



## Quantinuum Helios advances

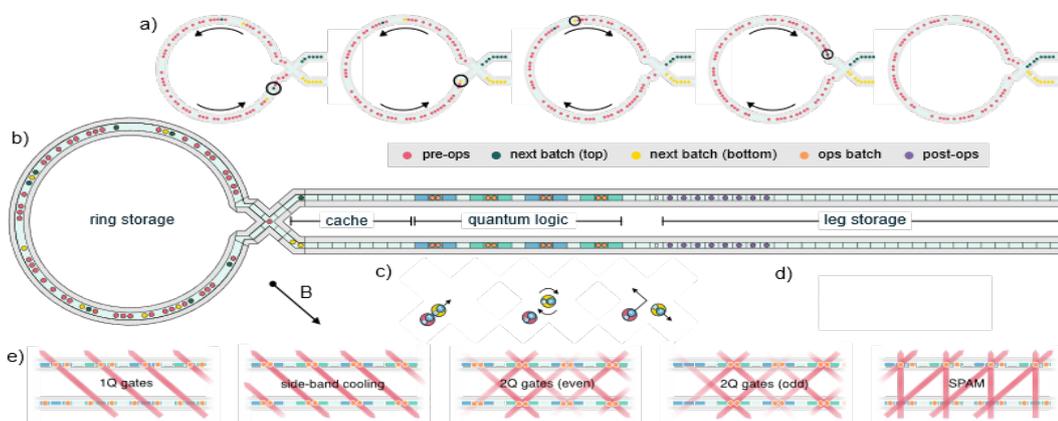
In November 2025, Quantinuum released its Helios QPU with 98 qubits. It also announced it was sold to **Singapore** as part of a partnership with A\*Star and to be deployed there in 2026. Helios got less attention than the recent news coming from Google (**Willow, Echoes**) and Microsoft (**Majorana 1**) but it deserves it. It is one of the best QPUs ever produced so far in the NISQ regime. I will look here at its key disclosed technical characteristics. It's a first for me given I've mostly always published such posts only on superconducting and Majorana qubits.

Helios is fairly well described technically in **Helios: A 98-qubit trapped-ion quantum computer** by Anthony Ransford, M.S. Allman et al, arXiv, November 2025 (25 pages). Most of the information below comes from this preprint paper.

### Helios ion trap

Like all Quantinuum past QPUs (quantum processing units), Helios is using trapped ions. These ions are trapped in vacuum above a semiconductor circuit which contains various electrodes that emit an electric field. This field traps the ions and also drives their displacement. Lasers and light forming devices are used to cool the ions, handle qubit gates and qubit readout.

Helios is using different qubit zones than the previous “race track” H2 system which supported 56 qubits. As shown below, the ion trap contains a circle shaped “ring storage” memory region for qubit storage used as a random-access memory, connected through a four-way X junction to an intermediate cache zone and a gate zone for single, two qubit gate and qubit readout operations, and finally, a leg storage used as a first in, last out memory. The readout operation zone can accommodate for 16 ions and supports 8 two-qubit gates in parallel at a time. So we have in total 98 qubits but can handle two-qubit operations on only 16 qubits at a time.



In Helios, Quantinuum switched its ion elements. The data qubits are now barium ions ( $^{137}\text{Ba}^+$ ) while ytterbium ions ( $^{171}\text{Yb}^+$ ) are used for sympathetic cooling. In previous generations, barium ions were used for cooling while

ytterbium ions were used for computing. As a result, the energy transitions to control qubits are not the same. It means that all the used lasers must have different wavelengths compared to the previous H2 system. Ions positioning is supported by 1,228 electrodes in the ion trap.

### Laser drive

Helios uses at least 7 lasers in total if I count them well:

- 369 nm for ytterbium ions sympathetic cooling.
- a pair of 515 nm lasers with 8 GHz frequency splitting for single gate and Mølmer–Sørensen two-qubit gates ( $R_{zz}(?)$ ).
- 614 nm, and 1,762 nm for state preparation, 1,762 nm being used to avoid measurement crosstalk and leakage errors.
- 493 and 650 nm lasers for fluorescence readout.
- Plus maybe, some laser used for ionizing the barium and ytterbium atoms.

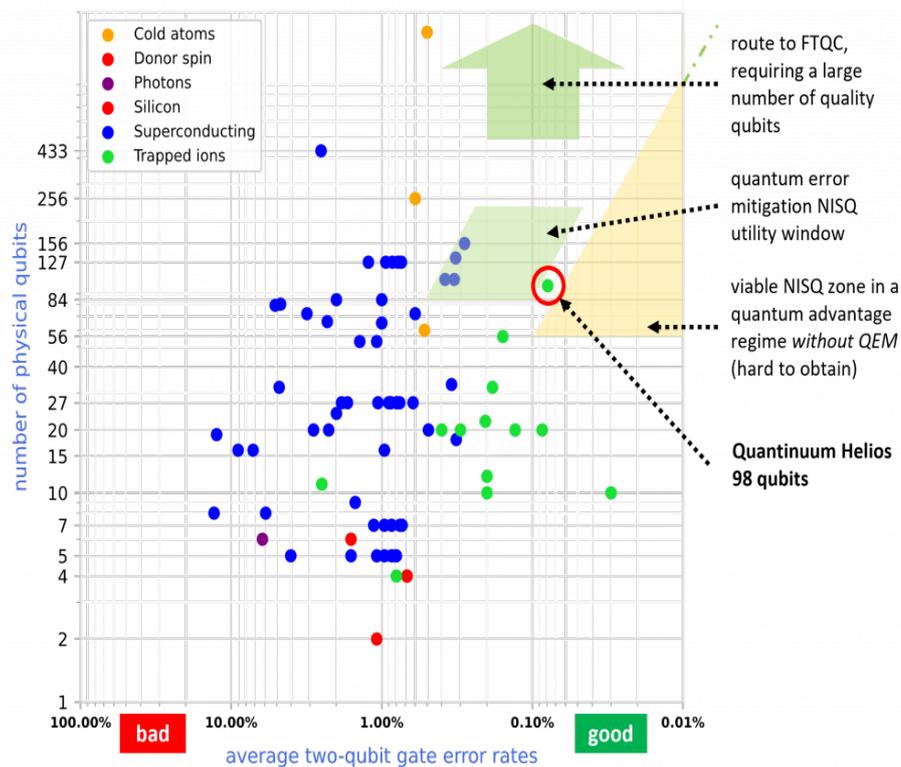
The missing information here is the technology they are using to drive the laser fields onto the ions. You see in the picture above that laser beams are sent in free space on the ions at static positions in the “logic zone” for all active operations (gates, readout). There are not so many zones to control so this looks like reasonable. One can then ask how that will scale with a much larger number of qubits in multiple logic zones.

### Fidelities

Qubit operations fidelities are very good with two-qubit gate fidelities at 99.92% and readout fidelities of 99.952% which is excellent to implement error correction. Single qubit gate fidelities are better, as usual with most qubit modalities, at 99.9975%.

The two-qubit gate used in their benchmark is a  $R_{zz}(?/2)$  gate, which is equivalent to a CNOT gate with applying single qubit gates to the affected qubits. Given the excellent single qubit gate fidelities, it means that synthesized CNOT or CZ gates using this  $R_{zz}(?/2)$  gate and single qubit gates would still exhibit an excellent fidelity.

This positions Helios as having the best qubit count/fidelities balance as shown in my updated fidelities scatter plot chart below. It is nearing the arbitrary viable NISQ zone that I defined a while ago, starting with 56 qubits at 99.9% fidelities, which shrinks as the number of qubits is growing. We can say that this is it with Helios.



It looks it can support quantum circuits with a couple thousand gates, before even starting to use quantum error mitigation. It puts it in a favorable position vs IBM's Heron r3 156-qubit QPU, that a slightly lower fidelities in the bucket of blue dots in the slanted light grey zone, along with Google Willow and China's Zuchongzhi 3.0.

## Speed

Gate operations are quite slow, given the ions shuttling required between zones. It takes about 40 ms to move ions from the storage to the computing zone. The time per layer with 98 qubits is between 40 to 55 ms. 300  $\mu$ s of cooling time is also needed after two qubit gates (IonQ/Oxford Ionics recently **found a way** to escape this). It is true that this chip enables many-to-many connectivity. But at the price of long ion shuttling operations. These shuttling operations are also needed for implementing single qubit gates.

In a **podcast** that I recorded earlier in 2025 with Chris Langer from Quantinuum, he said that *"because we have this tight storage ring, we have higher density storage. It operates probably two to three times faster than H2, for twice the number of qubits."* Also, the good qubit fidelities enable simpler error correction codes which will be faster in runtime and require fewer physical qubits per logical qubits.

Whatever, this low quantum circuit execution speed may become problematic. It's about 100,000 times slower than superconducting qubit gates, which fare between 20 ns and 300 ns in most QPUs. The many to many connectivity is a plus here, but at a serious cost that is not entirely offset by this difference in gate speed. This problem will probably not show up in NISQ circuits execution, but will become a handicap for FTQC when error correction kicks-in with very large circuits, even when being more efficient. You need to do some math to incur for all these aspects and see their practical impact.

## Supremacy experiment

In the recent arXiv describing Helios, Quantinuum conducted its own supremacy randomized Clifford circuit experiment. It which would consume  $10^{73}$  watts on a supercomputer with using a tensor network contraction classical simulation. Their Guppy programming language is used in their tests. This sort of comparison doesn't make more sense that their equivalent with Google Sycamore and Willow since these quantum circuits are not

generating useful data and have no input data at all. When solving a real problem, the gap between classical and quantum computing is not that large, although with such a NISQ QPU, it may be possible to enter a quantum advantage regime for some physics simulations like with some Fermi-Hubbard model simulations, as demonstrated in **Superconducting pairing correlations on a trapped-ion quantum computer** by Etienne Granet, Sheng-Hsuan Lin, and Henrik Dreyer, arXiv, November 2025 (70 pages).

We can still suspect that many interesting experiments will be achieved with Helios for solving various problems, particularly in quantum physics simulations.

Quantinuum also announced it supported 48 logical qubits with Helios, but without any scientific paper. Whatever will appear in such potential paper, you can guess that such logical qubits will increase fidelities only marginally. It may support deeper algorithms based on that. We will see later.

### Form factor

Quantinuum did publish some pictures of its Helios QPU. It is quite bulky for a system supporting “only” 98 physical qubits. It’s got a lot of racks, on the left below, which are probably dedicated to their many lasers control.



The box on the right contains the “physics experiment” with an optical table, probably from Thorlabs. I don’t know what are the other racks made for. Probably some electronics control support for their laser shaping devices which they don’t document well. If this form factor remains the same while Quantinuum increases its number of supported qubits, it will be acceptable. It depends mainly on the number of lasers needed to drive a larger number of qubit gate operations on larger ion traps.



## Power consumption

What is the power consumption of this QPU? Quantinuum is not communicating on it. You need some lasers, control electronics and an ultra-vacuum pump which consumes less than a couple hundred watts. Still, it uses some liquid helium 4 cryogenics. I would guess that the total power consumption is of couple 10 kW.

In the various estimates that I have consolidated for future FTQC QPUs, trapped-ion systems were always in the low range power consumption. It is one of their strengths even though circuit execution time may be much longer due to the shuttling mechanism they are using. That's energy, not power. Energy can be accounted only for the execution of a given task. It is the role of standardized applications oriented benchmarking to define such task.

By the way, Chris Langer from Quantinuum is one of our invited speakers at the **Quantum Energy Initiative workshop** that is planned in Barcelona, Spain, in May 18-22th, 2026 (see them *below*).

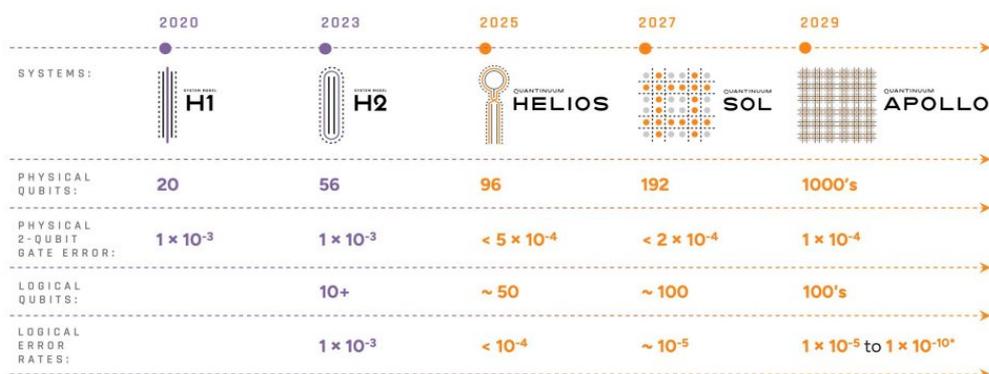
### Invited Speakers:

- Fernando Brandão (Keynote), California Institute of Technology & AWS Center for Quantum Computing (USA)
- Frank Wilhelm-Mauch, Institute for Quantum Computing Analytic, Forschungszentrum Jülich (Germany)
- Gerard Milburn (Keynote), National Quantum Computing Centre, NQCC (UK)
- Juan Parrondo (Keynote), Universidad Complutense de Madrid (Spain)
- Chris Langer, Quantinuum (USA)
- Oleg Mukhanov, SEEQC (USA)
- Yasunobu Nakamura (Keynote), University of Tokyo (Japan)
- Michael Vasmer (Inria)

## Scalability

While Helios qubit fidelities are impressive, you can wonder about how it will scale with using a larger number of physical qubits. Their plans is to release a new QPU in 2027 named SOL. It will host 192 qubits in a 2D frame. We can suppose that it will use multiple zone for gate logic. Quantinuum is quite reasonable compared to many other vendors, with planning new generations every two years instead of every year. They still plan to support tens of thousands of physical qubits and thousands of viable logical qubits on a single chip before having to interconnect multiple traps with whatever photonic links and quantum memory. This is a similar scaling strategy than with IonQ/Oxford Ionics.

## Development roadmap



© 2024 Quantinuum. All Rights Reserved.

\*analysis based on recent literature in new, novel error correcting codes predict that error could be as low as  $10^{-10}$  in Apollo (ref: arXiv:2403.16054, arXiv:2308.07915)

By the way, you can notice that in its 2024 roadmap, Quantinuum was planning to have 96 qubits in Helios. It finally has 98. This is not a big deal given the space they have for storing qubits in their ring.

All in all, Quantinuum Helios is a key milestone for the development of NISQ (noisy) quantum computers. It will enable the execution of larger experimental circuits and will advance the knowledge in the buildout of early stage FTQC (fault-tolerant quantum computers).

Who's next to impress us? Maybe IonQ/Oxford Ionics with their 256 qubit QPU planned in 2026. If it has the same fidelities as their 2-qubit recent experiment, it could further extend the reach of NISQ experiments, while we wait for larger systems and full FTQC support.

Cet article a été publié le 12 novembre 2025 et édité en PDF le 9 décembre 2025.  
(cc) Olivier Ezratty – “Opinions Libres” – <https://www.oezratty.net>