Goldman Sachs Talks
Nigel Toon, CEO, Graphcore
Clif Marriott, Moderator
Recorded: May 14, 2024

Nigel Toon: Today, we make chips with a 100 billion transistors. So if your car had improved by the same amount, you would travel everywhere at 200 times the speed of light.

Clif Marriott: Thank you, everybody, for being here. My name is Clif Marriott. I'm co-head of our Telecom, Media, and Technology Group within Investment Banking in Europe. Very pleased to have Nigel Toon here. Nigel is an experienced entrepreneur, engineer, and leader in technology. Prior to founding Graphcore, he led two successful VC-backed processor companies and is widely recognized as being a leading innovator in his field.

Alongside his work, he's a member of the UK Research and Innovation Board. And if that wasn't enough, he recently authored a book, which I hope all of you have read prior to this talk. I definitely have. We'll get into portions of that book. *How AI Thinks: How We Built It, How It Can Help Us, and How We Can Control It.* Questions that are at the top of I think everyone's mind right now, given what's happened in particular over the last couple years.

It's a really exciting time in AI. We're really pleased to have you here, Nigel. Thank you for joining us.

Nigel Toon: Thank you so much.

Clif Marriott: Maybe we can start, just for the people in the audience and then people listening in who may not know you or know your background, can you go through a little bit of your background? What led to the founding of Graphcore and where you started before that?

Nigel Toon: Yeah, sure. So I'm an engineer. And then more recently an entrepreneur. I spent 14 years at a company called Altera. So I joined them just before they went public. We did a 256x in the public market, which is pretty good. It became worth 20 billion. Ended up being bought by Intel, and now Intel's I guess spinning it out at some point. Seems to be the path that Altera is on. So that was a really interesting journey, going through that. I joined, there was about 100 people in the company. And by the time I left, there was two and a half thousand and, you know, billion and a half of revenues, etc.

So going through that whole growth experience, you know, was just amazing. And figured why can't we do this in the UK? So together with a couple of other people, we started a company called Icera that built a baseband processor for cellular data on cellular devices. This was before the iPhone. We started in 2002. IPhone came out in 2007.

Clif Marriott: '07, yep.

Nigel Toon: And that was basically our premise, that devices would come out that would be, you know, mobile and data rich and that they needed a modem to go with that. And so that worked out. We sold that company to Nvidia. In today's money, that would be $100 billion. Obviously, it was a bit less then because Nvidia wasn't quite worth as much as they are today. But I would call that 100 billion. That's pretty good.

Clif Marriott: I hope you held onto those shares.

Nigel Toon: And then helped VCs on a few different companies, process companies. And then we started Graphcore back in 2016, basically having recognized from about 2012 that this whole AI thing was going to be a thing. It would need new types of processors. Nvidia was still pretty small at the time. We thought we could compete with them. And, you know, we're probably one of the few companies that really is because it's not just about the hardware. It's actually also very much about the software as well.

Clif Marriott: So I had a couple questions on chips, but I just think, given what you talked about in terms of GPT 4.0 being released in the last 24 hours, maybe we start more high level about AI overall. We're clearly in this interesting moment. I mean, reading your book feels like you put a very positive spin on AI. I think there's probably some nervousness also in the room around AI, so I think we'll go through those themes.

I just wanted to read this because I thought it was very interesting. It was very early in your book, but you said, "I've come up with a set of rules that describe our reaction to technologies. Anything that is in the world when you're born is normal and ordinary and is just a natural part of life. Anything that's invented when you are between 15 and 35 is new and exciting and revolutionary, and you can probably get a career in it. And then anything invented after you're 35," which unfortunately is a bucket I sit in, "is against the natural order of things and dangerous and scary."

Nigel Toon: Wait, so that is not my quote. That is Douglass Adams, *Hitchhikers' Guide to the Galaxy* author.

Clif Marriott: You're right.

Nigel Toon: But I think it's very prescient in terms of how people think about technology. We tend to be scared of it, you know? It's interesting. One of the things I was -- I've been giving thoughts about the book, and one of the things I always open with is this year 150 million intelligent machines are going to be delivered. And they're going to end up being as intelligent as you, in many cases more intelligent than you. They're going to end up running companies. They're going to advise governments and will eventually end up running governments, these intelligent machines. They will beat all of our existing sporting champions in every sporting category. And when you get old, you need to be sure that these intelligent machines are going to look after you because they're going to be in charge. We call them children.

So, you know, we're already -- you know, we become obsolete. You and I will be obsolete perhaps before some of the other people in this audience. Our children will take over. We are an intelligent machine, you know? And I think the other thing which I find fascinating, having kind of dug into this, is -- I don't want to pick on anybody in particular -- but you share nearly 50% of your DNA with a cabbage. Actually, we all do. And we're all kind of evolved from this single microbial cell that split and created every living thing on this planet. And so these are all intelligent machines that we're surrounded by.

You know, plants are intelligent. You know, bats are intelligent. Bats are actually pretty close to us in the evolutionary spectrum. And interestingly, AI is helping us to understand more about them, and they actually have names for each other, have unique call signs that they use. So, you know, they actually call each other by a unique name.

So all these things we think of as being particularly human, you know, actually we share a lot of it with all of these devices. And now we're developing this technology which is sharing other aspects of what we do. And I guess the point of my book is this is an opportunity because it's not going to replace us. It doesn't have to replace us. It's a tool that we can use to make our lives better, to make us more intelligent.

Clif Marriott: Yeah. I mean, I'm going to skip around a lot here as we go through this, but, I mean, there was another part of your book that I found really interesting and refreshing or made me feel more comfortable was I think you talked about in the US in 1950 there was a list of occupations. And if you look at which occupations have actually been replaced by technology, only one had been replaced, which was a lift operator.

Nigel Toon: Elevator operator, yeah, yeah.

Clif Marriott: Elevator operator in the US. So that made me feel pretty good. That, you know, after whatever, 60, 70 years of automation technology, only one job has actually been replaced. But --

Nigel Toon: Completely replaced, yeah, yeah.

Clif Marriott: Completely replaced, exactly. Many others have been complemented and everything else. But you went very thoroughly through your book on all of the different aspects of society that can be improved by AI in a very powerful way, in a very positive way. Maybe you can pick one or two where you think the change will be profound and complementary to what we're all doing in the audience here.

Nigel Toon: Well, you know, obviously, I'm getting a bit older, and so you start to worry about your health maybe? I don't worry too much about my health, but I guess you should at some point. And we actually don't know much about cells. We actually don't know much about molecules and atoms. We actually don't know much about proteins. You know, in spite of all of the research we've done, our knowledge is pretty superficial at this point.

And one of the things that AI is already starting to do, for example, in terms of proteins, you know, work that DeepMind has done, has allowed us to model every protein in the human body. And now we're sort of working on every protein in every leaving thing.

Now, the models are not yet very complete, and there's more work that needs to be done. But we're starting to build a picture of what these proteins look like. And why is that important? Because a protein will connect to a cell. And a protein you can attach a molecule to that might have some medicinal impact on the cell. And so if you can work out which protein is going to attach to which cell and which molecule is going to interact with the disease in the cell, you can actually cure diseases, you know, like cancer. No side effects.

And so think of it, you know, the other work that DeepMind did on playing computer games, where you understand the environment and you can start to learn more about the environment and you can actually have an impact. So just like a computer game, you can work out, okay, which are the proteins, which are the molecules, how do they interact with the cells to make us better and to cure diseases?

And that is going to -- and the next ten years, those types of breakthroughs are going to happen. There are scientific breakthroughs that we're going to achieve that are just out of reach at the moment. And that I think is the power of it, you know? So I think that's an amazing area.

We're all obviously worried about our planet and all this carbon we're throwing into the atmosphere. And we can't just halve the amount of carbon because actually that still means the carbon is going to keep increasing in the atmosphere, so we've actually got to get rid of it all. So everything has to go green. All of our energy needs to go green. And the only way we're really going to achieve that is probably with nuclear. People might not like the idea of that, but that's probably the only way we're going to do it. And the only way real way we're going to do that when nuclear in a safe way is probably with fusion, which has been one of those projects which is always 20 years away.

Clif Marriott: Yeah.

Nigel Toon: But AI is actually starting to make an impact in understanding how these fusion experiments work, how we can control them, how you can control the plasma.

Clif Marriott: When do you think we have fusion energy?

Nigel Toon: I don't want to say 20 years, right? Within 20 years. I think, you know, the experiments that we're currently doing have already been able to create more energy than we put in, which was the first sort of major breakthrough.

Clif Marriott: That was recent. That was a year ago.

Nigel Toon: That was, like, two years ago, yeah. And now there's a big project in France that are working on actually starting to make this stuff more commercially viable. Still is an experiment. So maybe that's another five to ten years of work. But maybe after that we then have a basis on which we can actually start to roll this stuff out. And it's going to be AI that will help us understand that really complex problem and how do we actually solve it.

And this is -- you know, so I like you said, I sit on the board of UKRI, which is the UK Research and Innovation. We're doing all the funding into universities and the research councils and all this piece. And the way in which we'll use compute to solve some of these scientific problems is going to be incredible.

Clif Marriott: So we're going to cure diseases. We're going to have unlimited waste-free energy. What about finance, which also is maybe not as important but for the people in the room pretty important? What do you see happening to finance? Will we actually invest or will robots invest for us? How does this work?

Nigel Toon: Well, that's an interesting one, isn't it? So here in London, we have cab drivers, taxi drivers, you know, who go through this incredibly long process of the knowledge to learn all the streets around London. You know, incredible skill. And yet almost obsolete now because you have satnav.

But the thing that's interesting is it was really what Uber did in terms of combining satnav with locations of customers and getting people to interact in a system that actually reshaped how taxi systems could work. So it's really a system design. And if you go back to the start of the Industrial Revolution, at the start of the Industrial Revolution, the mill owners actually put machines in to replace the tasks of all skilled workers and use the machine plus an unskilled person to replace the skilled worker. And that's probably the fear of maybe, you know, many people in your organization thinking, you know, "A machine is going to replace my job," and an unskilled person could become a trader.

Now, what happened in the second phase of the Industrial Revolution with electricity is we actually got mass production. We put electric machines in line to actually lay out factories, and then we completely reshaped how factories worked. Actually, the wages of staff went up. The distribution of wealth was much better distributed.

And so when you actually change the system, the way in which you do things, you use technology to maybe change the way in which you work. Rather than just automating what you're currently doing, that's when the big wins happen. That's when the big gains happen.

So what I would say is your business is probably going to look very different from what it does today. If you just go and automate the tasks that you're currently doing, you'll probably lose. But if you can work out, "How can I use this technology to reshape what we do and reshape how we do it?" then that's where the big win. And that will be a combination of human intelligence and machines working together to create new systems and new ways of working.

Clif Marriott: I think you actually said you see AI actually like electricity.

Nigel Toon: Yeah.

Clif Marriott: And that ties into another thing from your book in terms of people tend to overestimate technological changes early on and underestimate technological changes long term. Electricity was probably the same thing, right?

Nigel Toon: Yeah.

Clif Marriott: The invention of electricity started first and was most effective in factories, but then impacted people and their homes much later.

Nigel Toon: Put it in a context that's maybe very relevant for people here. In electricity, initially the value went to the producers of the electricity, the Edisons and the Westinghouse. You know, they accrued the value. And then over time, it was all the applications that electricity got used for where the value actual created.

When the personal computer came out, it was Compaq and Apple that accrued the value for the personal computer. And then it's all Microsoft and all the applications that now capture all the value.

And where are we today in the AI cycle? Well, yes, Nvidia is capturing the value today. They're worth two trillion today, but where is the value going to be in five, ten years' time? It's going to be in the hands of all the applications that get developed on top of this infrastructure that we're building. So it's just another infrastructure build-out. It's probably much bigger than people realize. The amount of money that's going to be invested in this infrastructure is enormous. You know, there's going to be other layers have value that get created on top that are even bigger and are very different from the current ways that value is created today.

Clif Marriott: You were talked about the impact to different sectors, to jobs. So what should one's kids study or practice in order to be prepared for the world that's coming?

Nigel Toon: I'm not sure it changes dramatically. I have a son who's just in his final year at Oxford University doing engineering. And, you know, he's very good at maths. And he's going to go off and work for a hedge fund doing quantitative trading, using -- you know, he doesn't want to go and be an engineer. I'm, you know, slightly disappointed in that. He wants to go off and use his math skills to do this. So it's like the fundamentals. Maths, sciences, actually I think arts and humanities are going to be very important in this world.

Because it's really it's about using those skills to apply them because a computer, you won't need a program a computer anymore. We're already getting to the point where we're starting to stop programming computers. Most people don't program them. You tell it what to do through a prompt, and it creates some new output for you. And that's going to get more and more complete.

So think of, like, creating a computer game. You have a framework where you start off and you say, "Okay, so I want some characters, and so I'll give it some prompts to generate some, you know, give me, you know, horrible tortoises or something like that," and so you get your characters. And then, you know, some other framework that creates the background for you. And then some other sort of storytelling pieces that create the gaming environment for you. And within a few steps, you will have a computer game that you might enjoy. And then you can throw it out into the world, and somebody else can come along and change it. You're not having to program that anymore.

And the same will be true in what I talked about in terms of biosciences and drug discovery. The scientists won't have to learn how to program the computer. They'll have frameworks that will allow them to use the computers to solve the problems that they have. That's the potential here.

You know, your traders will have a framework in which they can use their analytical skills and their understanding of the markets to get -- you know, rather than having to go to the quants necessarily to always come up with the models, they can actually create some of the models themselves and use their creativity to create some of those pieces. And that's the potential here that we have.

But let's not get too far over our skis because this is a really powerful technology as well. In the same way that there's lots of good stuff we can do with it, you know, bad actors can do some bad stuff with it. But that's no different from any powerful technology we've had before. You know, a gun is a powerful piece of technology and, in the wrong hands, can do some terrible harm.

Clif Marriott: That's right.

Nigel Toon: AI will be the same.

Clif Marriott: We're going to get to that in a second and definitely about how to regulate it, governments' impact, etc. So let's go back to just more to the technology behind this, which obviously is more where you're at. You're providing the pick access for this new world. I mean, a lot in your book around the evolution of the chip set, the fact that Moore's Law is slowing on silicon. Maybe put some numbers in people's minds around where we've come from, where we've gotten to on silicon. And then what are the limitations from here in terms of further progressing? You talked about a couple different alternatives, either molecular or quantum or other options.

Nigel Toon: So 1947, first transistor. 1960, the first integrated circuit that actually encapsulate multiple transistors in a chip, so that took quite a long time. First integrated circuit came from Fairchild. You could argue Texas Instruments similar time. First Fairchild chip had four transistors in it.

Today, we make chips with 100 billion transistors. So that was 1960 to today, you know, 60 odd years later. 25-billion-fold improvement. So if your car had improved by the same amount, you would travel everywhere at 200 times the speed of light. So we just do not understand or even realize the level of improvement that has happened in semiconductors and how pervasive they are in everything we do.

You know, obviously it's in our mobile phones and in our laptops and in our computer screens. It's in your wallet, in your credit card. It's everywhere. And our life basically runs on chips, but they're sort of hidden away. And people aren't aware of the incredible impact. So that has allowed the platform on which AI has become possible.

The second big thing is information. So same time as the transistor, a year later, 1948, a guy called Claude Shannon, working at Bell Labs, came up with this idea of a binary digit which we call the bit. And the idea that all information can be held as combinations of ones and zeros. You know, switch on or off.

And the real insight he had was that a transmitter and a receiver can exchange this information. And what you're trying to do is to reduce uncertainty at the receiver of what the transmitter has said. And he came up with a bunch of maths that describe that whole process, and it's why we have the Internet. It's why we have WiFi. It's why we have mobile phones. All the communication technology we have is because of the work of Claude Shannon.

So now we've got compute. We've got information. And then actually in about 1980s, there's a bunch of people who came up with methods for how we can create AI in a sort of artificial neural network style, you know, mimicking a little bit how our brain worked. Didn't work. We didn't have enough information. We didn't have enough compute.

And the breakthrough was 2012, Alex Krizhevsky, Ilya Sutskever -- Ilya is now the founder of OpenAI -- they came up with this AlexNet model that was the first deep learning AI model that would recognize an object and an image better than any other type of computing approach. And within a year or so, it was better than a human on images that it had been trained on.

So literally, since back to the 1950s, we've been trying to build artificial intelligence, and it's these breakthroughs in the Internet and semiconductors and these methods that now around about 2012 kind of kicked off the deep learning process and then kind of hits everybody's public consciousness in 2022 --

Clif Marriott: With ChatGPT.

Nigel Toon: -- with ChatGPT. Interestingly, though, in China, it was --

Clif Marriott: AlphaGo.

Nigel Toon: -- AlphaGo, 2016. 240 million people in China watched that live on television.

Clif Marriott: It was just a headline here. I don't think anybody [UNINTEL]. So maybe going back to what's next. So Moore's Law, massive growth in terms of [UNINTEL].

Nigel Toon: We all know hopefully that things follow an exponential curve, but they don't actually keep going. They follow an S curve. They slow down over time. So things improve, improve, improve, and then you can't improve anymore. There's limits you reach, and it starts to flatten off.

We're in the S curve. We're at the top of the S with semiconductors. So it will still improve over the next few years, but much more slowly, take much more effort to improve. You look at the improvements in compute that have happened in GPUs, for example. Most of that is not because of Moore's Law. Most of that is because we've gone from 64-bit floated point down to 32 then 16 then 8 and now 4 with the latest Blackwell chips. And that gives us much more compute for the same energy at lower precision, but AI works with low precision. So that works for us.

That's the main reason we've got more compute. Not because of Moore's Law, in spite of what Jenson will tell you. We're starting to run out of steam on chips. And so then you look beyond, okay, so what comes next? Well, everybody's talking about quantum. It is not yet certain that quantum will work. We have examples of quantum computers that do things and are interesting, but it's not clear we can scale them.

I don't know the answer, whether we will be able to or not. It's an interesting area of research, but it's an area of research at the moment. Maybe AI will help us work out how to do it, but it's very much about understanding the materials that we build these things in.

Clif Marriott: Can you spend a minute just to explain what is quantum computing for people?

Nigel Toon: So if you think of a binary digit is a one or a zero, a quantum state is like a spinning coin that is always spinning and could be a one or a zero at any time. And what quantum computers try to do is they try to get these quantum bits, these Qubits, to spin together.

Clif Marriott: These are sub molecular?

Nigel Toon: Yeah, the -- you know, it could be an ion state. It could be a spin transistor. It could be a light source. Everybody's trying to work out what the right way to build these Qubits is. The thing that's interesting about Qubits is you double the number of Qubits and you actually grow on an exponential law the scale of compute. So with a few million Qubits rather than hundreds of billions of transistors, you could build a really powerful computer.

The problem is, when you try -- it's the Schrodinger's cat problem -- when you open the box and look inside to try and work out which way this Qbit is spinning, often you will destroy the experiment. You will create noise that causes the experiment to fail. And so you can kind of get it to work with a few Qubits -- built computers with a few hundred Qubits so far. But to get to a million qubits, it's not clear how you scale it up with all this noise that doesn't the noise just end up, you know, killing the way the machine works?

So it's kind of unproven at this point. People are working on it because the potential is enormous. You know, you end up with a computer that just, like, you know, a maze. It doesn't just try and work out how to do the maze, it just follows all the paths all at the same time instantly --

Clif Marriott: Instantly.

Nigel Toon: -- and will tell you which is the way through the maze. You know, playing a quantum computer at chess won't be much fun because it will just work out every move in advance and will know how to win.

Clif Marriott: So quantum is one way and molecular is the other? Or?

Nigel Toon: So quantum -- I'm -- it's unclear to me whether quantum is going to work. I think the thing that's really interesting is molecular because that's how we build us and that's how plants work. And the difference is actually transistors and silicon actually uses a lot of power, as we know. We're ending up with megawatts of power required to power these AI computers.

Your brain is about an exaflop and about 25 watts. So the same as a laptop. An exaflop of compute, which in a computer today, would probably be about 20 megawatts. So your brain is about a million times more efficient than a silicon brain. Not that we've really been able to build a silicon brain yet. But to the extent we could, your brain is about a million times more efficient.

And so if we could use molecular approaches, potentially we could build much more powerful computers much more energy efficiently. And as we start to understand more about how molecules work and atoms work and cells work, potentially that is a way in which we could build computers in the future that would be small, incredibly powerful, and we could get much more performance from. That's at least 20 years away.

Clif Marriott: Well, Nigel, thank you very much for the conversation.

Nigel Toon: Thank you.

Clif Marriott: Thanks for joining us.

Nigel Toon: Thank you very much. Thank you, everybody.