



Opinions Libres

le blog d'Olivier Ezratty

Decode Quantum with Jay Gambetta from IBM

In our third Decode Quantum episode in English after **Simone Severini** from AWS and **Tommaso Calarco** from Julich, we are with **Jay Gambetta** from IBM. And he welcomed us since we recorded this episode near his office at IBM Yorktown Heights Research lab in New York state. This is the 68th episode of Decode Quantum. I spent a week in the area visiting IBM and other companies in the quantum ecosystem along with Fanny Bouton. This episode is also broadcasted on **Frenchweb**.



Jay Gambetta is a quantum physicist. Born in Australia, he did his thesis there at Griffith University in a quantum foundations theme. He then worked on superconducting qubits as a post-doc at Yale University and the Institute of Quantum Computing of Waterloo University in Ontario, Canada. He then joined IBM in 2011 and became in 2019 the VP in charge of all things quantum computing: hardware, software and business development. He is also an American Physical Society fellow, an IEEE fellow, and an IBM fellow.

Here are the topics we covered during our discussion.

How did he land in **quantum science** during his studies in Australia? He started with doing an undergrad in laser science which he found to be cool. He decided to do his undergrad “honors” in quantum physics with shooting lasers on atoms and chose a topic in quantum foundations at the intersection of physics and mathematics.

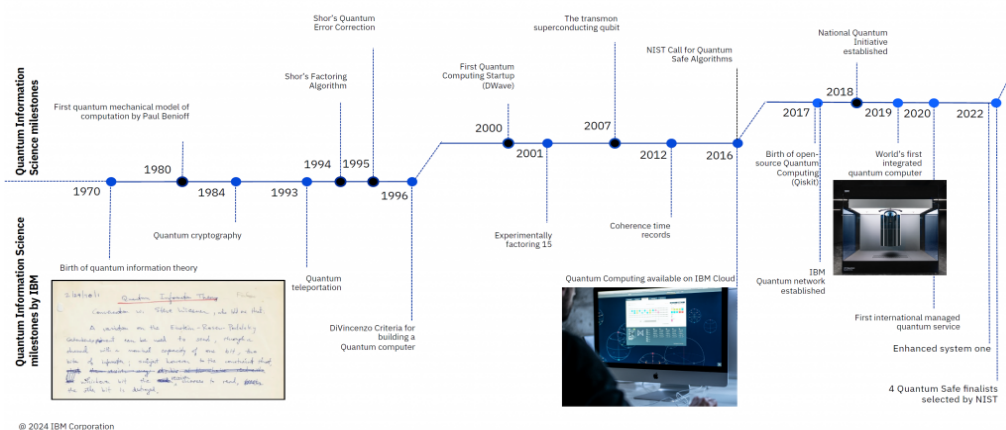
His **thesis title and topic** was **Non-Markovian Stochastic Schrodinger Equations and Interpretations of Quantum Mechanics**, 2004 (211 pages). He was looking at the potential need for a non linear Schrodinger wave equation which landed him looking at quantum trajectories and the likes. Such a debate started in 1994 about whether the world could be Markovian, with some memory. His work, however, remains an open envelope, it didn't solve the debate. He did his thesis with Howard Wiseman after having followed a course on stochastics differential equations. He didn't believe he could be a theoretical physicist. He had to prove this after being an experimentalist.

The link between his thesis and what he is doing now at IBM. It helped him understand quantum optics and it happens that superconducting qubits are about quantum optics and quantum electrodynamics. It couples microwave pulses which are a category of photons and artificial atoms made of superconducting current flowing through Josephson junctions.

His tenure as **post-doc in the "Yale gang"** which saw the creation of circuit quantum electrodynamics (cQED) in 2003-2004 by Andreas Wallraff, Alexandre Blais, David Schuster, Jerry Chow, and many others, under the auspice of Michel Devoret, Rob Schoelkopf, and Steve Girvin.

He then worked in **Canada** (Waterloo) with Joseph Emerson, Frank Wilhem-Mauch and Raymond Laflamme, working on benchmarking quantum systems.

The long history of **IBM** in quantum, which started in the 1970s. Charles Bennet defined the whereabouts of reversible computation in **Logical Reversibility of Computation**, November 1973 (6 pages) and worked on quantum teleportation. He also created the first "prepare and measure" QKD protocol in 1984, along with Gilles Brassard. IBM created a first quantum computing system using the NMR technology, which factorized the number 15 using Shor's integer factoring algorithm, with the participation of **Matthias Steffen** from NIST and UCSB, and who works at IBM since 2006. They then worked on flux superconducting qubits. They published a first paper around the 2012 APS march meeting. See **Superconducting Qubits Are Getting Serious** by Matthias Steffen, December 2011. Until 2015, all of this was pure research. Then, they launched the first multi-qubits QPU in the cloud in 2016 so that many people outside IBM could test it.



The history of **IBM's Yorktown Heights** research lab. The NMR qubit group was in Almaden, south of San Jose in California. It was a testbed of qubits, with some lessons learned on how to build functional qubits. Then, we cover the progress in stability time for superconducting qubits which was in the nanoseconds range with Yasunobu Nakamura's experiment at NEC in 1999. IBM recently reached 3 ms stability times (T1), although it is not yet in production.



IBM's state of the art current technology is their **Heron** processor with 133 qubits. It is using tunability in coupler for two qubit gates. It has fast gates, low crosstalk, and high coherence (see Heron's current figures of merit below). Their transmon qubits are not tunable by design, which helps obtain good figures of merit. Heron is an enabler for "quantum utility", the ability to run quantum computing workloads bringing some form of quantum advantage. They are now focused on modularity and scale-out approaches to assemble multiple such chips using microwave couplers (with Flamingo). Regarding quantum utility, the 100×100 challenge became the challenge to use 5K gates with quantum error mitigation (QEM).

ibm_torino

Exploratory OpenQASM 3

i You do not have access to this system in `ibm-q/open/main`. Switch instances or [explore plans](#)

Details

133

Qubits

Status: ● Online

Total pending jobs: 1623 jobs

Median CZ error: 3.927e-3

Median SX error: 3.105e-4

0.8%

EPLG

Processor type ⓘ: Heron r1

Version: 1.0.5

Median readout error: 2.050e-2

Median T1: 163.88 us

3.8K

CLOPS

Basis gates: CZ, ID, RZ, SX, X

Your instance usage: --

Median T2: 126.54 us

The **quantum software infrastructure** built around Qiskit, which helps programming and send gates to a quantum computer. It then transitioned to make it more reliable. How to make compilation and transpilation scalable? The notion of composable software and dynamic circuits with measurements done in real time that enable conditional code. Qiskit 1.0 was just launched. It took 7 years to release a "production grade" version. Qiskit is now more stable and reliable. One future extension will be to build in quantum error correction.

Progress with **qubit quality**? Qubit fidelities do not measure qubit crosstalk, thus the use of **errors per layered gate** benchmarking that was proposed in November 2023 which assesses the ability to run a large number of gates and cycles. Quantum volume was the first attempt to create a benchmark. It is important to have good fidelity across all the device and with a minimum standard deviation. It depends on fab quality. To do that on a regular basis is challenging. Quality depends on the Josephson single junction tunnel thickness. It is improved post-fab with using a laser that is shooting the junction. See **Laser-annealing Josephson junctions for yielding scaled-up superconducting quantum processors** by Jared B. Hertzberg, Jerry M. Chow et al, arXiv, September 2020 (9 pages). They also make sure to separate well the control frequencies of neighboring qubits. IBM also looks at reducing the turn-around manufacturing cycle to quickly test new qubit chips. The company

has tested a large number of chips over time. Obtaining 99.99% fidelities in whole QPUs is in their roadmap.

The **learning cycle and choice making**. They initially investigated the use of surface codes using square lattice. It was abandoned. Cross resonance gates made that not possible. It led them to pick LDPC error correction codes that are more efficient with heavy-hex lattice (that enables large scale entanglement) and long range connectivity.

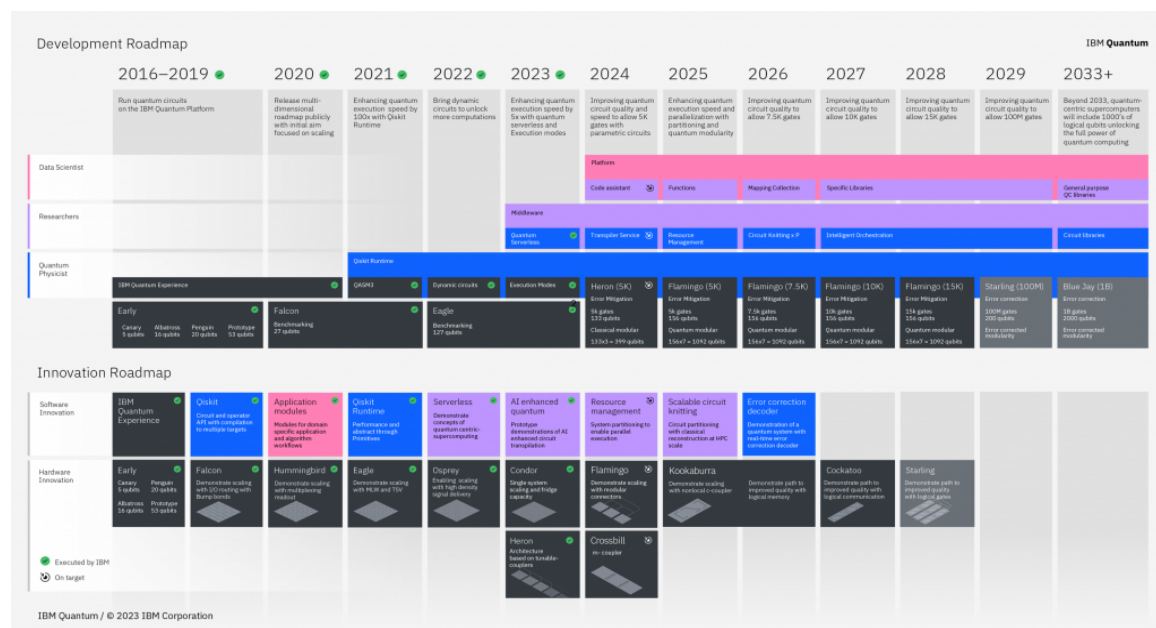
We then discuss about the **quantum utility** pushed around the kicked Ising model paper in **Evidence for the utility of quantum computing before fault tolerance** by Youngseok Kim et al, IBM Research, RIKEN iTHEMS, University of Berkeley and the Lawrence Berkeley National Laboratory, Nature, June 2023 (8 pages). The paper was followed by papers from many research labs describing various classical simulations of the circuit: **Efficient tensor network simulation of IBM's Eagle kicked Ising experiment** by Joseph Tindall, Matt Fishman, Miles Stoudenmire and Dries Sels, PRX Quantum, June 2023-January 2024 (16 pages), **Fast and converged classical simulations of evidence for the utility of quantum computing before fault tolerance** by Tomislav Begušić et al, Caltech, August 2023 (17 pages), **Effective quantum volume, fidelity and computational cost of noisy quantum processing experiments** by K. Kechedzhi et al, Google AI, NASA, June 2023 (15 pages), **Simulation of IBM's kicked Ising experiment with Projected Entangled Pair Operator** by Hai-Jun Liao et al, China, August 2023 (8 pages), **Efficient tensor network simulation of IBM's largest quantum processors** by Siddhartha Patra et al, September 2023 (6 pages). Jay explains us what these papers didn't get well. IBM is pushing two things:

- Find interesting problems to be solved on current systems (≥ 127 qubits with quantum error mitigation and the use of constant depth dynamic circuits), which requires some algorithmic research. The kicked Ising model used in the Nature paper is not the only one possible. You can have some symmetry in the problem but not too much, otherwise it could be classically simulated. The scope of these algorithms is not determined yet, thus the absence of quantum advantage claims from IBM.
- Execute the algorithm on a QPU of large enough scale and obtain a reliable result that is not classically emulable.

The mix of these two is a “quantum utility”. So far, 13 different demonstrations were done seen in multiple papers, and not restricted to a kicked Ising model, but in fundamental physics research and not yet industry-grade use cases. Chemistry simulation is the hard part since fermion to qubits mapping is usually expensive. See **Uncovering Local Integrability in Quantum Many-Body Dynamics** by Oles Shtanko, Zlatko Mineev et al, July 2023 (8 pages) on spin lattices simulation with 124 qubits, **Scalable Circuits for Preparing Ground States on Digital Quantum Computers: The Schwinger Model Vacuum on 100 Qubits** by Roland C. Farrell et al, August 2023 (14 pages), **Quantum Simulations of Hadron Dynamics in the Schwinger Model using 112 Qubits** by Roland C. Farrell et al, January 2024 (53 pages) running on Heron with 133 qubits, **Quantum reservoir computing with repeated measurements on superconducting devices** by Toshiki Yasuda et al, October 2023 (12 pages) and **Machine Learning for Practical Quantum Error Mitigation** by Haoran Liao, Zlatko K. Mineev et al, September 2023 (11 pages) on using machine learning to improve QEM.


On **IBM long term roadmap**, Jay discusses the way to scale out QPUs with assemblies of qubit chips and different interconnect solutions. Eagle allowed TSV (through silicon vias) and multilevel wiring. Osprey and Condor were about I/O scaling. Heron was about implementing a new two-qubit gate (CZ) and a tunable coupler. Modularity is important. IBM plans to use couplers to interconnect qubit chips: the M coupler (for short-range chip to chip connections, with Crossbill), L (long range of about one meter, with Flamingo) and C (cross coupling, with Kookaburra). It enables the implementation of qubit blocks with a certain number of

logical qubits. Blocks are connected with other blocks. They are working on reducing the cost, footprint, and energy consumption of qubits control. One solution being to get away from the (energetically costly) FPGA and move on ASIC controls.



On **IBM customer outreach approach**. Quantum computing as a service. IBM builds a hardware and software platform. He tries to be transparent on wording with no overpromises. IBM puts a scientific tool into the hands of people. In the enterprise segment, IBM targets early adopters who know they won't have an ROI today but maximize their future success. Customers don't want to miss a technology disruption. He believes the wait won't be that long. 2030 could be a reasonable horizon to obtain some business benefit from quantum computers. He advises customers to use the free offer, which since September 2023 includes several 127 qubit QPUs and is not limited like before to <15 qubits QPUs (they are actually all retired). Customers should create a small team on quantum computing to maximize the success of using it when it shows up. Five people is a sweet spot for such a team, on top of some AI investment of course, which can bring shorter term return on investment.

In the end, we found out that a one hour audio format is still short to cover all these topics!



Above: IBM Quantum System Two installed at Yorktown Heights. Another one is installed at the Poughkeepsie Data center. This system hosts three independent Heron 133-qubit chips. The casing can accommodate future systems with thousand qubits given many parts are still empty.

Cet article a été publié le 20 mars 2024 et édité en PDF le 1 avril 2024.
(cc) Olivier Ezratty – “Opinions Libres” – <https://www.oezratty.net>