MIT CSAIL Alliances | Peter Shor Will Oliver Podcast Export 4

Welcome to MIT's Computer Science and Artificial Intelligence Labs Alliances podcast series. My name is Steve Lewis. I'm the assistant director of Global Strategic Alliances for CSAIL at MIT. In this podcast series, I will interview principal researchers at CSAIL to discover what they're working on and how it will impact society.

Peter Shor is the Morss professor of applied mathematics and the chair of the Applied Mathematics Committee in MIT's department of mathematics. His current research interests are in theoretical computer science, error correction, and fault tolerant quantum computing. He is known for his work on algorithms and quantum computation, in particular, for devising Shor's algorithm while at Bell Labs. Shor's algorithm is a quantum algorithm for factoring exponentially faster than the best currently known algorithm running on a classical computer.

He is a fellow of the American Academy of Arts and Sciences, a member of the National Academy of Science, and a recipient of the MacArthur Foundation Fellowship. He earned a BA in mathematics from the California Institute of Technology and a PhD in Applied mathematics from MIT.

Will Oliver is a jointly appointed professor of Electrical Engineering and Computer Science, professor of Physics and a Lincoln laboratory's fellow at MIT. He serves as the director of the Center for Quantum Engineering and an associate director of the Research Laboratory of Electronics.

He is a principal investigator in the Engineering Quantum Systems group and the Quantum Information and Integrated Nanosystems group at MIT. Will's research interests include materials growth, fabrication, design, and measurement of superconducting qubits, as well as the development of cryogenic packaging, and control electronics involving cryogenic CMOS, and single flux quantum digital logic.

Will is a fellow of the American Physical Society and senior member of the IEEE. He serves on the National Quantum Initiative Advisory Committee, the US Committee for Superconducting Electronics, and is an IEEE Applied Superconductivity Conference board member. He received his PhD in electrical engineering from Stanford University, his MS in electrical engineering and computer science from MIT, and a BS in electrical engineering from the University of Rochester in New York.

Well thank you gentlemen, for joining us today. Will, I'll ask you first to bring our listeners up to speed on what is the current state of quantum computing, and what do you feel are the most exciting developments in the last five years?

Yeah, thanks Steve, really pleasure to be here. 20 years ago, experimentally, people were in the lab trying to see if these qubits can be coherent, and act like quantum logic devices, and do single and two qubit gates, and really laboratory curiosity-type experiments. To my mind, what's really changed in the last five years is we got to the point where, for many of these qubit modalities, we know that they work.

And now, we need to improve their performance, and make them better, and build larger systems out of those qubits. So we've moved from a laboratory curiosity to a technical reality. And in light of that, in the last five years, a number of companies are starting to see how quantum computing could potentially, in the future, affect their bottom line as users of quantum computers.

And so many of those companies are jumping in. And at the same time, we have a few companies who build and make quantum computers. And those companies are really pushing forward at high speeds. And what we've seen is that now there are very small scale, but nonetheless, existing small quantum processors that people can access online in the cloud. And that to my mind is really quite remarkable. People are actually trying out very primitive small-scale algorithms, but nonetheless, they're trying them out themselves online.

Peter, what's your take on this?

I am much more tilted towards software than hardware than Will is. So I think the really exciting developments in the last few years have been in quantum error correction. So for quantum computers to ever work on a large scale, you're going to have to fix errors, because you're never going to get quantum computers accurate enough to do more than a few million steps without making an error.

And once a quantum computer has made an error, your computation is probably irreversibly lost. However, there are fault tolerance techniques that can protect quantum computers from errors, even with billions or trillions of steps. And there are two very exciting things.

One is that we're actually starting to see these protocols tested. And everything has been working according to theory so far. And the other exciting thing is that there have been a number of recent developments in quantum error correction codes that look like they'll reduce, eventually, the very substantial amount of overhead that these fault tolerance techniques require.

So just for clarification sake, just go back, what is the promise of quantum computing? Why would people want to make this investment and shift toward this new hardware, software design?

OK, so I have an analogy that is-- you can think of classical computers as cars and maybe quantum computers as theories. So for some problems, there is a shortcut from your starting point to your destination. It involves taking a path through Hilbert space, which is where quantum computers natural live in. And which classical computers cannot take this path. So this is a computation that could only proceed in quantum space.

So there are places you can only get to with boats. And there are some places where going from Connecticut to Long Island, it's much, much faster if you take a ferry than drive all the way around through New York City. So this is what quantum computing is like.

There are some problems where there are shortcuts through Hilbert space, which get you to the answer much more quickly. Even though your clock speed isn't any faster, you're using fewer gates to get to the same-- or fewer computational steps to get to the same answer than you would on a classical computer.

I really like that analogy. That's a great analogy. Maybe it's worth noting that a quantum computer, in fact, can also run algorithms that are classical Boolean logic algorithms. But it's important to note that they often do it no better than a classical computer. And often they would do it worse.

So of course, yeah, like Peter said, you can, of course, go to New York City and around to get to Long Island with your quantum computer. But it's probably faster to go in this classical computer car, but for certain problems, there are some problems which we either know or strongly believe are really hard on a classical computer. And quantum computers may provide a shortcut, in this case, the fairy to just jump across the Long Island and so on.

Yeah, they're not going to replace classical computers in the future. And in fact, we believe that classical computers will likely be needed in tandem with a quantum computer to operate it, to understand the error correction, and to feedback to make those corrections to the errors as they occur. So it's very likely you'll need both running in parallel. But the problem that's running on the quantum computer, those problems that are meant for quantum computing will vastly outperform what we could do on classical computers. That's the promise.

Great, Will, thank you very much for that clarification. In what ways can we expect quantum computing to change everyday life in the future?

Yeah, well, that's the $10 million question, I guess. At this point, it's very early on in the game. And I like to often think of this by comparing to classical electronic computing in the 20th century. And we're at the 1940s, maybe the early 1950s. And we're asking ourselves, how is classical computing going to change our lives in the future?

And people are thinking about flying cars and things like that. But of course, nobody really knows. And in the case of quantum computing, what we do know is that, of course, we're here with Peter. And his algorithm on a quantum computer at scale can attack a commonly used public key cryptosystem, right?

And so we know, immediately, and this is happening right now that we want to switch over to new cryptosystems that we believe to be immune to attack by a quantum computer. And that's ongoing work. Maybe you've heard that NIST is holding a competition of sorts. And they're starting to release some of those early candidate cryptosystems so people can test them out.

So that's one concrete thing that we know is coming. And people are switching over to these new cryptosystems. But there are many other potential applications of quantum computing, but those still need to be borne out. And some of those areas relate to simulation of quantum systems. So that might have impact on materials, quantum chemistry, pharmaceuticals.

There are ideas for algorithms that will improve optimization, for example, sampling solutions to systems of linear equations. But as Peter said, we need to build hardware that can operate billions or trillions of gates or clock cycles. And to do that, we're going to need error correction. And really, what we can do before we have error correction is highly speculative, in my opinion.

And Peter, so are we seeing the end of RSA security public private key encryption with quantum?

Well, you need a quantum computer that can do a billion gates without errors before you can break RSA. And we are not anywhere close to that. Right now, I guess, in the last month or so, NIST rolled out the first proposal for post-quantum cryptosystem that cannot be broken by a quantum computer. And I think it looks very secure.

I mean, there are-- of course, any time come up with a new cryptosystem, you have to worry that you didn't make it big enough. I think the principles that this works on, which is some problem called learning with errors. I think learning error with errors, really, is hard for quantum computers. But you have to hope that you chose a large enough learning with errors problem for the specific cryptosystem. They cannot be broken. And that is probably true, but we'll just have to wait and find out.

And Steve, one thing maybe to add here is Peter's spot on. We don't have a quantum computer anywhere near the capability that would attack today's RSA systems. But I think it comes down to a risk analysis that each of us, and companies in particular, would need to take, which is that there's probably material today that you want to maintain its security or remain secure for a long period of time. And that might be 20, 30, 40, even 50 years or more.

And so then it comes down to, well, what's the likelihood you're going to have at scale quantum computer in say, five years? Well, maybe not that high. And you're not too worried about it. But how about 20 years, or 30 years, or 50 years? And at some point you realize, gosh. Maybe we better be switching over to a new cryptosystem now, even though we don't have a quantum computer that's large enough to attack RSA today. And you don't want to wait for that crypto-- or for that quantum computer to come because by then it's too late.

That's a good point. So in addition to building software that can handle error correction for quantum computing, what are the next big challenges to tackle when it comes to quantum computing, is it the commercialization of this technology? Is it the developing software tools? Is it the hardware? In your opinion, what is the next big challenge to tackle? Peter.

All of the above. So I mean, we're not going to get a practical quantum computer without software tools. And people are developing the software tools, but there's a lot of work. We're not going to get it without hardware. We're not going to get it without efficient error correction. There's probably more things we need to worry about too. Yeah, I mean, quantum computer architecture, we definitely need to worry about.

Yeah, sometimes I like to think about it as, really, from the perspective of a university. You have lots of departments both in the School of Science and the School of Engineering. And I think going forward, to develop a new field, and a new technology, a new business sector, we need to draw on talents from all of these departments. Almost literally all of these departments.

And it's not-- of course it includes physics, and it includes mathematics, and electrical engineering. But it includes chemistry, and biology, and even the business school, of course, because we want to eventually commercialize these technologies. And that, of course, falls squarely in that domain. It is a vertically integrated system that needs to be developed all the way down from the hardware, which we're making progress on. But some of the challenges there is this hardware needs to have much better performance than it does today.

As Peter described, quantum error correction is basically building resilience through redundancy. And the amount of redundancy we need goes down if the fundamental constituent elements, individually, are better. Put a bunch of kindergartners together and trying to do something that's resilient is really challenging. But putting a bunch of adults together in the same room, OK, there's a lot that we can accomplish.

With that sort of intuition, we need to improve our qubits quite dramatically. But from there up the stack, there's quite a bit of work that needs to be done, including developing software tools, design tools, EDA, the manufacturing processes that we need. All of the dual-use, off-ramp, classical technologies that are going to make this real. The lasers and the microwave gear, that's an entire sector that is pivoting currently to help address these problems.

And we need to train a workforce, frankly. That's one of the bigger challenges we have right now is that there just aren't enough people in quantum or being trained to be in quantum. And along those lines, Peter, and I, along with Aram Harrow, and I. Chuang developed an online course through MIT xPRO to help people already in industry learn what quantum is and pivot to quantum.

So those types of classes are starting to become available. And at MIT, in the EE department, for example, we're starting a quantum track so that EE students-- they'll still get an EE degree or EECS degree, but they'll be able to have on their diploma, that the track was quantum. We need to develop these curricula. And we need to identify what it means to be a quantum engineer, for example. And what are the courses we need to teach. And what do these people need to leave MIT with so they can go out be hirec by companies and add productively to this endeavor.

So do you think we're at the place where, for example, with punch cards with mainframes? Is that where we're at with quantum today? And how fast do you think we can get to the personal computer from mainframes to something that's usable in everyday use or in applications in quantum?

Maybe we're not even to the punch card. And talks I give about quantum computing, I had this great slide where there are two pictures. One of them is of an early particle accelerator. And one of them is of one of the first computers. And they both are these big cabinets, with lots of wires, with a technician dressed in 1950s clothes, doing something with these wires.

So that's probably the stage where we're now. It's not really mass produced thing yet. It's laboratory experiment, which each of them is-- each of them is unique and needs lots of technical upkeep. It's generally very expensive because you need an army of technicians to make them run.

And actually, it's not quite as bad as that because we saw over the pandemic that you could actually automate them and have them run mostly in place. But still, you need an army of technicians to build each one individually. And this is going to come to an end when people start mass producing them. And you can see that that's on the horizon, but it hasn't happened yet.

How many years of decades are we from that, in your estimation?

Well, I think as soon as people-- as soon as there is a real use for that. People are going to start mass producing them and selling them. So the question is, when is the first-- when are they going to be big enough to be really used?

Yeah, and maybe I could add to that. That really motivates what I think is-- and I'm sure others would agree that one very important direction is that we need more people thinking about quantum algorithms that are useful, and practical, and address real world problems. We need more such algorithms. Now, many of them may be in the fault tolerant era, when we have error corrected machines. But we would like to have more. And I think getting more people on task is quite important.

So if you look back at the history of classical computing, one way that algorithms were discovered is people playing around with computers, for example, the simplex algorithm was dreamed up in the 1950s. And it worked wonderfully. And people programmed it and used it. But it took somewhere between 30 and 50 years before people actually understood why it worked.

So right now, we're at the stage where we can come up with these algorithms, but we can't program them up and see whether they really do work. And we also are having a lot of time proving that-- a lot of difficulty proving that they actually work. But some of them may work and we can't prove it. And we can't experiment. So we don't know that they work.

So there are other limitations of quantum computing?

Well, I mean, another thing is that quantum computing algorithms are really very hard to come up with because, well, basically because a quantum computer is not going to run any faster than a classical computer unless you use the principle of interference. And interference means that you program up the classical computer so that all the paths through the computation that produce the right answer constructively interfere. And all the paths through the computation that produce the wrong answer interfere destructively. And this is something which is really very difficult to have any intuition about.

Yeah, I was going to just add to that there may be some misunderstanding, I'd say, in the broader public domain about thinking that a quantum computer is trying all different possible inputs in parallel and then finds-- and sends the answer to all of them. And you pick out the best one. But that's not really what's happening.

I think what Peter just described is what's happening, which is that you may start with an equal superposition, as we say, of all input states. But then through this process of quantum interference, both constructive and destructive over the life of the algorithm, the answer, so to speak, grows in likelihood. And the other answers decrease in likelihood. And eventually, there's really, ideally, just one result that when we measure it with very high probability, we get that result.

And that's really how they work. And the limitation is that we don't yet-- and I'd be curious if Peter agrees with this. But I think we don't yet have the intuition or the abstractions needed. Basically, we don't have the ingredients yet to be able to add water, and add these ingredients, and stir, and say, here's our next quantum algorithm. That those abstractions have yet to be discovered or codified. And we have certain examples. And Peter Shor's algorithm is a great example of one. But we would like to have more so that we can understand how to generate new quantum algorithms.

Can you talk a little bit about some of the mistakes, maybe, that you're seeing in the application of quantum computing in the industry or what can business leaders do to think about the future in implementing a quantum strategy? Could be any particular vertical like finance or anything.

So one of the misconceptions that a lot of people had in the early days of quantum computing, when I was talking to journalists and I think a lot of people still have, is that it's not the case that a quantum computer is just like a classical computer, only faster. You need a new kind of algorithm to make quantum computers run faster. And a lot of problems, we don't have such an algorithm. And for a lot of problems, it's quite likely that such an algorithm doesn't actually exist.

So quantum computers are not going to be general purpose computers that run faster for everything. There's going to be some class of problems that can speed up. And what we really would like to know is what is the class? How do you tell whether a problem is in this class? And how do you find algorithms for problems in this class? And we really don't have a very good handle on this yet.

Yeah, and I'd add there that right now, there is a lot of hype. It's quite frothy right now. And of course, the reason there's hype is because there's a lot of promise. And there's hope. And there's promise to quantum computing. But to realize that promise, we need to build large-scale, robust quantum computers. And the time scale over which that's going to happen, generally, is longer than I think is being acknowledged in some of the hype that we're hearing about. So it's going to take longer than two or three years, right? That's just the way it is.

Somehow, I find that when we think back to classical electronic computers in the '50s, we look back with a history knowing what happened and how it unfolded. And we're like, of course, it's going to take decades. But when we think about quantum computers today, because we're so excited about them and think maybe it's going to happen. And we want to realize that promise-- and because businesses are trying to stand up, then the timelines become much shorter.

So I think misunderstanding what we're able to do in two or three years is a misunderstanding. At the same time, I think humans are notorious for also misunderstanding what can be accomplished in a 10-year time frame or a 15-year time frame. We often underestimate that. So of course, I don't have a crystal ball. So I don't know.

But I think that one thing that we can do and companies can do is don't bet the farm, but start assembling researchers who are thinking about how a quantum computer could affect my businesses, bottom line, in the future. And have a small team.

Interact with folks at a place like MIT, for example, through the MIT Center for Quantum Engineering and our industrial consortium, the Quantum Science and Engineering Consortium or the QSEC, which provides a way for companies to dip their toe in the water, and to get started, and to understand what these implications might be in the future, understand what algorithms could impact their future, bottom line. And also understand the caveats, right? Where we actually are and try to cut through some of that hype. I think that's one of the benefits of being in a consortium and interacting with folks at a university.

Can you speak to some of the companies that are involved in the consortium?

Yeah, sure. For example, BMW is exploring optimization algorithms that would apply to manufacturing processes. And the way it works, of course, BMW joined our consortium a couple of years ago. And they're working with one of my students. Tim-- or two of my students actually, Tim Menke and Will Banner.

And in fact, also working with Zapata, which is a startup company that performs algorithms. And they also are members of the consortium. And so at this point, it's researchers from all three organizations-- MIT, BMW, and Zapata, working together to first try to identify what is the best classical algorithm that BMW can practically use today on their problem. And that's a benchmark.

And then look at what is a tensor network approach to that problem, which often is referred to as quantum inspired. And then third is, let's look at what quantum algorithms would be available for this problem. And then the fourth thing is that Zapata has access to some of the small-scale online quantum computers. And we could even try that algorithm out on a small scale quantum computer.

So again, I don't want to add to the hype to say that that's the solution, right? That's not the case, but what is important about doing this is that BMW and their researchers become familiar with the types of algorithms that are available, also the types of algorithms that are not available. And then also how this could affect their bottom line in the future. And as they build their team over time, the idea is that they would be ready to take advantage of a larger scale quantum computer when it becomes available.

And where could folks who are listening find out more about this consortium?

Yeah, so of course, please check out our website, Center for Quantum Engineering, it's cqe.mit.edu. And from there, we have links to the quantum consortium.

We look forward to seeing where your research takes us and seeing the evolution of quantum computing. Gentlemen, thank you very much for your time today. Appreciate it.

Thank you.

If you're interested in learning more about the CSAIL Alliance Program and the latest research at CSAIL, please visit our website at cap.csail.mit.edu. And listen to our podcast series on Spotify, Apple Music, or wherever you listen to your podcasts. Tune in next month for a brand new edition of the CSAIL Alliance Podcast and stay ahead of the curve.