UNIQUE AT SCALE

A roadmap to useful fault-tolerant quantum computing

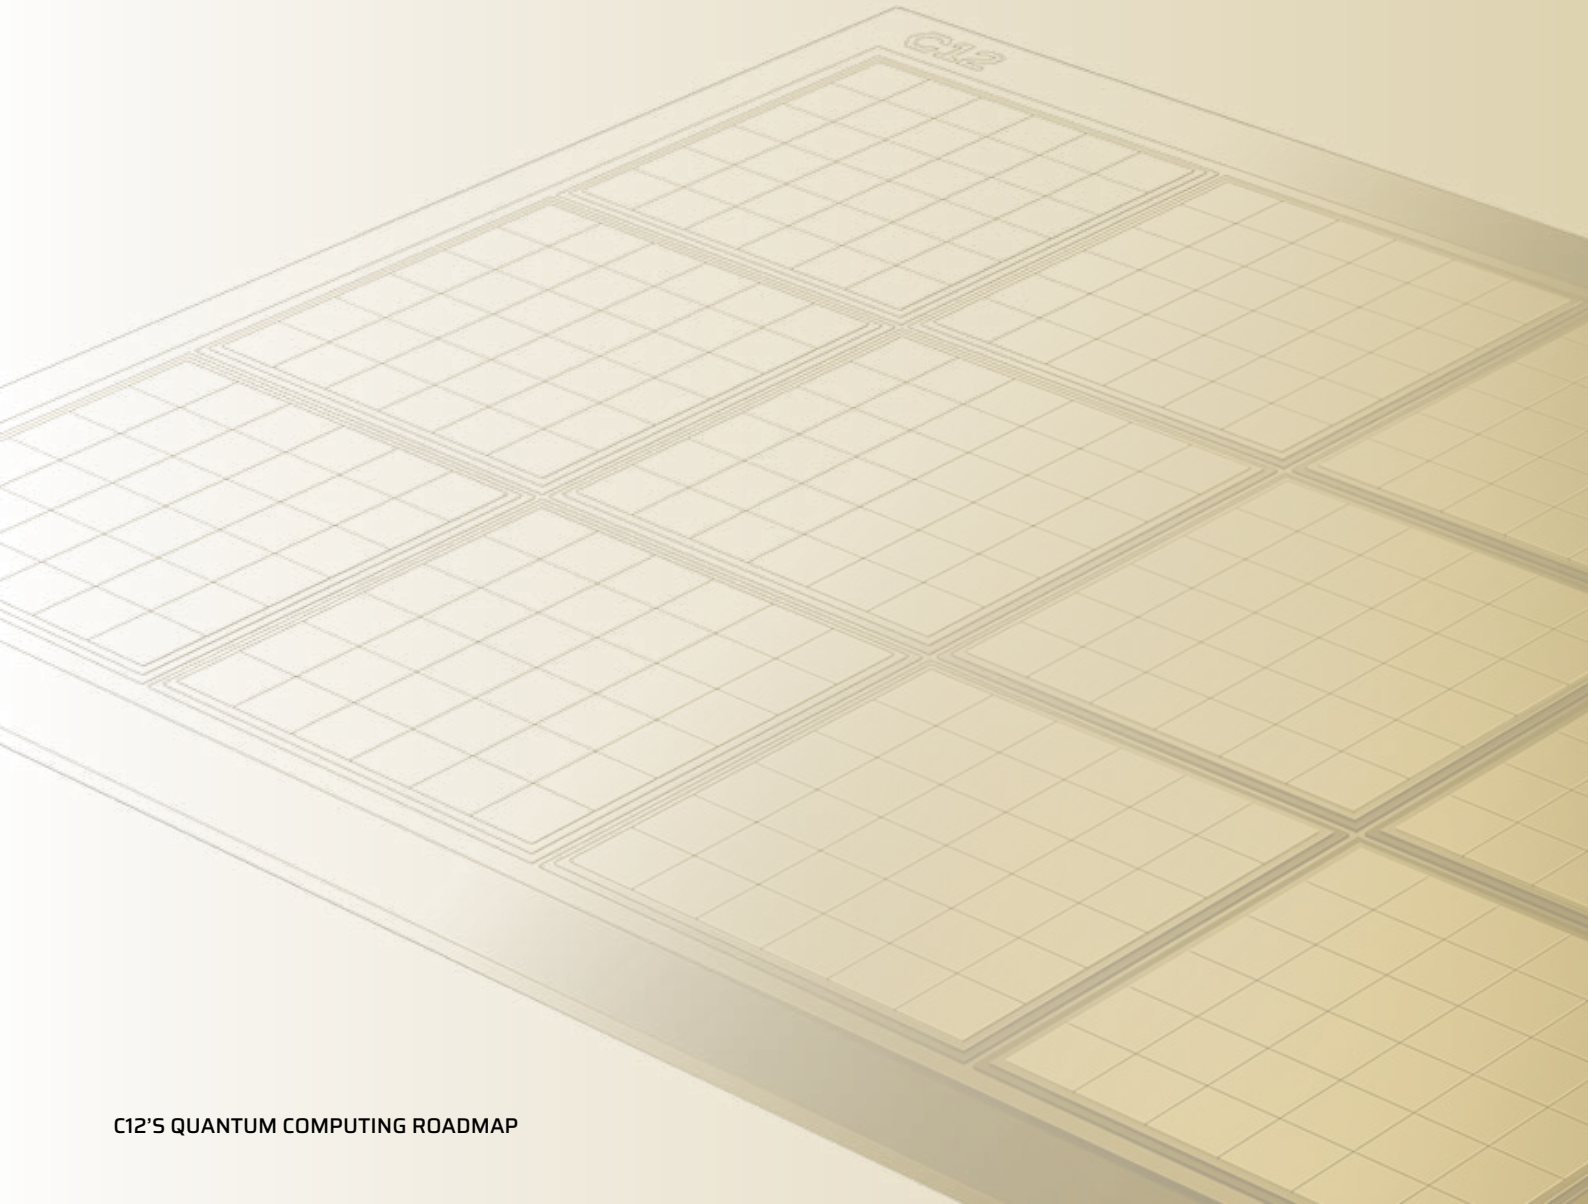

The roadmap

C12	2027	2030	2032	2033
	AÏDÔS	ZÉLOS	STYX	PANOPEIA
	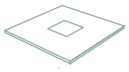		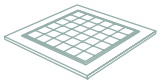	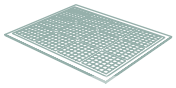
Logical qubits	1	8	128+	792+
Physical qubits	16	236	8,500	100,000
Logical error rate	10^{-3}	10^{-5}	10^{-6}	10^{-7}
Watts per physical qubit	1,500	100	6	0.5
Qubits per square meter	1.4	2.1	500	6,000

Logical error rate applies to Clifford gates.

On-premise delivery within 12 months of first demonstration system.

QPU specifications

Aïdôs 2027

Aïdôs, the Greek goddess of humility. She reflects the discipline and precision driving C12's engineering.

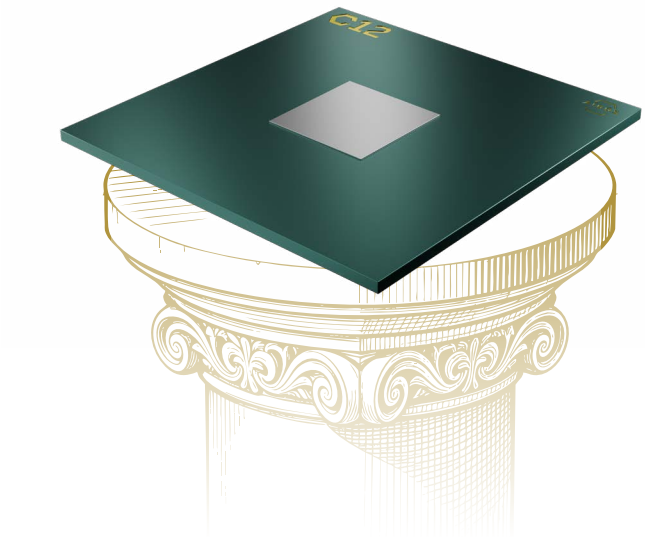

First logical operations

Aïdôs marks C12's introduction of foundational quantum error correction on a compact, next-generation solid-state architecture using spin qubits.

What it introduces

16 physical qubits
1 logical qubit
10^{-3} logical error rate
Universal physical gate set
Sub-µs gate speed
Long-distance resonator couplers
Cloud-accessible system

Zélos 2030

Zélos, the embodiment of ambition and pursuit of excellence, mirrors the drive to scale quantum systems with fidelity and control.

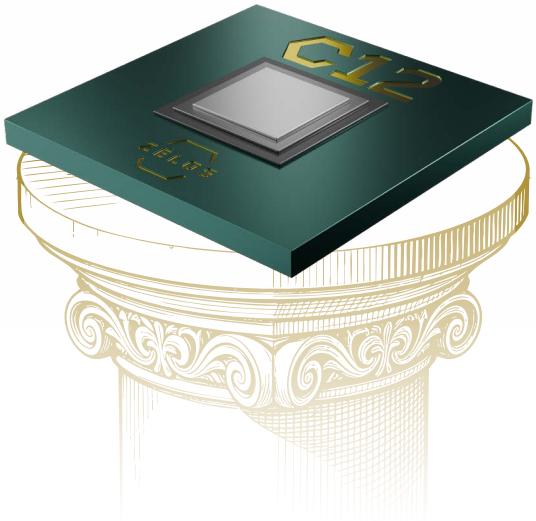

Modular architecture begins

Zélos introduces a chiplet-based architecture designed for replication and system-level scaling. This is when modular integration becomes central.

What it introduces

236 physical qubits
8 logical qubits
10^{-5} logical error rate
Modular chiplet packaging
Cryoelectronics
All-digital control signals
On-chip qubit bias array

QPU specifications

Styx 2032

Styx, the goddess and river dividing life from death, represents the threshold to quantum advantage, when quantum hardware becomes truly resilient.

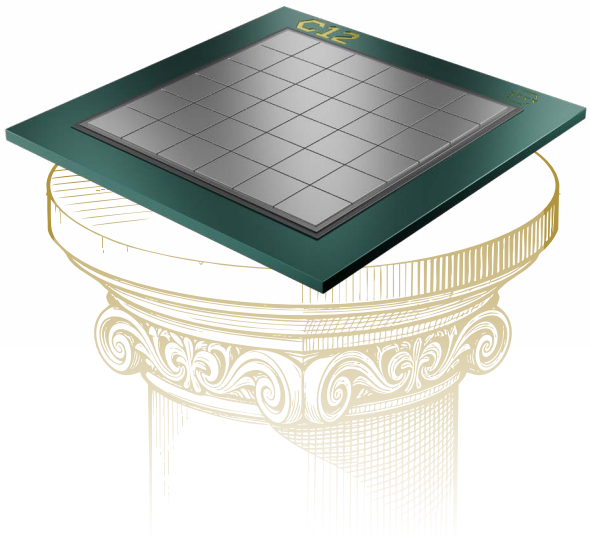

Scaling through replication

Styx combines many Zélos chiplets via inter-chiplet coupling, significantly increasing logical performance while improving efficiency.

What it introduces

8,500 physical qubits
128+ logical qubits
10^{-6} logical error rate
Multi-chiplet module
Major increase in power efficiency per qubit
Chiplet-chiplet couplers

Panopeia 2033

Panopeia, “she who sees everything”, is the moment we reach universal quantum computing.

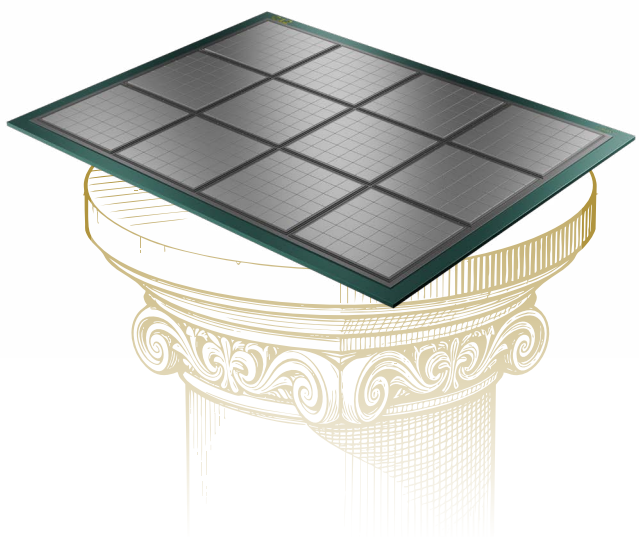

Integrated utility-scale system

Panopeia combines Styx modules to complete the transition to an integrated, deployable quantum system capable of sustained logical computation.

What it introduces

100,000 physical qubits
792+ logical qubits
10^{-7} logical error rate
Multi-module platform
Cross-module couplers
Sub-watt power per qubit at scale
High qubit count, still within a single cryostat

A semiconductor breakthrough opens a new odyssey in computing

The rapid integration of AI into business is placing unprecedented demands on hyperscale data centers to deliver exponential compute growth. But this growth is still limited by classical physics. To reach AI's true potential it will increasingly require the ability to evaluate an enormous range of possibilities at once to identify optimal solutions. This is a class of problems that quantum computing is uniquely positioned to address.

This is where C12 comes in. We are building a useful, fast, compact, utility-scale, fault-tolerant quantum computer. We're doing so by taking advantage of the unique opportunities provided by carbon nanotube based spin qubits. These are its speed, scalability, and real-world viability.

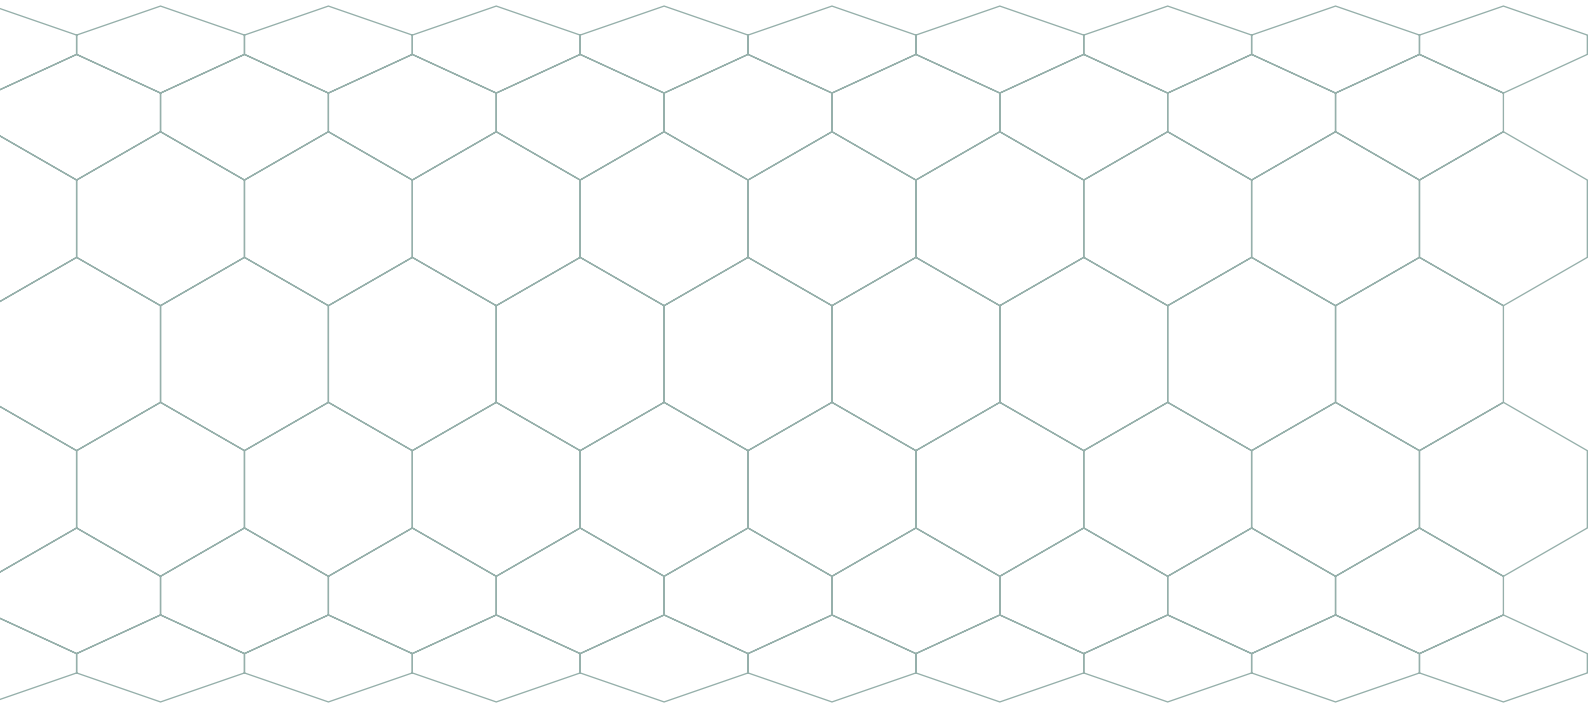

Speed

Speed is unimportant until one finds that the algorithm one wants to run will take five years – then it becomes central. C12 isn't leaving it as a future issue. From inherently fast qubits, to an entangling approach that requires moving nothing, to an architecture that appreciates that both connectivity and parallelization are essential, C12's system is built for speed.

Fast gates

C12's solid-state qubit supports speeds of tens of ns for physical single-qubit gates and hundreds of nanoseconds for physical two-qubit gates, similar to superconducting qubits, the industry speed standard.

Low latency by design

C12's two-qubit gates work by exchanging virtual photons long distances across a quantum bus.[1] No physical atom movement is necessary.

The electrons used as our qubits have negligible movement as well. C12's architecture avoids the need for electron shuttling, an approach some solid-state qubits use that imposes scaling challenges in material uniformity and control complexity.[2]

For utility-scale algorithms, being able to skip a step of moving the qubit relatively vast distances makes an orders-of-magnitude difference in total gate speed.

Efficient computation

C12's quantum bus approach allows local all-to-all connectivity zones that support more efficient algorithms, quantum error correction codes, and gates that entangle many qubits at once.

High connectivity means that within a zone, distant physical qubits can interact without extensive message-passing through a chain of intermediary qubits. That avoids the risk of errors from each step of passing the message along, leaves those qubits free for other operations, and saves compiler complexity.

The ability to adjust the all-to-all connectivity zone size from two qubits to a whole chiplet's worth across a long-distance quantum bus[3] allows C12 to optimize the balance between connectivity and parallel operations. It also unlocks the ability to use multi-qubit gates involving any number of those qubits simultaneously. These provide the flexibility necessary to support not just today's fastest and most efficient error correction algorithms, but also many yet to be developed.

1 https://www.nature.com/articles/s41534-019-0169-4.

2 See, for example, https://www.nature.com/articles/s41467-024-49182-4 or https://www.nature.com/articles/s41467-024-49358-y.

3 https://journals.aps.org/pra/abstract/10.1103/PhysRevA.111.042621.

Scalability

Quantum computers are useful when they solve real-world problems, at scale. The traditional path of incrementally scaling up a qubit lab experiment doesn't work for quantum computer development because one can't achieve 4-5 orders of magnitude scaling in linear increments. An honest evaluation of what an at-scale system looks like reveals fundamentally different technology from that of a lab experiment.

At that scale, the constraints that dominate system design are no longer qubit coherence in isolation, but wiring, control, connectivity, and execution time. C12's design approach centers around scaling, and carbon nanotube based qubits make it possible for many different reasons.

Modular chiplet architecture

C12's system centers on 3D-integrated chiplets to support a modular, scalable approach. Each of our systems after 2027 is designed as a module that can be replicated and used many times in the next larger system. 36 Zélos chips of 236 qubits each are used in Styx, and 12 Styx modules are used in Panopeia.

The benefit of modularity is apparent in the number of systems on the roadmap; the modular design of Zélos requires only two more iterations to get to utility-scale fault-tolerant computing.

Semiconductor fabricated circuits

C12 is able to integrate better with CMOS than most quantum technologies because we leverage modern semiconductor fabrication techniques in our in-house nanofabrication facility and introduce carbon nanotubes only as the last step of the integration process. C12 has spent five years meticulously refining the groundbreaking, patented nanoassembly technology necessary to integrate carbon nanotubes into semiconductor-fabricated circuits.

Carbon nanotubes are first grown atom by atom through chemical vapor deposition with an ultrapure carbon-12 isotope. A non-invasive method is used to characterize nanotubes so that optimal ones can be selected for integration. The selected ones are placed and then electrically fused into the circuits on a fabricated silicon wafer, all while maintaining isolation from sources of contamination.

Does it work at scale? High-throughput pick-and-place technology is already able to precisely incorporate tiny surface-mount devices into circuits at a rate of tens of thousands per hour,[4] and C12 is quite confident its own pick-and-place process will scale accordingly.

4 For example, https://www.maximsmt.com/asm-pick-and-place.

Efficient error correction

C12's qubit connectivity supports error correction codes with at least a 10x lower physical qubit requirement than the surface code.[5] Depending on the code, the high connectivity supports multiple options for non-Clifford gates, such as magic state cultivation or code switching with transversal gates.

C12's product roadmap takes conservative assumptions for error correction. It's based on simulations using our own noise model and state-of-the-art decoders. Aïdôs assumes a rotated surface code[6] with distance-3, weight-4. Panopeia assumes a bivariate bicycle/generalized bicycle code[7] with distance-8, weight-4. Our estimates do not see non-Clifford gate implementation as the limiting factor for the logical error rate.[8]

But these are mere examples. Error correction will be far advanced by the time these systems are available, and the high connectivity of C12's architecture gives it future-proofing to exploit many new techniques for even better results.

Uniform, all-digital qubit control

The requirement for a unique microwave signal to be generated and delivered to every qubit, as is common on quantum computing platforms, presents an enormous barrier at scale due to the wiring and heating challenges imposed.

The solution for this wiring bottleneck is to minimize the amount of information that must be transferred from room temperature for each qubit operation, and rely on cryoelectronics near the qubits to create the signals the qubits need. But the qubits themselves need to possess special characteristics for this to happen.

Experimental validation has shown that carbon nanotube devices allow very fine electrical control of the electron spin qubit's environment.[9] By using an on-chip qubit bias array to retain calibration settings, C12's qubits can be tuned at a millisecond scale to match each other. This avoids the huge burden of customized timing and microwave frequencies specific to every qubit for every gate. Precision is granted by low-frequency updates, allowing very simple qubit instructions to be sent with high-frequency updates.

Through co-integration with cryoelectronics at the level of the quantum chip, high-speed control is achievable sending only digital signals into the cryostat.

High fidelity for efficient qubit use

Carbon nanotube based qubits maximize fidelity by offering supreme noise isolation, especially from nuclear spin noise thanks to the purified carbon-12 isotope used. C12's qubit is not embedded in silicon rich with sources of charge noise like silicon spin qubits, but held within a nanotube far from that surface: a much cleaner interface. The amount of material near the qubit is close to the minimum conceivable, as it sits on a nearly one-dimensional bridge. Since the number of qubits to correct errors drops as qubit fidelity increases, this allows C12 to execute algorithms using fewer physical qubits.

Thanks to these advances, C12 targets 99.99% single-qubit gate fidelity and 99.9% two-qubit fidelity for the physical qubits, and these are used in the error correction assumptions. Likewise, a 99.8% readout fidelity is targeted.

5 A head-to-head result shows 288 qubits for a C12-compatible qLDPC approach versus about 3,000 in surface code for the same 12-logical-qubit memory target (https://www.nature.com/articles/s41586-024-07107-7).

6 https://journals.aps.org/pra/abstract/10.1103/PhysRevA.90.062320.

7 https://www.nature.com/articles/s41586-024-07107-7.

8 For example, see https://journals.aps.org/prxquantum/abstract/10.1103/9kys-3whh.

9 https://www.science.org/doi/10.1126/science.aaa3786.

Real-world viability

Scaling a quantum computer by brute force is conceptually easier, but a whole data center shouldn't be required for a few hundred logical qubits. C12's design is compact and deployable.

Small system footprint

C12 is able to target 100,000 physical qubits (792+ logical) with only a single dilution refrigerator and a few standard server racks. That's an overall footprint of 17 m^2, comparable to a dozen standard server racks, whereas a superconducting or photonic system might be the size and complexity of a data center or factory.

Low energy consumption

Carbon nanotube based qubits lead to affordable system running costs. A relatively high operating temperature in the low hundreds of mK, control electronics efficiencies, and tiny qubits needing low drive power allow C12 to target 100,000 physical qubits within a single dilution refrigerator at a cost of less than a watt per qubit, for a total of 50 kW.

Reduced cooling requirement

C12's spin qubits can be operated at temperatures in the low hundreds of mK,[10] which is an order of magnitude higher than superconducting qubits and supports two orders of magnitude more cryoelectronics. The roadmap is designed based on modern cryogenic system capabilities, rather than relying on future improvements.

Helium-3 efficient

C12's single-cryostat approach requires only a few tens of liters of helium-3. This is a real advantage over systems that scale by adding cryostats; a paper by IBM Quantum, Bluefors and others recently stated that multi-cryostat "Systems of 100k qubits would demand large fractions of the yearly world production [of He-3] based on naive linear scaling."[11]

Reduced maintenance challenge

By relying on one cryostat, long-term operation of a C12 system is straightforward. This avoids the escalating maintenance challenge that would result from scaling by adding cryostats.

10 https://www.nature.com/articles/s41467-025-60952-6.

11 https://doi.org/10.48550/arXiv.2512.15001.

From the classical world into the realm of quantum advantage

This roadmap reflects a singular focus: to build a fault-tolerant quantum computer that can solve meaningful problems at scale.

C12's first-principles approach starts from the constraints of large-scale systems and the architecture is designed accordingly. By prioritizing speed, scalability, and real-world viability from the outset, we define an ambitious but realistic pathway to a genuinely useful
quantum computer.

Conceptual rendering