



Opinions Libres

le blog d'Olivier Ezratty

Assessing IBM Osprey 433-qubit quantum computer

On Wednesday 9th November, 2022, IBM announced the “release” of its Osprey superconducting qubits quantum computer with 433 qubits. It was part of a broad scientific, technology, software and partnership announcement package, as IBM practices about twice a year in the quantum space. These news are well described in **Quantum-centric supercomputing: The next wave of computing** by Jay Gambetta, IBM’s VP covering all quantum activities from R&D to field operations. Here are the videos with the **keynote** from Dario Gil, the 54 mn long **presentation** with Jay Gambetta and Jerry Chow and a **presentation** on quantum error mitigation by Sarah Sheldon.

Osprey will be made available to IBM customers on the Quantum Experience cloud (paid) offering by the beginning of 2023. Current prices of Quantum Experience quantum computers beyond 7 qubits are \$1.6 per second of computing time which is in the range of what some small quantum code may consume.

While this milestone doesn’t yet mean that IBM created a quantum computer with a generic quantum advantage vs classical supercomputers, it represents a technology intermediate step that deserves attention. The team led by Jay Gambetta is probably the largest in the world designing in a fully-stack manner a quantum computer. They build their own chipsets, control electronics, and of course, all their software stacks around the Qiskit framework. In the superconducting qubit realm, they are at this point the most advanced vendor.

This post consolidates various information sources about these announcements. And it is somewhat technical. If you want additional background, you can browse my free ebook **Understanding Quantum Technologies 2022**, and particularly the parts on superconducting qubits and their related enabling technologies (cryostats, control electronics) as well as on the energetics of quantum computing.


What IBM did improve with Osprey?

IBM’s last qubit number record was reached with their Washington system announced in November 2021, based on the Eagle processor containing 127 qubits. The Osprey chipset is over three times larger with 433 qubits but has the same 3-layer architecture than Eagle, with one layer containing the qubits themselves, one layer for the measurement controls, and possibly the related resonators, and one level for the wiring between the qubits. The chipset is installed on a very large PCB board.




To triple its number of qubits, IBM had to improve and scale its control electronics outside the processor. They designed their own “room temperature” custom made control electronics with AWG (arbitrary waveform generators), DAC (digital-analog converters, used for qubit gates controls), ADCs (analog-to-digital converters, used for qubit readout), and a new generation “3” custom-made FPGA (for qubit drive and readout data processing). They also use their own TWPA, these low temperature readout microwave amplifiers used for qubit microwave readout at the 15 mK stage in the cryostat. These amplifiers enable multiple qubit readout multiplexing using different microwave frequencies.

You may wonder why these chipsets have these weird numbers of qubits like 127 and 433. Is it because these are prime numbers? Well, no! It is related to the hex lattice structure that IBM is using to connect its qubits as shown below.




IBM is using its own manufacturing lines in Yorktown Heights (New York State) where they can produce both superconducting qubits circuits and more classical electronics circuits using CMOS technology. The whole Osprey room temperature electronics fit into a single rack which is impressive given Google’s Sycamore 53 superconducting qubits need about 3 racks of control electronics. All of this is about classical electronics, not quantum stuff.


The most striking advance could be perceived as minor but is a very interesting one: they packed the qubit drive wiring reaching the quantum processor into 3 flexible “Cryoflex” cables. It replaces part of the cable forest that you usually see in IBM and Google cryostats, as shown just below in a 2020 generation IBM system of around 20 qubits.



The numbers IBM provided are a 70% increase in wire density and a 5x reduction in price-per-line. When you know that a niobium-titanium superconducting cable costs about \$3K per meter, and you may need as many as 1400 of them to control 433 qubits, you understand the impact in cost structure of this new technology feat. As a result, the space within the cryostat is much cleaner as shown below. These flexible cables seem to be homemade. One independent flexible cable provider is Delft Circuits (The Netherlands) but no sign is (yet) indicating they are involved in this design.



Still, you have about two chunks of about 25 individual cables left, which seem to correspond to the qubit readout signals that exit the qubit chipset as shown above. Why 50 cables and not 433? Qubit readout signals are frequency multiplexed, with about 8 to 9 qubit per cable, thank to large bandwidth amplifiers (the TWPAs). The TWPA amplifiers are not shown in the picture. They are placed below the lower cold plate stage in the cryostat. A cold plate is the large gold cylinder in the above picture. It is made of deoxidized copper covered with a small gold coating that protects the copper from oxidation and also has a very good thermal conductivity.



It is nice that IBM disclosed these pictures of the Osprey “chandelier”. In 2021, the company didn’t publish any picture of its Washington 127-qubit system! Likewise, the various China teams creating superconducting qubit prototype computers with a record of 66 qubits with the Zuchongzhi One system have never published any picture of their chandelier or even of their full system.

Why it doesn’t matter (yet, much) for developers?

IBM has not yet published any qubit fidelities or quantum volume data on its new Osprey processor. Either, they have not yet done their benchmark, or the data is bad.

Some data are available, though. Osprey has a T1 of about 70-100 μ s for this first revision R1 ([source](#)). It is not stellar, but still about the triple of Google Sycamore’ T1 which is of 25 μ s ([source](#)). T1 is equivalent to the amplitude stability of a qubit. Another metric is T2 which is the phase stability time of the qubits and is usually smaller than T1. Osprey’s T2 value was not provided by IBM. A second revision of the Osprey processor (R2) is to improve these coherence times further.

Let’s look at their report below as of early November 2022 with the first Eagle 127-qubit generation from November 2021. Two-qubit gate error rate is 8% and readout errors are 3.7%. And the number is different each and every day as they benchmark their system. It means that after just 10 entangling quantum gates, you have less than 50% chance to get a good result. That’s not very useful for obtaining any quantum advantage, let alone have a computing capacity equivalent to the emulation capacity of your very own laptop!

127 Qubits	Status: ● Online - Queue paused	Avg. CNOT Error: 8.159e-2
	Total pending jobs: 1100 jobs	Avg. Readout Error: 3.569e-2
64 QV	Processor type ⓘ: Eagle r1	Avg. T1: 96.47 μ s
	Version: 1.4.0	Avg. T2: 90.41 μ s
850 CLOPS	Basis gates: CX, ID, IF_ELSE, RZ, SX, X	Providers with access: --
	Your usage: --	

The last v3 generation of Eagle processor, labelled “Kyiv”, shows better fidelities of 2.07% for two-qubit gates and 1.42% for qubit readout (at least the day I extracted the data). It is in line with the fidelities of many smaller systems available on IBM Quantum Experience (27 qubits and below). On top of that, with a record T2 of 160 μ s, Kyiv is best in class for a “live” superconducting qubits processor. Still, 2% errors for two-qubit gates is prohibitive, particularly for deep algorithms.

127

Qubits

Status: ● Online

Avg. CNOT Error: 2.078e-2

Total pending jobs: 0 jobs

Avg. Readout Error: 1.426e-2

Processor type ⓘ: Eagle r3

Avg. T1: 301.67 us

Version: 1.0.0

Avg. T2: 160.92 us

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

Providers with access: --

By the way, there are not many papers published with algorithms that were run on Eagle's processor. Developers prefer the latest (online) generation of Falcon 27 qubits processor which have much better fidelities.

On paper, a 433-qubit processor handles a 2^{433} Hilbert space of complex numbers, way beyond any classical system. But to handle this vast amount of data inside the processor during quantum computation, which is labelled a "state vector" with 2^{433} values being complex numbers, you need to use single and two-qubit gates. And most quantum algorithms require requires running at least a similarly large number of gate cycles. And errors add up quickly. With a 99% two-qubit error rate (without error mitigation), you have 1,28% chance to get a good result when using an algorithm with 422 two-qubit gates. With 99,9% it becomes interesting, at 65%. But a number of two-qubit gates is much larger than an algorithm depth since several gates can be executed in a gate cycle.

On a practical basis, both Eagle and Osprey processors enable only a very shallow quantum algorithms as shown by Eagle's Quantum Volume of 2^6 . Meaning, you can run only a quantum algorithm with 6 cycles of gates with 6 qubits to have 2/3 chances of getting a good result. Such small and shallow algorithms with fewer than 20 qubits don't bring any quantum advantage. You'd need to have at least a QV of 2^{55} to exceed the memory capacity of the largest supercomputers (which by the way is not a measurable QC given it requires a similarly powered classical computer to check the benchmark results). You may still run a shallow algorithm with a lot of qubits. Would that be useful? There are some algorithms in that class like VQE (variational quantum eigensolvers) for which it may be interesting.

So, when IBM writes: "*This processor has the potential to run complex quantum computations well beyond the computational capability of any classical computer. For reference, the number of classical bits that would be necessary to represent a state on the IBM Osprey processor far exceeds the total number of atoms in the known universe*", we are still in fantasyland. These comparisons are inconsistent, mixing a combination of states and a number of objects. Take your dining room and the chairs around your table. There is always many more combinations of the way to arrange your chairs than the number of chairs you have, and it scales exponentially with the number of chairs. Comparing the size of a quantum computer vector state with the number of whatever item in the Universe is nonsense and not homothetic.

If and when IBM will produce processors with better than 99% qubit fidelities (99,9% would be better), they will start implementing quantum error correction codes (QECC) using variations of surface codes adapted to their hexagonal geometries. Google is already testing internally a physical qubit logical qubit with a 100+ qubit processor. This logical qubit would be the first with Google having a better fidelity than its underlying physical qubits.

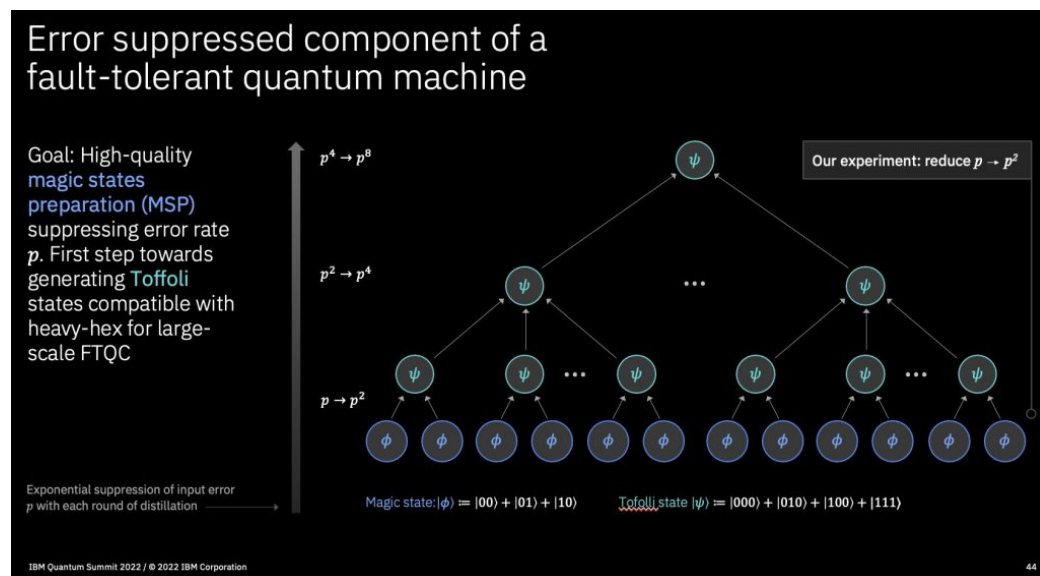
At last, IBM announced it improved its **CLOPS** (circuit layers operations per second) from 1.4K to 15K. But on a different chipset they're working on, a new version of Falcon, which has 27 qubits. Not with Osprey.

Other announcements

Software wise, IBM also announced the integration of error suppression and error mitigation features in

Qiskit. Error suppression happens at the quantum hardware level, using “dynamic decoupling” which modulates microwave control pulses to reduce qubit decoherence and crosstalk (interferences between qubits). This technique was analyzed by two researchers from LIRMM in Montpellier, France, in **Analyzing Strategies for Dynamical Decoupling Insertion on IBM Quantum Computer** by Siyuan Niu and Aida Todri-Sanial, April 2022 (6 pages). These suppression/mitigation control Qiskit features will soon start in beta, with a final release planned in 2025. It is quite a long ride!

IBM also touts “dynamic circuits”, a technique enabling a reduction of circuit depth and adding “threads”, allowing the control of parallelized quantum processors. IBM has indeed many hybrid techniques mixing classical and quantum computing to improve computing efficiency. At this point in time, these methods are difficult to benchmark.



IBM announced many new partnerships, with new academic and education partnerships in the USA, and India, and new customers labelled “industry partners”, like **Crédit Mutuel Alliance Fédérale** in France, and **Bosch** in Germany to develop new electric vehicle materials using less rare earths in fuel cells and electric motors (in the future, and for the future of the future...). In that case, it is about helping the customer start a quantum computing learning curve, create small scale pilot projects, before IBM’s quantum computers scale enough to reach “production grade levels”. According to IBM, this milestone could be reached by 2025.

They also announced a partnership with **Vodafone** in the deployment of post-quantum cryptography solution. At this point in time, this is the best way to generate some revenue although it is not some sort of quantum technology but instead a classical technology that is deployed to get protected against a very distant threat, a quantum computer with over 20 million well entangled physical qubits (and fidelities over 99,9%) that would break RSA 2048 bit asymmetric keys protecting a big chunk of classical communication on the Internet. Fear is a powerful tool to drive incremental business in cybersecurity.

Future IBM developments

Osprey is just another interim step in the long and public IBM roadmap. The next iteration will be Heron with “only” 133 qubits but seemingly, with much better qubit fidelities, leveraging what was learned with Falcon R10 27-qubit processor advertised late 2021, but not yet online in IBM Quantum Experience cloud service. Falcon R10 was said to have a 99,9% two-qubit fidelity thanks to a better control of qubit crosstalk (when operations on some qubits badly influence other nearby qubits) using tunable couplers between qubits (which seems similar to Google’s approach in Sycamore). These fidelities were announced in November 2021 and confirmed later in May 2022.

If IBM could obtain 99% to 99,9% two-qubit gate fidelities with its 133 qubits, it would be a real breakthrough and game changer. And even with taking into account that half of the qubits are usually dedicated to ancilla qubits with most algorithms, it would enable over 60 useful qubits, which may be above the quantum advantage threshold. This may happen in 2023. Then, with such fidelities, they can start creating logical qubits with surface codes if the fidelity is stable as the number of qubits grows. Logical qubits would have better fidelities than their underlying physical qubits and enable the execution of larger algorithms.

But here and now, in November 2022, IBM announced that Falcon R10 quantum volume reached 512, meaning it can efficiently compute with 9 qubits and an algorithm depth of 9 qubit cycles. It is fine but not stellar, and not consistent with 99,9% or 99,99% two-qubit gate fidelities. This is not very good news. It is still better than the qubit volume of 64 (6 qubits) of Eagle but lower than what can be currently achieved with trapped ion systems from IonQ, Quantinuum and AQT, who are all reaching 20 qubits with high fidelities (caveat: it doesn't scale well).

Still, IBM marketingized that *“By the end of 2024, we are pledging to offer our clients and partners **noise free observables** of circuits running on **100 qubits at 100 gate depth**”*. (source). It means that they have confidence that thanks to Heron, their quantum volume will dramatically improve, although, if you look carefully at the wording, not up to 2^{100} . Noise free observables are obtained with using quantum error mitigation techniques. The technique is described in **With fault tolerance the ultimate goal, error mitigation is the path that gets quantum computing to usefulness**, by Kristan Temme, Ewout van den Berg, Abhinav Kandala, Jay Gambetta from IBM, July 2022.



The graphic features the text "100 x 100 challenge" in large white font on a black background. To the right, a white text box contains the following text: "100x100 challenge. We want to build a tool, which beginning in 2024, is capable of estimating noise-free observables of circuits consisting of 100 qubits and depth 100 within a day."


This is somewhat connected to a recent paper coauthored by IBM researchers, **Towards Quantum Advantage on Noisy Quantum Computers** by Ismail Yunus Akhalwaya et al, September 2022 (32 pages) also discussed in **Quantifying Quantum Advantage in Topological Data Analysis** by Dominic W. Berry, Ryan Babush et al, September 2022 (41 pages) and contested in **Complexity-Theoretic Limitations on Quantum Algorithms for Topological Data Analysis** by Alexander Schmidhuber and Seth Lloyd, September 2022 (24 pages). It says that some quantum advantage could be obtained with 96 qubits having 2-qubit gates fidelities of 99,9%. We're not far from Heron here!

#QEI
the quantum energy initiative


If Heron is successful, it could not only bring some quantum computing advantage wrt classical computers but also potentially lead to create a quantum energetic advantage in the NISQ realm (noisy intermediate-scale quantum computers). This is an aspect related to the mission of the **Quantum Energy Initiative** that I co-launched with **Alexia Auffèves** (CNRS MajuLab Singapore), **Janet Splettstoesser** (Charlmers University, Sweden) and **Robert Whitney** (CNRS LPMMC Grenoble) along with a first set of **17 partners**. Its goal is to foster a cross-disciplinary research line around the energetics of quantum technologies, particularly quantum computing, with creating a methodology (described [here](#)) and benchmarking tools to optimize the energetic efficiency of these systems. In one [source](#), Oliver Dial from IBM mentioned that the room temperature electronics of IBM's quantum computers will soon switch from FPGA et ASIC and enable significant improvements on power efficiency, with single qubit control power drain going from 100 W (FPGA) down to 10 mW (ASIC). Provided these qubits have good fidelities, this could lead to some interesting quantum energy advantage.

IBM has also been quietly making other interesting inroads in the cryoelectronics front which can help reduce the energetic footprint of their systems, publishing several papers showing some routes they want to test to better scale qubits drive and readout.

- They prototyped a 14 nm **qubit drive cryo-CMOS chipset**, running at 4K. See **A Cryo-CMOS Low-Power Semi-Autonomous Qubit State Controller in 14nm FinFET Technology** by David J Frank et al, IBM Research. At the IBM Quantum Summit November 2022 event, IBM described this component packaging. It will be a fingernail size chipset handling 4 qubits ([source](#)). It is not that miniaturized. You could expect such a chipset to handle a larger number of qubits.




- IBM Research Zurich also created a **SRAM based AWG chipset** also in 14 nm cryo-CMOS. In plain language, it generates the formed microwave pulses that drive qubit gates and qubit readouts (example below). Since these microwave shapes are multiple, they need to be parametrized, thus the usage of a SRAM memory inside these chipsets. SRAM is the type of CMOS memory used in regular processors cache memory. It seems to complement the former cryo-CMOS. See **A cryogenic SRAM based arbitrary waveform generator in 14 nm for spin qubit control** by Mridula Prathapan et al, November 2022 (4 pages). It also runs at 4K and consumes only 2 to 4 mW for AWG and DAC, but at a low 8-bit sampling rate. It is however designed to drive **spin** qubits at >16 GHz frequencies, not superconducting qubits at 5-6 GHz. Why spin qubits? It seems it is IBM "plan B" qubits. This work is done at IBM Zurich and in partnership with the Swiss quantum ecosystem (and funding).



(d) Measured signal in time domain at 4 K


- The same team created a system approach integrating the former component, always specialized in driving spin qubits, with various electronics optimizations such as at the readout microwaves amplification stages. See **A system design approach toward integrated cryogenic quantum control systems** by Mridula Prathapan et al, IBM Research Zurich, November 2022 (4 pages).
- More interestingly, IBM developed a **superconducting electronics (SFQ)** chipset for implementing part of surface code error correction. It is very interesting since it could both reduce the energy footprint and cycle time of error correction, making it more efficient. See **Have your QEC and Bandwidth too!: A lightweight cryogenic decoder for common / trivial errors, and efficient bandwidth + execution management otherwise** by Gokul Subramanian Ravi et al, August 2022 (14 pages).

In the November 2022 announcement, IBM also described what they brand **Quantum System 2** which will accommodate the next generations of quantum chipsets like the 1121 qubits from Condor, that is set to be released in 2023. This system will use a new **Bluefors KIDE** cryostat ([Twitter](#)). This cryostat seems quite large as we can compare it with Cryomech compressors on the right. It seems to be about 1.2 m wide for the vacuum chamber. The stainless steel and aluminum doors are easy to open for loading stuff. It contains 3 dilution units (cooling the lower stage at 15 mK) and 9 pulse tubes (each corresponding to one Cryomech compressor cooling the cryostat at 4K on the right of the picture below, where we see only 4 of them).




IBM presents the Quantum System 2 as a modular offering that reminds us of the IBM 360 saga which started in 1964 with a very modular approach. This “mainframization” of quantum computing leads to very large size units assembling these KIDE hexagonal cryostat that are either surrounded by classical electronics racks (small rectangles in the schemas below) or connected together with long distance cooled microwave links connecting several superconducting chipsets ([vidéo](#)).

Below is a 2-system unit. The items on the left are probably the Cryomech compressors and the GHS (Gas Handling Systems) which manage the helium 3 circulation for the cryostat dilutions.



And here, we have a 3-unit assembly, that would consolidate 3 interconnected Kookabura processors of 4158 qubits. At this point, these are more science-fiction dreams than real stuff.



These are nice plans and roadmaps. It looks like Moore's law could be applicable to quantum computing thanks to the engineering gods. Rock'n'roll! It is not that simple. Again, this will work out well only if IBM (and others) can create high-fidelity qubits (>99,9%) at a very large scale.

Comparing IBM, Google and others

How can we rate IBM's work in the superconducting qubit space when comparing it to other vendors? They clearly lead the pack at this point. They probably have the largest team in the world, which seems at least 2 to 3 times bigger than Google's team (600 vs 200). Not only have they adopted a full-stack approach in designing their systems, but they have a clear roadmap and they are always delivering things on time (given they don't have an official roadmap for qubit fidelities...). On top of that, they have the largest field presence with developer evangelists in many countries, on top of having installed Quantum Systems in locations in Germany, Canada, Japan and South Korea. Qiskit is also probably the most used quantum development framework in the world, way ahead of Google's Cirq, Amazon Braket and Microsoft's Q#.

Quantum computing is probably way more important for IBM than for Google. Google is still toying with its system and mostly internally. It has set some 1/1 partnerships with US universities. But they didn't put their system on an open cloud like IBM. Maybe Google thinks it is premature to do that. But they still have 11 to 31 IonQ qubits in their own cloud offering, mostly since Google is an investor in IonQ. Google publishes some algorithms papers from time to time as well as on error correction and error mitigation, like in the recent **Overcoming leakage in scalable quantum error correction** by Kevin C. Miao et al, Google AI, November 2022 (17 pages) that shows a way to correct a particular type of superconducting qubit error.

How about the others? Rigetti has 80 qubits but with very poor fidelities. OQC has 8 qubits with unpublished fidelities (in AWS cloud). IQM has 5 qubits with unpublished fidelities as well. Unpublished usually equals... very bad. The most serious contenders seem in the short term Alibaba with its fluxonium qubits and in the long term, Amazon and Alice&Bob (France) with their "cat-qubits" whose main benefit is self-correction for flip-errors and a lower overhead to create corrected logical qubits. A&B plans to create logical qubits with only 33 physical qubits (for an error rate around 10^{-8}). This is a potential game changer in the making!

What you have to get accustomed to with quantum computing is that it's a very long horse race with many twists and turns. We're not out of surprised in that field. And I didn't mention all the various skepticism around the advent of useful quantum computers.

PS: here are some additional sources for Osprey announcement coverage: a good and detailed piecemeal

coverage of IBM's announcements in the **Quantum Computing report** by **Doug Finke**, another **coverage**, from **Francisco Pires** at Tom's Hardware which is based on an interview with **Oliver Dial** from IBM, and an **interesting review** with a detailed coverage of error correction by **Paul Smith-Goodson** and **Moor Insights and Strategy**, in Forbes, not the typical media for such a technical approach.

Cet article a été publié le 11 novembre 2022 et édité en PDF le 18 mars 2024.
(cc) Olivier Ezratty – “Opinions Libres” – <https://www.oezratty.net>