Decode Quantum with Christian Weedbrook from Xanadu

We are back after seven months of break. This is the 85th episode of Decode Quantum, with me, Olivier Ezratty, and Fanny Bouton.

In this new episode, we welcome **Christian Weedbrook**, the founder and CEO of **Xanadu Quantum Technologies** (Canada).

He is an Australian born physicist and entrepreneur based in Toronto, Canada. He has a PhD in physics from the University of Queensland, specializing in quantum information theory. He was a postdoc at the MIT and the University of Toronto. Initially, he went to a film school with the aspiration of becoming a director before returning to mathematics and science. He also explored several small-scale entrepreneurial projects including selling heat packs and trading at local flea markets. In 2014, he launched CipherQ, which was specialized in quantum security and cryptography, and he then created Xanadu in 2016.

Here is an edited transcript of this podcast, with some additional links and illustrations.

Fanny Bouton: our first usual question is how did you land in quantum? What created the flame to start in this strange topic?

Christian Weedbrook: it is very strange for sure. I was following things that I enjoyed and that I was good at. Originally at the university, I majored in math, and that was a bit too "mathy" for me. I decided to do more applied maths, which I enjoyed. It got even more applied and became physics. So, I did a major in math and physics. That's how I started in physics. I wanted then to do an honors degree kind of like a master's. I was really excited by things that I was hearing around at that time on quantum teleportation, quantum computing and quantum security. So, I applied and got into the honors program there, and I did a project in quantum security and enjoyed it. I wasn't sure if I was able to do a PhD, if I'd be accepted. But at the end of the honors, I got accepted. Each step really came naturally after the other, and I was just looking to continue learning, but also working on something that I was okay at. It just went from there really after the PhD. I wanted to continue with the same mindset and was a postdoc, and so forth.

Olivier Ezratty: what was this company CipherQ that you created?

Christian Weedbrook: it lasted between two and three years. I had contributed to the academic field in quantum security and quantum computing. As I was always an entrepreneur at heart, I thought that starting a company in quantum security would be a bit easier, just because this was 2014. The technology was like it is today but not as advanced. It's not needed to be as advanced as building a quantum computer, so from a technical point of view, it was much easier. I got some contracts and some revenue. And when I went out to raise money, investors were more interested in seeing if I had a background in quantum computing and indicated that they would be more likely to fund a quantum computing company. The reason I bring that up is both companies existed at the same time for maybe one or two years after. CipherQ was still going and I had some work to finish off. I didn't know if Xanadu would take off, because at the time it was a bit too ambitious to build a quantum computing company, particularly for me, as a sole founder and my background is not experimental. It's actually theoretical. In the end, I went with the Xanadu quantum computing side of things.

Olivier Ezratty: CipherQ was a QKD company or a PQC company, if you use the traditional segmentation?

Christian Weedbrook: yes, it was QKD. My background which was with continuous variables and homodyne detectors.

Olivier Ezratty: it was using the "prepare and measure" method?

Christian Weedbrook: indeed, it wasn't entangled-based. It was just really the simplest thing you could do based on using coherent states. That's the benefit of continuous variables. You don't need to have single photons, but I guess more importantly or just as importantly, the detectors were homodyne which were operating at room temperature. That was always a benefit, and the rates were much faster for a smaller set of distances.

Fanny Bouton: can you tell us about Xanadu story. How did it start?

Christian Weedbrook: it was this idea that investors were kind of like the initial customers and they were talking more about quantum computing. And at that point, Rigetti had started and raised something like one or two million dollars. I thought, "you can do that!" Like there is still a lot of work to flesh out. D-Wave was around before that. It was really just a case of it would be fun to see if the photonics-based approach and specifically Xanadu's approach could be a contender. I also knew the people in the industry that I'd like to hire, and I managed to raise around two and a half million dollars. I thought that was all the money in the world, and I still do. I was surprised that money appeared in the bank account one day. And it really just went from there and offering a different side of things. That's why we're now going public as well because the pure play photonics-based approach isn't represented in the public markets, and then ten years later, we are where we are now.

Olivier Ezratty: we'll probably discuss about that later before we look at the technology. But yes, indeed, yeah, the SPAC and IPO market is a weird place right now with many companies getting in that bandwagon. What was your vision initially when you created the company back in 2015?

Christian Weedbrook: it was really just that I loved the math and this photonics-based approach that I was contributing to, and it's really a simple idea. I wondered if we got enough money over the years, whether we could actually realize a large-scale quantum computer using photonics. It was nothing more or nothing less than that. It started with a simple idea: let's see if we can get some funding. And honestly, the first amount of funding of two and a half million dollars, it was like I can't believe they gave us any money at all to work on this. But fast forward ten years, we can actually we can feel it in our guts now that we know how to build this large-scale quantum computer. Like we can see it. We have the blueprints. It's still challenging but it looks very different to when we started 10 years ago. It was really just that simple idea. No one was doing it. We have the knowledge and, it would be cool to have our shot at building this type of quantum computer.

Olivier Ezratty: in the last couple years, there were a couple of defining moments, one being the GBS (Gaussian boson sampling) experiment in 2022, published in a **Nature paper**. It was a very visible moment for you because you had this kind of so-called quantum advantage. Before that, you invested of course in the qumode (qubit modality) and continuous variable photonic qubits. Explain us the benefits of photonic qubits and then within the photonic qubit space, the specific advantages of using CV qubits. It's very different from what PsiQuantum is doing, what ORCA Quantum is doing, and what Quandela is doing as well. CV qubits are a bit special. Explain us that.

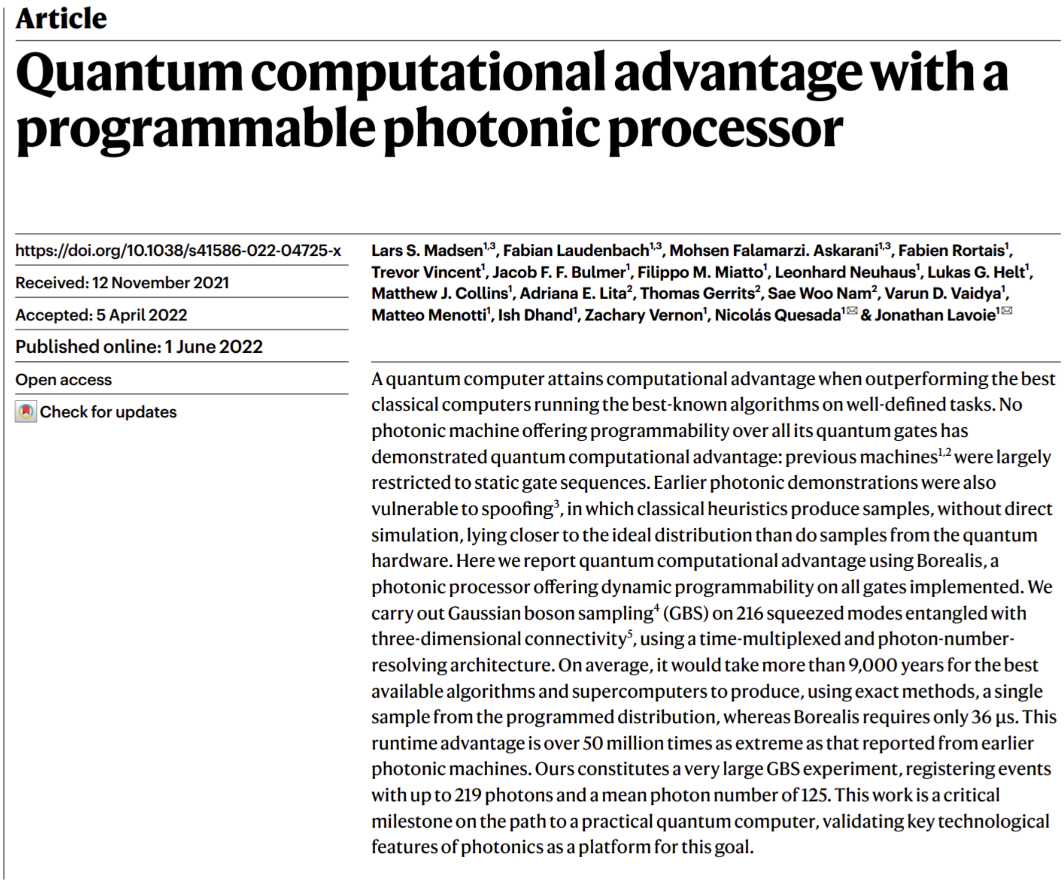

Article

Quantum computational advantage with a programmable photonic processor

https://doi.org/10.1038/s41586-022-04725-x

Received: 12 November 2021

Accepted: 5 April 2022

Published online: 1 June 2022

Open access

Check for updates

Lars S. Madsen[1,3], Fabian Laudenbach[1,3], Mohsen Falamarzi. Askarani[1,3], Fabien Rortais[1], Trevor Vincent[1], Jacob F. F. Bulmer[1], Filippo M. Miatto[1], Leonhard Neuhaus[1], Lukas G. Helt[1], Matthew J. Collins[1], Adriana E. Lita[2], Thomas Gerrits[2], Sae Woo Nam[2], Varun D. Vaidya[1], Matteo Menotti[1], Ish Dhand[1], Zachary Vernon[1], Nicolás Quesada[1] & Jonathan Lavoie[1]

A quantum computer attains computational advantage when outperforming the best classical computers running the best-known algorithms on well-defined tasks. No photonic machine offering programmability over all its quantum gates has demonstrated quantum computational advantage: previous machines[1,2] were largely restricted to static gate sequences. Earlier photonic demonstrations were also vulnerable to spoofing[3], in which classical heuristics produce samples, without direct simulation, lying closer to the ideal distribution than do samples from the quantum hardware. Here we report quantum computational advantage using Borealis, a photonic processor offering dynamic programmability on all gates implemented. We carry out Gaussian boson sampling[4] (GBS) on 216 squeezed modes entangled with three-dimensional connectivity[5], using a time-multiplexed and photon-number-resolving architecture. On average, it would take more than 9,000 years for the best available algorithms and supercomputers to produce, using exact methods, a single sample from the programmed distribution, whereas Borealis requires only 36 μs. This runtime advantage is over 50 million times as extreme as that reported from earlier photonic machines. Ours constitutes a very large GBS experiment, registering events with up to 219 photons and a mean photon number of 125. This work is a critical milestone on the path to a practical quantum computer, validating key technological features of photonics as a platform for this goal.

Christian Weedbrook: the bigger picture sort of photonics versus the electronic based approaches. With the photonics-based approach, you can work with large scale foundries. I think everyone says that, but it's actually true in our case. The reason is that large foundries have been servicing the telecommunication industry for many years now. They use, for the most part, the same materials and tools than the ones we need. We can literally use that. We are using a large number of foundries around the world. The other big great thing about photonics is networking. Most, if not all approaches will have to have, roughly speaking, a smaller quantum computer. You have hundreds of them, and you need to network and connect them together. Very early on, we realized that with electronic based approaches, if you have two kinds of server racks or cryogenics using electrons in both, you got to convert from electrons to photons and then convert back up to electrons again. It's very difficult to do, and no one's done that in a scalable way before. Our idea is, why not keep everything photonic? It's a much simpler proposition to connect discrete quantum computers. The other benefit is cooling. Not all photonic systems are created equally. Everyone needs cooling, including us. We believe we have the lowest requirements for cooling than anyone. The qubits themselves, the gates and the measurements they're all at room temperature. Meaning no cryogenics or any laser cooling, no nothing. It's literally at room temperature. We do need some cooling, but that's really for the analogy would be turning on the computer or initializing it. That's at the start, but the rest of the footprint is at room temperature.

Olivier Ezratty: can you explain that a little bit? What do you initialize at cryogenic temperature? Because I thought it was about the TES, the photon detector.

Christian Weedbrook: it is, it is, but they only help prepare the states, the qubits. We don't use it in any sort of measurement of the computation. Basically, a very simple example: you've got say two laser beams and they're squeezed states. You mix them on a beam splitter and they become entangled. If you measure one arm, one part of it, if you get the right answer, it puts the other state into a state that you want.

Olivier Ezratty: It's a kind of heralding process?

Christian Weedbrook: exactly. The other thing, too, is that we show the GBS demonstration (setup *below*) that you mentioned where we show quantum supremacy. It is continuous variables. We've always had the goal of going to full fault tolerance and error correction using continuous variables. The GBS system which served to show some quantum supremacy will become a part of our fault tolerant quantum computer. The GBS system is how you actually make GKP qubits or photonic qubits in our case. Very similar to what I mentioned, you have these squeezed states, we mix some using interferometers, and then you measure a subset, and then it pops it or collapses it into a GKP state. Then it's a two-level system. And we have the ability to do error correction and fault tolerance. The continuous variables is still an important part, but the overall 2022 demonstration was a stepping stone to full on fault tolerance.

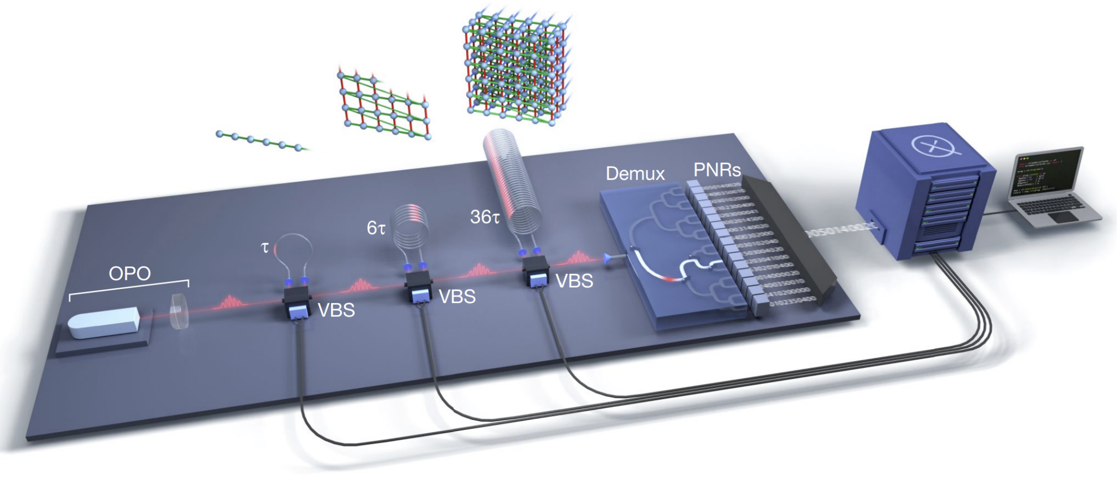

Olivier Ezratty: I would like to understand GKP qubits. What is GKP and what are the differences and similarities with other GKP qubits like the ones from Nord Quantique, which are solid state based. And how do you implement error correction on top of that? Can you describe the space and time aspect of error correction on GKP qubits?

Christian Weedbrook: GKP at the highest level is named after three physicists that discovered or created this type of qubit, and so therefore it's called GKP qubit (Gottesman, Knill, Preskill, in a **2000 paper**). One way to look at it is you can visualize these states as a superposition of squeezed states, and a squeezed state is just a quantum state. You are creating a superposition of quantum states, and they look like fingers or peaks (in the **Wigner quadrature representation**).

Olivier Ezratty: like the cat moustache?

Christian Weedbrook: yes, so you've got this here and you can kind of separate the quantum states by certain distance. And then you say, "Let's call this a zero state, logical zero". Then what you do is you put another one that shifts it a little bit. So now you got a logical "one". Provided these peaks are very, very small, you actually have the ability to distinguish between a logical zero (|0>) and logical one (|1>). And because they're logical zero and ones, you can create superpositions. You have some resistance to errors which in our case is due to photon losses. You have these zeros and ones, but you can add more tolerance to errors by adding another code on top. So, you've got a lower code and then an upper or outer layer code. Then, you can choose whatever code you want. We're focusing on **qLDPC codes**, and so you've got two levels of protection that you can actually do. You do protection for one type of error, say, a bit flip and a repetition code for a phase flip error correction, and then you can kind of combine them together as well. We're using that. For the physical instantiation of this, you can use photons like we're doing or you can use other type of light like microwaves, which is what Nord Quantique and others do. At a high mathematical level, they're exactly the same. It's just how do you actually make it physically. We are using 1,550 nm wavelengths, so telecom wavelength, and they're simply using microwave wavelength to do the same thing.

Olivier Ezratty: there is still a difference here. You are using flying qubits. It's different than solid state and static qubits where you send microwave to control the qubits in some location in a circuit. Where are photons flying from and to, and what are they traversing? And where is error correction handled in that process?

Christian Weedbrook: the good thing about the superconducting approach to GKP states is it's deterministic to create these GKP states, and it's been done for maybe 10 years or more now. With photonic systems, it was much harder to do. And we did it for the first time on chip using photonics. It took many more years to do that. With state preparation, there is always the "no free lunch" theorem. The state preparation is much harder for us. We've solved that now, but it's still easier in the microwave regime. The challenge for them to my knowledge is that the qubit readout can be quite challenging, whereas the readout for us is standard off-the-shelf photon detectors. You're kind of shifting the problem around. In terms of the error correction and fault tolerance, one of the key things for us, if you look at it, starts with a GBS chip to create it as mentioned. Loss is occurring wherever light is propagating. You have it in the rings that create squeezed states. You have it in the interferometers where you entangle the states, and you have it on the readout. On the one hand, what we do is we make every process through the chip design, as low-loss as possible, and we do that by going back and forth between the foundries, iterating on our own designs and based on their fabrication processes.

Olivier Ezratty: you don't mention the foundries. Is it confidential? You don't say it's Global Foundries or whoever.

Christian Weedbrook: we were pretty open, so we've used GF indeed. We also use the **Albany Nanotech from NY Creates**, as well as **Tower Semiconductors**. We work with those guys. We also work with **UMC** in Taiwan, **A*Star IME** in Singapore, and with **Imec** in Belgium. They're some of the major ones that we work with.

Olivier Ezratty: for different components.

Christian Weedbrook: some are doubling up, but some are used for certain things. Some are better at certain low loss for certain materials. In our computer, there's three different types of chips. We use silicon nitride, lithium niobate, and III/V materials for the detectors. Some entities are better than others at the low loss, and low loss is the key parameter that we're optimizing for.

Olivier Ezratty: I suppose you don't use the same line as PsiQuantum at Global Foundries when they have their own Applied Material MBE systems for specific components. That's really important because all the startups try to improve the efficiency of their manufacturing lineup, and they sometimes switch from one vendor to the other. It's, a mix of flexibility in the manufacturing process, but also, I would say professionalism industry grade, and it seems to be difficult in most cases for superconducting qubits. Which is more or less academic kind of fabs, even with the largest ones like SkyWater, which was just acquired by IonQ. What's your experience there on how you build a mix of industrialization on one side and also customization?

Christian Weedbrook: for a lot of the stuff, we're using the most advanced 300 mm wafer tools. We are already using them for a variety of different parts for our quantum computer. Other things we may use university foundries to test new ideas and be, you know, quicker about it to test them much faster. For the most part, we're using 200 mm and 300 mm wafers. It's very like you said superconducting and ion traps in that in principle, they say they seem most compatible. In principle it's true but if you go to TSMC, they're not going to say, "Let's do high volume manufacturing," because they're very picky with their production lines contamination, the tools and all this sort of stuff. For us, it's easier because they're already servicing the telecommunication industry. So, we can present them our chip designs. They'll be still using the same stuff you are already making for the telecom like beam splitters, interferometers, and detectors. It is much easier for photonics companies to use high volume manufacturing lines and foundries.

Olivier Ezratty: but is it the same for III/V because you don't have 300 mm fabs there?

Christian Weedbrook: that's true. For some of the detectors we work with Applied Materials. We use a whole spectrum of 6, 8 and 12 inches wafers. The main thing is is we've started with the most advanced tools. When we do need to scale up, we're trying to get the lowest loss possible. So that's what we're focused on. But when we do need to scale up, we know the foundries are there for us. For other approaches, it's much more difficult if not impossible.

Olivier Ezratty: coming back on error correction, you either have technology like what Quandela is willing to do or Photonic Inc in Vancouver is trying to do. They use static qubits where interactions are used using photons, but the qubit sits in the semiconductor. In your case, you are sending photons somewhere. Is it close to the so-called MBQC and MBQC measurement based model?

Christian Weedbrook: it's indeed measurement based. We create the GKP qubits and, ultimately, we connect them together in a cluster state. What's good about that is our entangling gate is essentially a beam splitter, which is a well-known optical element. Once you create all of these, this large, entangled state, then you perform measurements on the cluster. So, it is an MBQC with the nodes being GKP states.

Olivier Ezratty: but how large are the clusters in that case?

Christian Weedbrook: it depends on the algorithm, and ultimately, the large scale quantum computer will depend on what algorithm you need to run. One of the things about MBQC is you need to, for a variety of reasons, have two different sections where you've stored the qubit in fiber optics. That's our memory because they're flying around. You need to kind of contain them or store them as well. The great thing about having flying qubits is you've got another dimension that you can encode in, which is time. And that's what's really great about our approach, and to be fair, PsiQuantum's approach as well. If you can create 3D type structures, you don't just have to have two dimensions. And, what's even better about our approach compared to others is we can definitely have flexibility in the type of error codes that we want to have. We don't have to build our chip reliant on a specific code like the surface code requires certain characteristics. We don't have to be blocked into those characteristics straight away. That allows us in principle use qLDPC codes and a family of codes that we can leverage by creating the certain types of cluster states when needed.

Olivier Ezratty: so you say cluster states MBQC (from Briegel and Raussendorf) was initially a very theoretical model dreaming about very large cluster states. We know it's difficult to create large cluster states either in DV or CV qubits, so it ended up being coming more like an FBQC model, *aka* fusion-based, with small cluster state and entanglement gates between them. Are you envisioning smaller cluster states like with MBQC or something different?

Christian Weedbrook: with PsiQuantum and us, we use different encodings. We don't need the fusion gates because, once we create our GKP qubits, everything is deterministic. What happens now is you use space encoding and time encoding, creating a two-dimensional cluster. If you look at it, you propagate it through time, and you can picture it like the drag. It creates a square, then like a longer square and so forth. At any one point in time. You only have three slices of cluster states, and so basically, what that means is you don't need to create the whole cluster at time equals t. You are creating the cluster and generating the statistics over a period of time, and that's what's really cool about measurement-based stuff is that you don't need the whole cluster to exist. Say, you don't need this abstractly speaking, you know this whole cluster existing right now. You can break it up into smaller parts and create the smaller parts over time, provided they're entangled and that you're measuring it at each time step.

Olivier Ezratty: so you measure and you have some kind of feed-forward. There's classical computing, I presume to handle the code. I would suspect that you have a challenge in making sure that the speed of your classical error correction and detection is matching what's happening with the photon, even though you may use delay lines.

Christian Weedbrook: we do demuxing to slow it down when needed. Like you said, it's a good problem to have because compared to other approaches, which are relatively slow compared to us, we need to slow ours down. Like you said, so we can use fiber optics to demux and slow it down. What, we have reached at the moment is getting a decoder or speed of the FPGA, in our case around one megahertz. So, we need to extract, we need to measure, extract the data, analyze the data, see what errors occurred, feed it back on within one megahertz.

Olivier Ezratty: is it fast enough, given the speed of photons?

Christian Weedbrook: yes, it's fast. You can slow them down to whatever you want, basically, by increasing the length of the fiber. It's like keep going we'll call you when we need you kind of thing. Okay, but one megahertz is definitely fine. We want it to be faster just because we can. If it's faster, we can then reduce the size of the computer. But also, it's even you know in principle the detectors we have, the homodyne detectors that measure the clusters can be at gigahertz or more. They're telecom so to speak. So, the faster we can but the challenge is finding decoding algorithms that are fast. They're intensive. The fastest we've got now, which we believe is still the fastest in the industry, just because the decoder for other approaches, non- photonic approaches, is not the bottleneck for the speed. The gates are the bottleneck for the speed.

Olivier Ezratty: and qLDPC is known to require a decoder that's kind of costly, compared to surface code or color codes.

Christian Weedbrook: there are indeed always some trade-offs. The trade-off with qLDPC codes in principle, is you only need on the order of a hundred to one physical to logical instead of a thousand or ten thousand. That's in the ideal scenario. But then the decoder can be a little bit more challenging. So, it's always trade-offs with the type of noise that your quantum computer has.

Olivier Ezratty: Fanny, by the way, you visited Xanadu last year, didn't you?

Fanny Bouton: it was in 2025. I visited your office and met your team and the most interesting thing that I saw was the planned size of your quantum computer, with about a hundred logical qubits. It was so huge. I would like to know more about this. I'm working on the future of data-centers as I work for a cloud provider. It is really interesting to see the size of different types of quantum computer and how we can handle it in data-center.

Christian Weedbrook: I am sure you saw we called the quantum computer Aurora with its four server racks (pictured *below*), they are traditional size server racks, they are seven foot, that wouldn't look out a place in a normal data center. Aurora was published about a year ago in a **Nature paper**, The key thing is was a couple of things. We showed real time decoding, so we did it within one megahertz. We also showed all the parts and subsystems needed for error correction and fault tolerance. But the big thing was, is we showed how to network together individual quantum server racks. Each server rack would have a small quantum computer and you'll have hundreds of these that will be networked together.

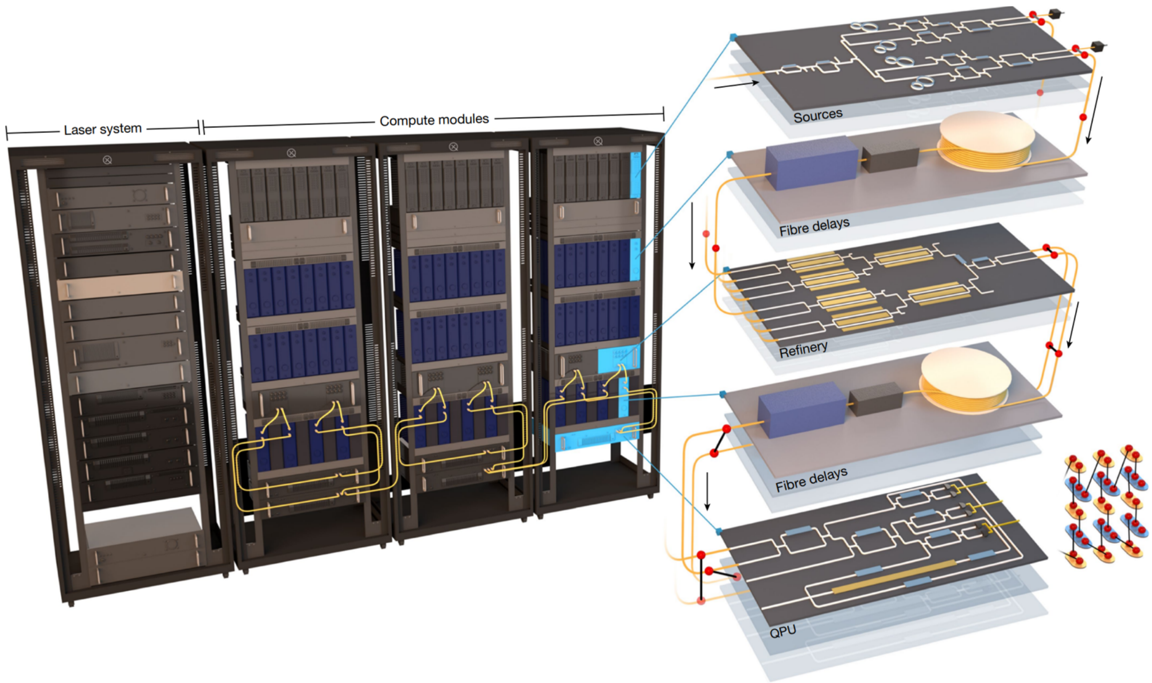

But what was really cool about that is you have four server racks. We could have gone up to forty or four hundred. The scaling only matters is what's nearest neighbor, so, which server rack is near you? So if you have one here, one either side, you only need to connect the nearest neighbors. You don't have to connect the first one to the four hundredth one because that would've been terrible with the loss. The nice thing about it is the ability to network because everyone's going to have to do it in some way. Getting back to this idea, Photonics allows you to do it much faster than others, just because you don't have to convert qubits. You can just keep everything photonic based.

Olivier Ezratty: in the Nature paper, the sizing of what you would need to scale was kind of enormous. It mentions thousands of fracks would be needed. Tell us about that because I presume that you have a lot of things in your roadmap to improve losses and efficiency all along the path of the photon to get down to a more reasonable number of racks.

Christian Weedbrook: the four server racks are really just that we didn't care about optimization or shrinking any parts and sort of densifying things. Our roadmap will really lead to having hundreds of logical qubits in about three to five tennis court sizes, so around that, which is very different to say, other photonic companies, which are talking about one and a half football field sizes for the same computational power. And so, it would have taken much more, but it was only an MVP, a sort of a minimum viable product demonstration. But as we scale up and densify, have lower loss chips, we can reduce redundancy. We're aiming for around three to five tennis court sizes for hundreds of logical qubits.

Olivier Ezratty: how many racks do they make?

Christian Weedbrook: it'll be a couple of hundred. That's what we're aiming for.

Olivier Ezratty: It's still a lot, no?

Christian Weedbrook: it depends. The data centers can have ten thousand server racks, so it really depends on what problems we're solving. It's probably like a small to medium sized data center compared to others.

Olivier Ezratty: I've got my other hat as a cofounder of the **Quantum Energy Initiative**. Have, you estimated the power consumption of such a setting at least in their initial version?

Christian Weedbrook: we have. With hundreds of server racks, we anticipate for our photonic system to have whatever power is being used by hundreds of server racks.

Olivier Ezratty: 30 kW is what I guess for a rack full of DGX servers roughly.

Christian Weedbrook: yes. We're thinking about tens of megawatts. The hope of quantum computing is one of these data centers could replace hundreds or thousands of data centers for very big problems. So, but that's really the back of the envelope sort of estimate. It would be roughly whatever it is for a normal classical one. But with the idea that it's solving bigger, bigger problems than what that could have at that same size.

Olivier Ezratty: do, you know how many HPCs in the world consume more than ten megawatts right now? It's ten! That's why I asked the question, because if if the market size of this kind of QPU is kind of similar to the largest supercomputers in the world, we are in bad shape.

Fanny Bouton: It depends what problem they can solve.

Olivier Ezratty: it boils down to better understanding what are the key technology investments, particularly in manufacturing and silicon, silicon and other components quality to reduce the footprint So when I saw your paper in on Nature last year. It did show that there were some improvements, particularly on losses, and which is typical with photon qubits. Do you expect some gain with better silicon or better semiconductors to reduce this footprint, which is very high?

Christian Weedbrook: you know as thinking about it like a physics sort of problem. If you take the limit of zero loss, that obviously doesn't exist. If we had zero loss, then we wouldn't need any error correction or fault tolerance. So, it's really that thinking because the more loss and errors we have, the bigger the system has to be because of this what I mentioned. So, as you tend towards lower losses, you need less resources to correct for the errors that come from those losses.

Olivier Ezratty: it boils down to what you call encoding rate. So, you need a lower number of CV photons to encode your data.

Christian Weedbrook: exactly, the encoding rate is another key thing, and that's why qLDPC codes are interesting.

Olivier Ezratty: when you mentioned the size of system, you talk about logical qubits. Can you account for the size of the circuit in a classical way? Because we now speak in megaquops or teraquops or whatever. Are you accounting the size of the circuit as well as with a solid-state qubits or in a different way?

Christian Weedbrook: what we say our end goal is : is we talked about hundreds of logical qubits that would be billions of gate operations, and then operating anywhere from ten to the minus twelve to fifteen error rates. Okay, that would be a hundred to one encoding rate of the order of magnitude. So, that's kind of all that gets put together as kind of the final complete system that we're aiming towards.

Olivier Ezratty: it covers well the expected applications we have in quantum chemistry, for example, and finding new ground state of whatever molecule or compound.

Christian Weedbrook: exactly. I don't know if two hundred or three hundred logical qubits are enough, But just getting, I wonder what things can be discovered once we have these large scale quantum computers as well. You see AI now, they don't care whether there is a speedup proof or anything that is getting their hands dirty.

Olivier Ezratty: they have the hardware and it works, and they have hundreds of millions of users. That's a different thing.

Christian Weedbrook: exactly. If we can get people using this and understanding, I think we can see the exponential growth as well. Just basically stuff that we didn't even think of. But you need to actually have access to these computers.

Olivier Ezratty: it's very it's highly connected to the improvements in algorithm design. They are showing that the resources required to do so, I mean chemical simulations like for the femoco ground state calculations, is decreasing over time. I remember a couple of years ago. We were with ten power twelve operations, and we had a couple thousands of logical qubits. And now it's about one and a half thousand logical qubits. Yeah, that's getting down. So, it seems that there is more progress in algorithms than in the hardware right now. I don't know how far it's going to go down in the requirements, but there's another thing which is the progress that's not as stellar in optimization algorithms where it seems to be harder. So, are you interested in that? Or are you focused on quantum chemistry like what your friends from Qolab and John Martinis are trying to focus on? What's your views on the kind of solutions that those systems are going to enable?

Christian Weedbrook: we have a team of algorithm folks that focus on quantum chemistry, material design, specifically with next generation batteries. So if you had a large enough quantum computer, how do you work on solving problems and creating new materials? So we definitely work on that. It's, a lot of the results are still the same algorithms, you'd use for, you know, anything else like chemistry, like pharmaceuticals and so forth. On compilation and optimization, we have a we have **PennyLane** and as at a subset of that, we have **Catalyst** which is part of the compilation side of things. We spend a lot of time on optimizing. That's really a lot of the work when you work with car companies or whoever else. A lot of the time is understanding where their problems are and then working on compilation techniques to really optimize the circuit and bring down the costs of how many T gates do you need. How many operations do you need in total? So that's definitely a big part of it for sure.

Olivier Ezratty: do you have a resource estimator in your software stack?

Christian Weedbrook: that's part of PennyLane and Catalyst. Catalyst does that,

Olivier Ezratty: it's kind of specific for Xanadu. If I put aside what IBM has been doing for a while with Qiskit, you are a very special company because PennyLane is very popular. When you look at the **Unitary Foundation charts**, you are in the top five. So why is that? Explain us the story behind that success with developers.

Christian Weedbrook: we do get a lot of good feedback that people really just love PennyLane, which is always awesome to see. You know, I think it comes back. We started this nine years ago, very early on in the company's lifetime. So, we've had many years to build a great product and get feedback from people. So, there is a longevity aspect that you know, the longer you've been around. If you are doing good work, people will notice it and start really enjoying the product. You know early on, we made it agnostic, so it could. It has plugins to a lot of the other major hardware vendors as well. So, if people didn't like photonics, it's like that's fine. If you want to use ion traps, we'll create a plugin. I think those things is built on Python open source, which also attracts people. We've had the **QHack** software developer conference over the years as well. So all of these things we just put together, and people love people internally in Xanadu love working on it. That makes a big difference if you actually enjoy the product you're making as well. It kind of, Comes out as well.

Olivier Ezratty: it rings a bell with me. More than 21 years ago, I was in charge of developer evangelism at Microsoft. I know the trick, which is how do you talk to developers, how do you build a community with your tool? You knew the trick, but not all the startups know about that. Where is that culture coming from? Is there some people you hired or some kind of experience from somewhere?

Christian Weedbrook: a couple of guys that really did PennyLane and pushed it, they came from academia. It wasn't like let's hire some experts to do this. They just love the product and intuitively, wanted to get into as many people's hands as possible. They also read and understood the developer side of things just by reading books and watching what happens in at the time, you know, seeing TensorFlow and you know learning with things like that.

Olivier Ezratty: do you think that the Silicon Valley culture permeates in Toronto? Because you are based in Toronto, aren't you?

Christian Weedbrook: I think so. Like many others, we've been influenced by it and love reading about successful startups. There is a Canadian culture too, but Silicon Valley has kind of permeated everywhere.

Olivier Ezratty: and you have the CDL (**Creative Destruction Lab**) in Toronto.

Christian Weedbrook: yes, and that's been going a while. We were actually in the first cohort and first quantum company in that in 2016. It was before they had the quantum stream.

Olivier Ezratty: have you got some opinions on AI and quantum? I'm sure you use AI to develop your systems, But the usage of quantum to develop AI applications, it's kind of it's a big buzzword, but is it serious?

Christian Weedbrook: I think long term it is serious. Right now, I believe us and others are still working out how it can be used in the near term. Like you said, we use a lot of AI and LLMs just for building our chip design, our circuits, optimize optimization in general. Like it's just but it's like a tool that everyone should be using. We don't see that as a distinction too much. I would say there is a lot of hype between AI and quantum. We have Maria Schuld who's working on, trying to push different directions in different ways away from everyone else. So I think her team is one of the pioneers in this.

Olivier Ezratty: She's always in South Africa? She wrote many papers on the way you benchmark quantum machine learning. She was kind of pushing back on the hype on quantum machine learning, explaining how you can do this seriously. She's kind of anti-hype academic. She's still an academic even though she works at Xanadu, which has an academic profile.

Christian Weedbrook: she's been with us for eight or nine years now and has done great work. And push back at all, you know, it's quite easy. Sometimes you see a direction that everyone else goes into that direction. She's always looked at ways with the AI hype. I kind of think of like a lot of people saying that quantum will be the next big thing after AI. And I look at it that way: AI is just solving you know, quantum is just doing the same thing, so I think sometime in the future people will be asked, "What's your quantum strategy?" You know, and so I don't really get caught up too much with the AI hype. It's like quantum is going to do something great. It's by using quantum and it's another big solution that hopefully will revolutionize hopefully many industries.

Olivier Ezratty: when comparing your with other photonic vendors space, I have of course in sight Quandela because they are next door in France. Some companies like them have multiple approaches, which enables them to have a kind of so-called NISQy approach, either using a GBS or using a dual rail photons or whatever. Would you position yourself as a company, which has still some kind of niche key solution like what you did four years ago, or you're directly reaching out to the FTQC realm?

Christian Weedbrook: GBS was always going to be a subcomponent of our fault, tolerant, quantum computer. So that's why we went through that path. For the last probably four years now, we've only focused on fault tolerant. Error corrected algorithms. So, what can you do with hundreds of logical qubits? And that was, you know, what we pushed a lot on when we were in the, you know, currently in the DARPA QBI program as well. We did investigate along the way if GBS could help in anything, but we didn't find any examples where you could actually convincingly outperform a traditional classical computer. And. So we haven't looked at NISQ era sort of algorithms for many years now.

Olivier Ezratty: is it possible to program GBS with some parameterizing, the phase, the MZIs or the interferometer? There were some experiments done in China by Pan's team. Somewhere kind of inconclusive on what they could do. Quandela is trying to do that with quantum reservoir computing as well. Do you see some outreach there?

Christian Weedbrook: I keep watching Jian-Wei Pan's team's papers. They do say it seems like there are actually applications there. We with 99.99% probability, we don't think so, but if there ever is, we could always capitalize on that anyway. It would be great if there is, but we tried hard and we don't see any sort of value there.

Fanny Bouton: you have an IPO upcoming. How much is your company valued now?

Christian Weedbrook: for sure, we're going public. We're on track to go public end of next month, so getting close now. And the pre-money valuation was three billion dollars. We're doing a SPAC, with a pipe of $275M, one of the largest quantum pipes in history and we have a SPAC trust. So up to half a billion dollars in fresh capital, depending on redemptions for the SPAC, so we'll see how it goes. Everything's on track at the moment.

Fanny Bouton: what do you want to do with the IPO money? We start to see a lot of consolidation and other companies buying startups to consolidate the market. What do you think about what IonQ or others do? And do you have this vision about consolidation?

Christian Weedbrook: it's going to naturally happen. As you said, it's happening already. IonQ is doing some nice acquisitions there. For us, going public does help do acquisitions. It's a cleaner process to do that. We do have a few ideas of possible acquisitions. But I think the main thing is really the funding will go towards achieving our milestones and roadmap to ultimately this large-scale quantum computer.

Olivier Ezratty: if you look at consolidation, we have multiple options there. So kind of if you take IonQ, they want to do everything okay? Yes. Manufacturing, sensing, communication, computing okay? You could have a vertical, Integration, let's say some company buys its upstream or downstream enabling tech components. Another option, which I see, which is part of what IonQ has been doing as well as D Wave or even Google. Is you drop everything you do on your own technology, and you buy another one, which replaces your own stuff. Like let's say if you take QCI vs the D Wave. It's a replacement, and Oxford Ionics and IonQ is about the same. So, you get rid of the old trap, old 1D trap, and you replace it by the 2D microwave driven traps. There are multiple options. You've got a favorite one?

Christian Weedbrook: we don't know yet. It could be component orientated or software orientated. We'll see how much money we end up raising. The stock price ultimately plays a role too. It's easier to do acquisitions if the stock price is going really well, like it was a few months ago for all the quantum companies. So, we'll take it one step at a time knowing that if timing is right there, there may be a few that we have our eye on.

Olivier Ezratty: if you look at the SPAC mechanism, it is followed by an IPO, it looks like magic from the outside. You've got a big valuation, big money. You can use the ATM process to get funding on the fly afterwards. What's the trade off?

Christian Weedbrook: I've been watching a similar sort of thing. The trade-off is that times can always get tough, and so things get more difficult. Having enough cash is really important. ATM (**at the markeet offering**) works if you've got a lot of volume trading and your stock price is going well. With something like quantum computing, it's going to fluctuate. I think the real time is 2029-2030, when in our case, we have a large-scale quantum computer. You know times aren't always good. And bad times aren't always bad either, so I will see.

Olivier Ezratty: how about dilution?

Christian Weedbrook: it's pretty similar. But you just like to make sure your stock you know, SPACs start at $10. You want to make sure your stock price is above a certain amount for it to make sense as well. So, but very pretty similar. Yeah.

Olivier Ezratty: another trade-off when the company gets public, is it has to communicate on a quarterly basis. It must update its roadmap, then people can analyze you. And sometimes, you have short sellers who do bad reports on you, like with IonQ, three times in a row. So, you expose yourself. It looks like it's leveling the field because there are many more companies being introduced in the stock market. You have Infleqtion right now,

Christian Weedbrook: yeah, and today, **IQM announced** that. These are exciting times. We'll see how it all plays out. But I think it's a good thing to do. There are there is more stuff to do as a public company. But the capital that you can raise is enormous at this point in time, at least.

Olivier Ezratty: you are going to be the second company from Canada after D Wave.

Christian Weedbrook: actually D Wave's American. They became an American company just before they went public. They're an American company now. The public company is technically American. So, we'll be the first Canadian company to be a Canadian publicly traded quantum company. We're also being on the Toronto Stock Exchange as well, which is kind of cool.

Olivier Ezratty: on top, is it on the NASDAQ as well?

Christian Weedbrook: NASDAQ and TSX indeed.

Fanny Bouton: thank you very much Christian. We close this 85th episode of Decoded Quantum and see you soon.

Cet article a été publié le 16 mars 2026 et édité en PDF le 16 mars 2026.