker 4

Q1:

I have two questions, one that is top down and another one that is bottom up. The first one is you have the total addressable market slide in your presentation. It seems to me that the cloud compute expectations or forecasts are well behind what your downstream customers are doing, and so is that going to change soon? I know you talked about with the IPO that you have to be careful about how you are guiding, but that slide probably needs an update.

A1:

(Thornton) You are absolutely spot on there. We developed this as part of the IPO process in March to May last year, so really ahead of a lot of the excitement, ahead of AMD forecasting $400 billion worth of AI infrastructure chip sales. Not all of that $400 billion is CPU based, much of that will be GPU based, but every GPU needs a CPU. I think it is fair to say that the $28 billion that we are showing here for the TAM in cloud compute probably would be a larger number. But we did the analysis before knowing what we know now. We are not yet going to be updating it today. We have the numbers that we use for the IPO and need to stick with them. But for May 8th, which is our full year results, we are looking to see whether to update our longer-term forecasts and

actually I am looking at putting in a number for 2030 to give some long-term targets to chase after. That conversation is still happening internally, but that is my plan.

Q2:
My other question is related to your customer channels, and it is more of a confirmation I suppose because I do not think you can give us a lot of detail. When you look at companies like NVIDIA, versus companies like Apple or Microsoft, are the latter economically more rewarding for Arm because of the software sales or because of the integration package that you sell them like Neoverse and things like that?

A2:
(Thornton) Different customers obviously license different amounts of technology. Someone like NVIDIA uses a lot of our technology in a lot of their chips and have done for many years. Their chips in automotive have Arm based, their latest AI chips are a combination of an Arm CPU and their own GPU. We are in a lot of their products which obviously help us to drive revenues. I think you can see they are our main customer. You can see pretty much all the companies you just spoke about are our top 20 customers, so they all pay us a fair bit of money. Maybe I am not answering your question right.

Q3:
If you look at Microsoft making chips on their own, and what you sell them in terms of licenses, in terms of the royalty units you get for the chips they make, and in terms of the support they get for the integration of all the things on the chip. Is there a clear economic, maybe make it more simple, is the revenue you get for each chip higher with a customer like Microsoft versus NVIDIA?

A3:
(Thornton) Let us take it a step back from individual names, but if you use our latest Armv9 technology you pay a higher royalty rate than for v8. If you use our subsystems, you pay a higher royalty rate than if you just use v9 on its own. A chip like Cobalt 100 which uses our compute subsystem will deliver a higher royalty rate per chip than one that was v9 based but not using our compute subsystem. That is not because they are Microsoft versus NVIDIA, it is just because if you use more of our technology in your chip, you pay more.

Speaker 5
Q1:
Thank you for this presentation. This is more or less a follow up question to the question just asked. With regards to your compute subsystem solutions, you currently deploy those for server CPUs, smartphones, automotive. My question is, in which of these segments do you expect penetration of the subsystem-based solutions to be highest, say five years from now? And is there an upper limit to the penetration you can achieve in smartphones? Your customers in

smartphones are very experienced in developing their own SoCs, but do you see a limit to how high your penetration can go with your subsystem solution in that particular market? Thank you.

A1:

(Thornton) There is an opportunity for subsystems to gain share even within smartphones, even within companies that are very experienced at building chips. That is because the smartphone market is maybe not unique, but every single year a smartphone OEM needs a new flagship smartphone to come out. Every single year, they need a new, increasingly complex high-end chip to go into that flagship smartphone, and that chip must be significantly more advanced than the one they had a year ago. So, the semiconductor companies are required to deliver a new and better smartphone chip every single year. And they have done that for many years, except that now, the time it takes to manufacture an advanced chip is taking longer and longer. For TSMC at five nanometers, it took 16 weeks for a chip to be manufactured. For three nanometers, it is taking 20 weeks. The amount of time that you have to design a more complicated chip has been reduced by one month. You used to have nine months to design the chip, three months to manufacture a new chip. You now have eight months to design the chip, and it is a more complex chip. We think that when we go to two nanometers, then it is going to take another bite out of your remaining design time. So, even for smartphone companies, having a better starting point than the components may therefore be able to bring you a significant benefit. We are working very closely with some of our large smartphone customers in order to try and make sure that our compute subsystems for mobile are going to be able to enable them to continue to hit that annual beat that they have to hit every single year after every single year. The more time taken in the fab, effectively the more valuable and the more useful our subsystems become. So, there is potentially a big opportunity there. Outside of the smartphone market, we have multiple design wins. Now with our Neoverse compute subsystem for cloud compute, we have four licenses already. Microsoft being the only one that is public therefore the only one I can talk about, but there are others. And then earlier this week, we announced our automotive compute subsystem that is targeting in-vehicle infotainment and ADAS chips going into cars. It is not available yet, it will be available next year, available to be delivered to our customers next year. Hopefully that will then start to appear in cars in three to four years' time. And again, that is targeting both traditional semiconductor companies who just need to build chips faster, and also non-traditional semiconductor companies like car OEMs.

Q2:

Thanks Ian, that is great context. As an unrelated follow-up, given the proliferation of AI on smartphones and in Edge in general, do you see a possibility for v9 to be adopted at a faster pace than v8?

A2:

(Thornton) Going back to when we introduced v8 about ten years ago, the v7 to v8 transition took

about four years in the smartphone market. v8 brought a very important innovation, which was the ability to run PC applications in your smartphone. Things like Excel, things like PowerPoint, could not run easily within a smartphone before but with the introduction of v8 you enabled PC applications to migrate across into the smartphone market. V9 brings additional big benefits, such as accelerating AI, so I think that they both brought big benefits and I think maybe it is less to do with the attractiveness of the technology as much as it is to do with the ability for semiconductor companies to roll out a new technology across all their product portfolio, and then for OEMs to do the same as well. I do not think that we will see a significant increase in the deployment of v9 versus v8, simply because I do not think they could go much faster, if that makes sense. In terms of where we are today, if we assume that it takes about four years, which is what happened last time, we are about one year in right now and so we still have about three more years to go before nearly all smartphones become v9.

Speaker 6

Q1:

Could you tell me what is your biggest risk in the five-year outlook that you have made? People talk about RISC chips, RISC architecture as being an alternative solution perhaps to Arm, do you perceive that as a threat at any time, or do you think that your install base advantage is so strong that that is not really a threat to your position? Could you tell us what concerns you, concerns the other management most in the outlook that you give? Because you give many positives which sound very persuasive, but what are the negatives that you are most scared of besides the end customer disappointing.

A1:

(Thornton) I guess there are two parts there, there is what do we find is the biggest competitive or technological risk, and then you asked a very specific question about RISC-V, the extent to which RISC-V is a competitive threat. Those are slightly different things. In terms of technological risk, we are in a period of rapid change, and clearly that brings both opportunities and threats. We know that AI algorithms are going to be introduced into a wide range of embedded markets, as we discussed from smart phones, and PCs, smart TVs will be using AI, smart cameras are already using AI, maybe even your washing machine will use AI to wash your clothes cleaner using less detergent or something. There is an opportunity for this new technology to be deployed across a very wide range of electronic devices. But as I indicated earlier, the algorithms, the AI algorithms themselves are changing very rapidly. We are only just starting to see the first deployments of AI into edge devices. It is highly likely that in five years' time the software that will need to be executed will be different, and in ten years' time it will be different again. We need to identify how that software is changing, and therefore making sure that we are building the right combination of technologies so that Arm is able to provide as much of the solution as possible. What we do not wish to happen is for Arm to repeatedly build the wrong thing, and therefore create an opening for a competitor to come along and build the right thing, and to partially displace Arm or at least,

have the value that this new opportunity brings accruing to them rather than accruing to us. Now, I think that we are very well placed in order to identify, compared to anybody else on the planet, to identify how that software ecosystem is going to evolve. We have the largest software ecosystem today. We have deeply embedded relationships with Microsoft, with Google on the Android side, with Apple of course, with iOS. We develop a lot of technology with the Linux and open source software community, so we have the best sensing organization for any technology on the planet pretty much. But that does not mean that there is not a lot of hard work required, and a lot of decisions to be made in terms of how we go to market and with what technologies. Therefore, there is always a risk of being wrong and somebody else being right. That is probably the biggest technology threat and the thing that makes us most worried most of the time. And then, specifically on RISC-V, although RISC-V is an alternative technology to Arm in some markets, RISC-V is actually a different technology that is trying to solve a different problem to Arm. RISC-V is a processor architecture similar to Arm, but the RISC-V architecture is modular and scalable, allowing people to make lots of changes to the instruction set. If you remember, the instruction set is that relationship between the processor and software. With RISC-V, for each RISC-V design you can change the instruction set, and that means that if you know what software you want to run you can create an instruction set that is very optimized for that particular software algorithm. The problem is that that processor will then not run on another piece of software, which would need a different design. At Arm, we define the architecture, and we fix it, so software runs the same across all Arm processors. Whereas with RISC-V each one is unique, and so software developers need to optimize their software for each individual RISC-V implementation. We tend not to find RISC-V in markets where large amounts of third-party software are needed, because it is just a lot of cost to support multiple different RISC-V implementations. Arm is a better choice for something like a smartphone, which wants to run Android across multiple different companies' chips and apps across multiple different companies' chips. RISC-V tends to be used in more deeply embedded applications like in a Bluetooth protocol stack chip, some kind of wireless connectivity chip, that is where we see more RISC-V. So, there is some overlap, but it is 10% of our business, not 90% of our business.