Research @ Citi Podcast Episode 60: What's Next for Quantum Computing With Noelle Ibrahim From IBM

Recorded: Dec. 4, 2025

Published: Dec. 30, 2025

Host: Rob Rowe, U.S. Regional Director of Research and Head of Global Strategy and Macro Group, Citi

Guest: Noelle Ibrahim, Technical Client Advisory Executive, IBM Quantum

Transcript:

Rob Rowe (00:02)

Hi, everyone. I'm Rob Rowe, U.S. Regional Director of Research here at Citi. With me on the podcast today we're thrilled to have Noelle Ibrahim, who is by her own right a quantum physicist, but also one of the senior folks in terms of sales of quantum to IBM's clients. Noelle, thanks for being with us today.

Noelle Ibrahim (00:24)

Thanks Rob. It's a pleasure to be here and thanks for the lovely introduction.

Rob Rowe (00:28)

Noelle, maybe we can start with two simple questions just for our listeners, especially for those that are unfamiliar with quantum computing. What is a quantum computer vs. traditional classical computers today?

Noelle Ibrahim (00:41)

Absolutely. A classical computer, as you all may know, has a bit that can take on values of either zero or one. The quantum computer has a qubit and it can take on the value of zero, or one, or anything kind of in between. You can think of it as being in both states at the same time.

So, what that means is if I'm in 50% zero and 50% one, I can perform one operation on that qubit and I can manipulate two states at once, both the zero and the one. Now I also need to leverage the other two properties of quantum mechanics that are used in quantum computing: One is entanglement, and the other is interference.

So, entanglement allows me to introduce correlations. That can be useful, for example, — and I oversimplify — but if I want to talk about, say, default rates on a loan and unemployment rates, I can have a register of qubits representing the unemployment rate and a register of qubits representing my default rate, and I can entangle those two registers so that I design my algorithm: If the unemployment rate goes up, my default rate also goes up.

And then I can use superposition to enumerate over all the different possible economic scenarios of different unemployment rates, and I can store those all properly in my computer.

Finally, I use interference. And interference with an intelligently designed quantum algorithm allows me to manipulate those qubits such that the scenarios or answers of interest are the ones that come out of the computer and the wrong answers, or the scenarios I'm not interested in, cancel each other out or interfere with each other and are not produced as an answer in the computer.

So that's a fundamentally different way of operating than classical computer, because with the classical computer, you have to enumerate all those scenarios sequentially. And that's why you can design quantum algorithms that can have a speed-up over these classical algorithms by leveraging quantum mechanical laws of nature.

Rob Rowe (02:45)

So much going on in the quantum world, I'm very excited about it. I've gone down the rabbit hole on quantum computing, and we've certainly had a lot of discussions in the past.

I think there are several things that are probably on our listeners' minds. One is, “What is the current state of quantum computing?” And when I think about that, I think about quantum advantage. What is quantum advantage? What does that really mean? And then how far away are we? And where I'm leading to is also IBM's road map, the definition of where you're all going in the progress of quantum computing. We've heard a lot of things about its capabilities, but where are we on that road map? Maybe I can start there.

Noelle Ibrahim (03:34)

So, I think a good place to anchor the conversation is really around the definition of quantum advantage, because that's what we're all aiming toward. And a quantum advantage will occur when a quantum computer can compute a useful calculation either faster, more cost-effectively, or more accurately than the best classical algorithms available.

Now, it's important that this calculation can be verifiable as well. We can throw up our hands and say, “we think we've got something,” but unless we can actually prove and benchmark that properly against the best classical algorithms, it's hard to really determine that we've actually reached quantum advantage.

So, in terms of where we are in the field right now, we're at a point that IBM calls (and has called since 2023) “the era of quantum utility.” And what that really means is that we now have quantum computers that can compute calculations that are not accessible to brute-force exact classical calculations. So, we can calculate things that have never before been attainable in terms of calculations in the history of computation. And now we're getting to the point where we're saying, “OK, well, when will these be useful calculations?” And that's how we get to the point of benchmarking against the best classical algorithms and defining quantum advantage.

Linking this back to the IBM quantum road map, of course, the road map is very hardware-centric because underlying all of these algorithms is obviously the hardware capability that allows us to take advantage of the properties of quantum mechanics that will facilitate these more-efficient, cost-effective or more-accurate calculations.

And the hardware is a rapidly developing area of quantum computing. It's considered more of an engineering challenge at this point than a challenge in fundamental physics. And as we saw with the Quantum Developer Conference, IBM is rapidly progressing: We announced the new Nighthawk processor. We can delve into that a little bit later and Loon, etc. So, we're rapidly ticking the boxes on that road map, moving forward toward the ultimate goal of large-scale, fault-tolerant quantum computing.

Rob Rowe (05:54)

When we say a lot of it is happening in the hardware development, what does that mean? I know that some people say a milestone is 100 logical qubits, and I know that maybe 2025 could be considered the year of error correction, but I'm not sure if people understand what that means in terms of the development of these computers, and is there a simple way to look at the development of the hardware?

Noelle Ibrahim (06:22)

So there are many, many milestones. I'll hit a few of the high-level milestones here. One of the focal points of the field is the attainment of error correction and the logical qubit. So let me take a step back and just explain for the audience what a logical qubit is.

So we have physical qubits. They may have different … you may have heard of the word “modalities.” So these are objects that can actually carry quantum properties and be manipulated to take advantage of quantum properties. We call them superposition, entanglement and interference. And those are properties that aren't easily accessible by classical computers, that allow for quantum advantage to develop.

Now, I can have a device that can allow me to take advantage of superposition, entanglement, and interference, but it's noisy So, it needs very carefully constrained environmental conditions in order to maintain those quantum properties. You have either dilution refrigerators or you have special advanced laser systems, etc., to make sure that we maintain those quantum properties in the physical qubit. Now, a logical qubit is made up of several physical qubits. And instead of encoding a zero or a one in one physical qubit, we encode a zero or a one in a bit string composed of several physical qubits. Those bit strings have properties that allow us to basically correct errors that are introduced by noise, which is what happens in the physical qubit when the environmental conditions are not fully controlled.

So when we are able to correct those errors that are induced by that environmental noise in a string of several physical qubits, we're able to achieve error correction and able to build logical qubits that are essentially able to more robustly carry forward those quantum properties.

I mean, classical computers also have error correction — it's just a little less complicated. But in any computing system, you want to make sure you're fully in control of what you're calculating. You don't want random errors to come into it.

Rob Rowe (08:45)

I understand. And a lot of people talk about the “five nines,” right? That we have to achieve the 99.999%. Are classical computers already there? Like, how many nines do they have? I guess for the audience's sake, it's 99.999% accurate, right? Is that correct?

Noelle Ibrahim (09:05)

Yeah, I think it's some very large number of nines for the classical computers. And we're working toward that as well with the quantum computers. And it's true that one of the requirements of creating a logical qubit is that each physical qubit that comprises a logical qubit has to be of a certain quality. It has to have a very low noise threshold. And so when you hear about people talking about the accuracy and the noise threshold that they've achieved with their physical qubits, it's because they're working to get below that threshold that will allow them to create a logical qubit out of several physical qubits.

Rob Rowe (09:44)

And Noelle, how are the recent developments at IBM affecting this process? You mentioned Nighthawk; what I've noticed is when I see a superconducting quantum computer, I see a big chandelier with some wafer at the bottom that has the qubits. But then when I'm reading about Nighthawk, I see a processor.

So how have Nighthawk, Loon, and error correction been developed at IBM recently to improve or to get that accuracy that you're talking about?

Noelle Ibrahim (10:14)

So Nighthawk is really exciting. It's still not fault-tolerant yet, so it's still a utility-scale … kind of our utility era processor, but very exciting because it's introducing more communication between qubits, between nearest-neighbor qubits, which allows it to run much more complex circuits than the Heron that we have now. And Nighthawk will be available at the end of this year, which we're pretty much there already.

So that's going to allow more complex calculations to be accessible that have never been accessible before in the history of computation. And as we keep developing that hardware and reaching more of those milestones on the IBM quantum road map, we're going to be able to access more and more complex and larger and larger calculations that can take advantage of these properties of quantum mechanics. And thereby, we will increase and improve the separation, or the potential for the separation and performance between quantum algorithms and classical algorithms.

So that's all taking place on the foundation of these hardware milestones. And then IBM Quantum also announced the Loon, which is an experimental processor. And it will have all of the hardware components that are required for error correction.

So, for example, it will not only have these nearest-neighbor communications, but it will have some longer-range communications between the qubits that will help facilitate quantum error correction. It will also have the ability to more quickly reset the states of these qubits. And then we recently, you may have heard, were able to perform part of error correction on an AMD FPGA. So that's today's hardware, we're able to do some of the decoding that's required for correcting errors in real time on today's classical hardware.

So all of these things involve both quantum and classical computations and measurements. And so these milestones on the road map are really facilitating the working together of all these different pieces to build that error-corrected pipeline.

So we have Nighthawk, which is usable today for utility-scale calculations, where the number of qubits is your physical qubits. And then we have Loon, where we're working toward building those logical qubits and error correction with the longer-range connectivity, etc.

Rob Rowe (12:46)

So, it's interesting, because you're mentioning some other firms — one thing I've noticed is that a development within the industry seems to be integration with classical computing, integration with AI, but also integration with certain partners.

I think about NVIDIA and Google, and now you had an announcement with Cisco recently, as well as another one with AMD. What is the state of the business here? There are integrations with partners. Are we now moving on to quantum computers being integrated into the traditional computing network? Is it being integrated into the classical computing network via these various groups, Microsoft, others?

Is Cisco there to help development of the processors, or there to help development of integration?

Noelle Ibrahim (13:47)

So the Cisco relationship is more on the integration side. And this has to do with networking between different quantum processors, and then ultimately we'll move toward longer-scale network-like communication between networked quantum devices and ultimately internet.

So that goal is a bit farther away, but this interconnectivity between devices that, for example, may not be contained in the same dilution refrigerator, is really essential for scaling up our quantum computing capabilities beyond what IBM has on their road map today.

So if we're looking at very large-scale quantum computing, we've got to have multiple machines and multiple processors connected together. And then linking back to what you were saying about AI and other classical computing, the ultimate vision is not for quantum computers to ever replace classical computers at all. The ultimate vision is quantum-centric supercomputing where you will have AI — like bits, qubits and neurons, essentially.

So you'll have high-performance computing, you will have artificial intelligence, and you will have quantum computing together. And you will be able to ultimately route jobs and tasks or subtasks to the best processor.

And of course, that's an ongoing area of development. So nobody necessarily knows today which algorithm would necessarily always be best routed to which processor. And there's always going to be exciting developments in these algorithms that actually might change even which processor they could be best run on, etc.

But it's going to be a very fruitful and exciting endeavor. I did my PhD in classical computing, and the algorithms are always developing. And so you're never at a static point where we've reached the pinnacle of success. You're always developing something new.

And so now we have these multiple dimensions in which you can develop these new algorithms. So I'm kind of an algorithm geek: You can target it, but it's never going to be static, it's always going to be evolving. So that's the vision: all these things linked together, and all these network players will all have played a crucial role, such as AMD and Cisco and others.

Rob Rowe (16:13)

You mentioned algorithms and you said you're an algorithm geek. Good, because the next questions really have to do with algorithms.

When I was at the IBM Quantum Summit, I noticed a particular slide, which I found very interesting because I think the listeners are probably wondering about the timing of these things and how these things will be implemented. And one thing that everybody is very scared of is this idea of Shor's algorithm, the potential breaking of encryption, and also how that can affect crypto markets, etc. On a very simple level, people are worried about that.

But the slide that I saw that I'd love you to expound on is one which showed where quantum computing capability is. And on the other side in various clouds was where these algorithms required capacity in order to be implemented. And it looked like actually there's still some distance between the capacity of these quantum computers and say something like Shor's algorithm, which may need an even stronger capacity than what would be defined as quantum advantage.

Is that true? And how does that affect the timing of things, do you think?

Noelle Ibrahim (17:27)

For running something like Shor's algorithm, we're not estimating that that will be possible, like in terms of breaking our saying, for example, on the first generation of fault-tolerant quantum computers. So, on the IBM road map we have that coming in like 2029.

It'll be a bit larger-scale because of the number of logical qubits required to actually break the keys at the sizes that they're used today. But it is coming. We do think that there'll be other advantages that will evolve before that point. And that's why we're talking about the era of utility. We don't even think that quantum advantage will require fault tolerance, necessarily. And IBM is predicting that there will be some meaningful hypotheses about quantum advantage before the end of 2026, even on today's devices.

So that's why we've launched, for example, the quantum-advantage benchmarking initiative. We've collaborated with the Flatiron Institute, Algorithmic, and also BlueQubit to create a repository where groups can put forward their hypotheses for quantum advantage and then benchmark them, in this open-source environment, against the best classical algorithms. And that's something that we're actively working on today. So we don't expect that we'll have to wait for Shor's algorithm to be mature before we will start to see the first quantum advantages.

However, we do have the problem of “harvest now, decrypt later.” So in terms of quantum safety, we do recommend that organizations get started today protecting their data with quantum-safe algorithms, because bad actors can actually harvest your important data today, store it, and then decrypt it later when a cryptographically relevant quantum computer that can break RSA encryption is available.

So if you have any data with long lifetime value, it's best to get started updating your quantum algorithms now. In essence, the time to get started for both actual application advantage and quantum safety is now for those reasons.

Rob Rowe (19:48)

Maybe we can talk a little more about that, Noelle, because I know IBM has this program called Quantum Safe. What's interesting to me was — and correct me if I'm wrong — we've discussed previously that you can do things on the classical side to protect yourself now into a post-quantum encryption world. Is that right?

Noelle Ibrahim (20:11)

Absolutely, absolutely. So there are new algorithms that are lattice-based algorithms that run on a classical computer, but they use a mathematical framework that cannot be cracked or undone by either a classical or a quantum computer.

So RSA encryption is based on the factoring problem, and that factoring problem is something that is easily solved by Shor's algorithm, which is very difficult to solve using any classical algorithm. So we've moved away from factoring in terms of quantum-safe encryption to lattice-based algorithms that are run on a classical computer, but which cannot be cracked by a quantum computer.

Rob Rowe (20:55)

Wow. And maybe we can talk about two more subjects, time permitting. There are tons of modalities out there. Well, maybe not “tons” — I keep saying “tons of modalities,” but there's five, six, seven, maybe, and others in development. And by modalities, for the listening audience, IBM's computers are a modality called superconducting. There are others that have modalities such as neutral atom, ion trap, photonic — to name a few — but there's still some others in there. I think there's spin and a few others.

There's a lot of differences there, so integration seems challenging in one sense, because I think classical computing can all be considered one modality, and yet we’ve got like five or six here. Do you see this evolving as eventually there'll just be, as you'd say in “Lord of the Rings,” one modality to rule them all? Or do you think that there are essential applications for all these different modalities where one can excel over the other? And how will all that get integrated, do you think?

Noelle Ibrahim (22:00)

There are definitely some different views in the industry. IBM does have a partnership with PASQAL, which is public, and they have a neutral atom modality. Neutral atoms and superconducting qubits are complementary in certain ways — the superconducting qubits have faster gate speeds, but the neutral atoms have a lower noise threshold. So, depending on the problem, one or the other of those properties may be more important for achieving a quantum advantage. So in terms of that vision of the future of quantum-centric supercomputing, IBM and PASQAL are working together. We're not necessarily placing our bet on just one modality. And we do have that road map in play.

Rob Rowe (22:49)

Interesting. And back to algorithms for another second. Algorithms are developing also at the same time. So, in a sense, the chronology on this or the road map is a matter of algorithms getting developed and becoming more efficient while quantum computing is also developing and becoming stronger. So that middle meeting of the ground somewhere, it's very hard to figure out that timing.

But can we talk about some of the other potential applications for quantum computing? What do you foresee happening, say, in the financial sector? What about medical, pharmaceuticals, bio, pharma, also chemistry and material science? How is quantum computing going to affect that, do you think?

Noelle Ibrahim (23:33)

So there are some very interesting applications. Let's start with the financial-services sector, areas such as optimization. And if you're looking at more complex optimization problems, that's an area where there's potential for quantum advantage. So IBM published a paper a while ago called the Intractable Decathlon, and it lists like 10 different types of optimization problems that are very challenging for classical computers to properly treat. And one very interesting one for finance was a portfolio-optimization problem looking at rebalancing over several time periods. And then they looked at what we call “the gap to optimality,” which is basically, “How far away is the solution that is found by the algorithm from the best possible solution?”

And there are ways of estimating, for example, your risk and reward trade-off for the best possible solution without knowing what the best possible solution is. When we had risk-averse clientele that we're modeling, essentially, the gap to optimality was increasing rapidly for the classical algorithms, but not so rapidly for the quantum algorithms — which basically means that if you're doing portfolio optimization in a complex enough scenario, and you're looking at a risk-averse scenario, you may have an opportunity for quantum advantage in that kind of setup.

So that's one of the areas we're looking at. There's higher-order binary optimization, which basically would come into play if you're looking at characterizing risk with more complex mathematical structures, multi-objective optimization. So there's just many different areas that we're looking at right now with financial-services firms.

In terms of chemistry and materials science, there are some very interesting algorithms that actually look at estimating the energy levels of different chemical compounds. And when you think about these energy levels, you can think about it in terms of the energy level being important for understanding the progress of a chemical reaction that might be of interest.

So let's say I want to synthesize certain products and I start with certain reactants. If I want to know what proportion of different products will come out of that reaction, it's helpful to know the energy levels of the different starting reactants and the different end-point products. And sometimes those can be difficult to calculate with a classical computer.

With a quantum computer, you can actually better model some of the quantum interactions, etc., to better predict what your outcome is going to be in a chemical reaction. And if you know what the outcome is going to be, you can better engineer that chemical reaction to produce the products that you're interested in. There are a lot of applications like that, like surface absorption of materials that cause corrosion, and things like that are areas that we're looking at.

Rob Rowe (26:39)

And with corrosion, with something like that, that's an area where you can actually run simulations or molecular simulations to the degree that you can identify ways to reduce corrosion?

Noelle Ibrahim (26:50)

Yeah, I mean, let's suppose you wanted to look at different surface coatings or something like that and see which one was more effective for reducing corrosion, or calculate the rate at which corrosion might happen in different environments, etc., for your type of material. That's something where you could look at, for example, the energy levels, the corrosive material on its own, and then what the energy level would look like if it's adhering to the surface of your material of interest. And then you can understand how those energy levels might change in different environments, and then you could control the environment potentially if you can understand it better. So, that's the high-level idea behind it.

Rob Rowe (27:43)

Wow. And lastly, Noelle, because I didn't specifically hear about the road map, when does IBM anticipate having a fault-tolerant computer?

Noelle Ibrahim (27:55)

So in 2029 our road map has the Starling processor, which will be the first fault-tolerant quantum computer running error correction. And all the milestones up to that point — what we're seeing, for example, with the Loon, etc. —are leading up to that fault-tolerant quantum computer.

And one thing that's really interesting to note is that even the error-mitigation techniques that we're developing today will have relevance, especially in the early stages of fault-tolerant quantum computing, because we expect there to be a mix of error-mitigation and error-correction techniques in use.

So everything we're doing is really culminating in that fault-tolerant quantum computer, the first generation of which will come out in 2029.

Rob Rowe (28:43)

Noelle, thanks for a really great and engaging podcast session today. And we really look forward to watching the developments at IBM.

Noelle Ibrahim (28:53)

Excellent! Thank you for having me. It's been a great discussion and excellent questions. And I hope everyone enjoyed it.

Rob Rowe (29:00)

Thank you.

Disclaimer (29:01)

This episode was recorded Dec. 4, 2025.