Unique at Scale

A roadmap to commercial fault-tolerant quantum computing at scale

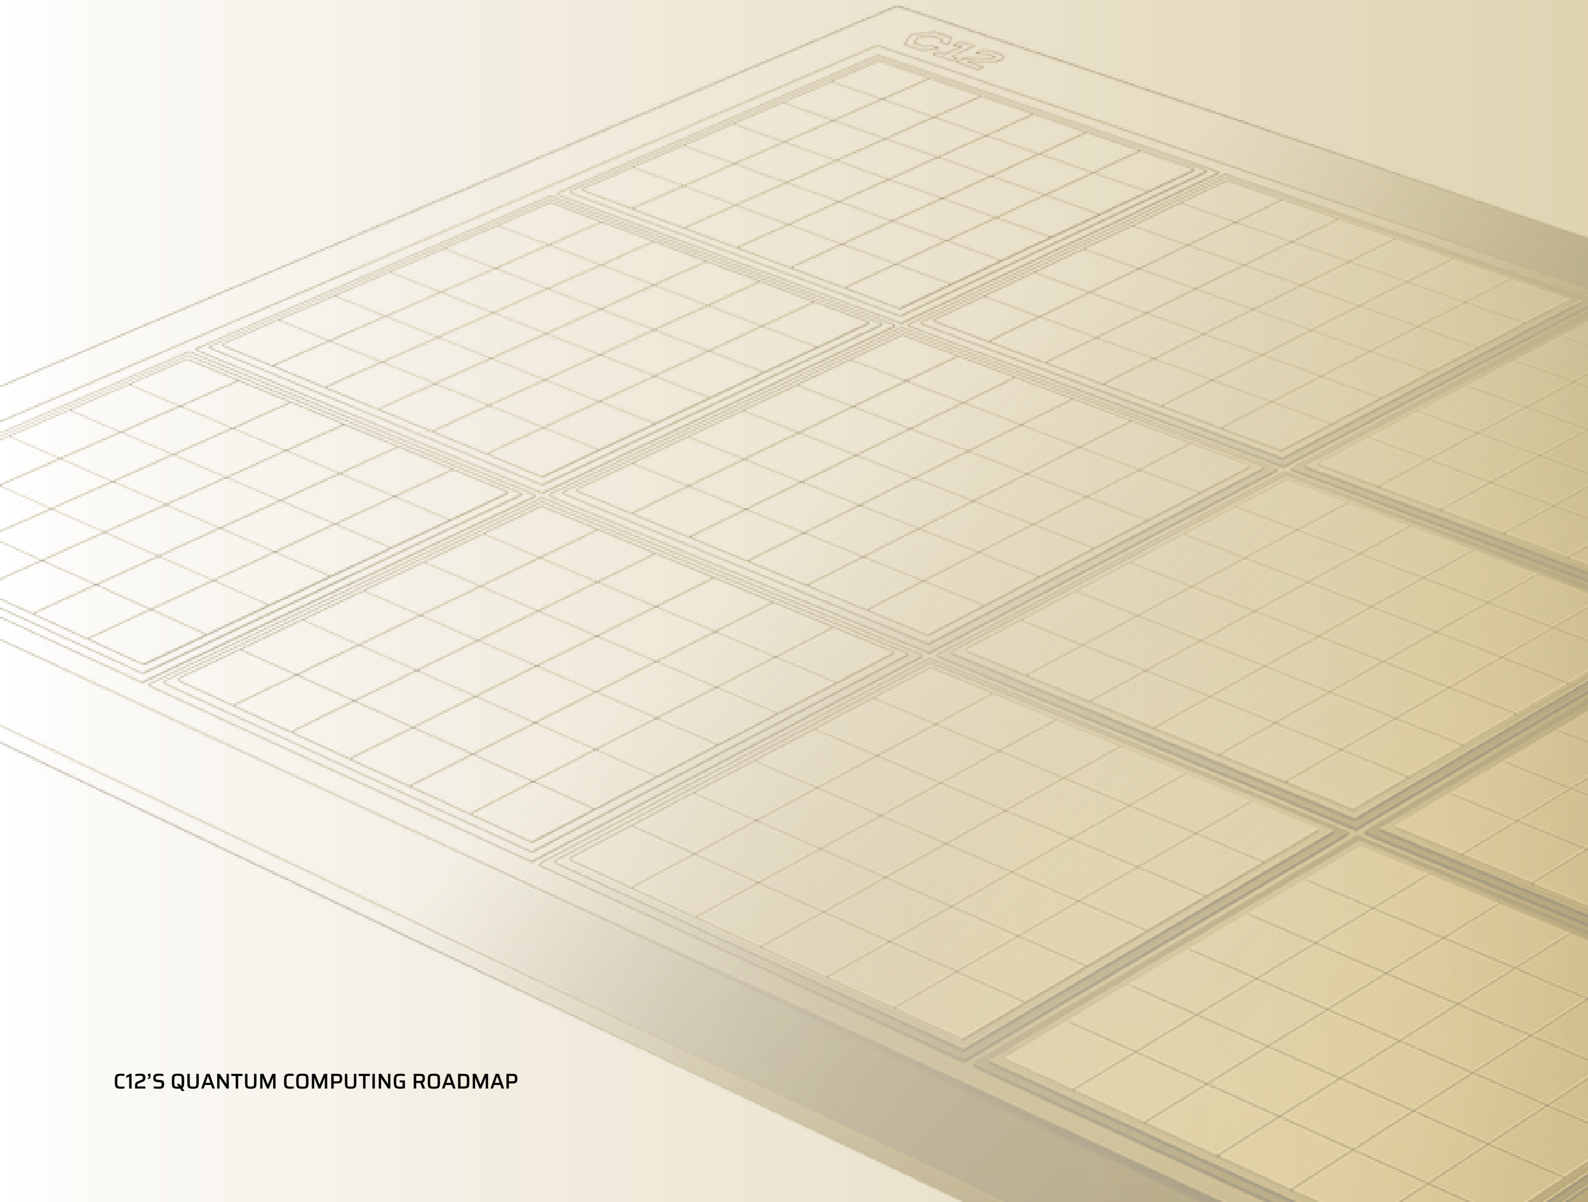

C12'S QUANTUM COMPUTING ROADMAP

SECTION 1

The roadmap

	2027	2030	2032	2033
	AÏDÔS	ZÉLOS	STYX	PANOPEIA
Logical Qubits	1	8	128+	792+
Physical Qubits	16	236	8500	100,000
Logical error rate	10^{-3}	10^{-5}	10^{-6}	10^{-7}
Watts per physical qubit	1500	100	6	0.5
Qubits per square meter	1.4	2.1	500	6000

Logical error rate applies to single-qubit Clifford gates.

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

SECTION 2

QPU characteristics

Aïdôs 2027

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

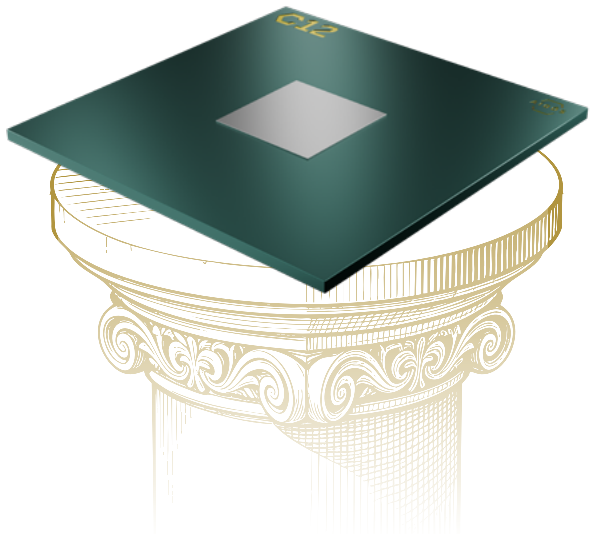

First logical operations

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.

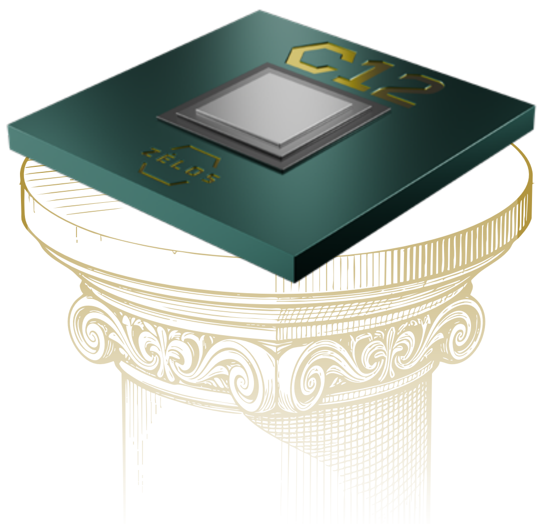

Modular architecture begins

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
- Qubit bias memory network

SECTION 2

QPU characteristics

Styx 2032

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

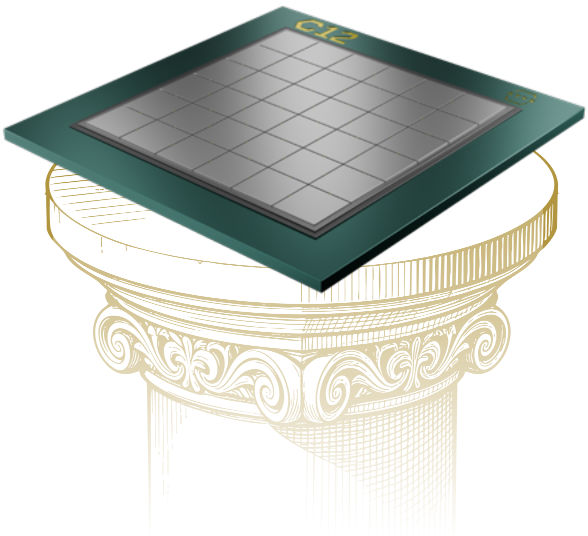

Scaling through replication

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

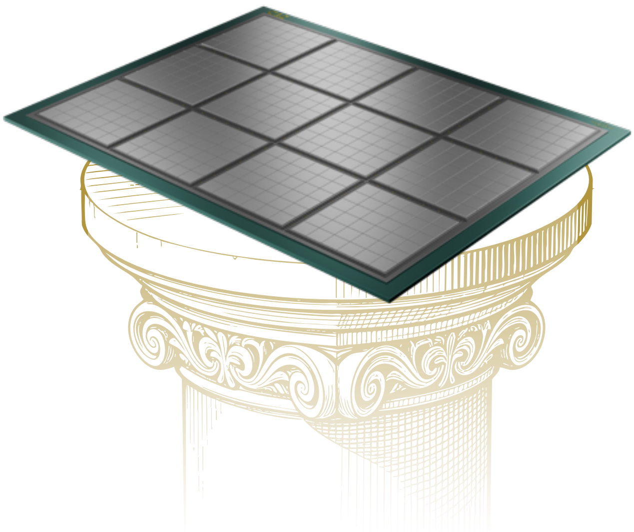

Integrated utility-scale system

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

SECTION 3

A coherent path to useful quantum computing

Quantum computers are useful when they solve customer problems, at scale.

Yet the approach companies in the field tend to take is to incrementally scale up lab experiments. This cannot deliver the 4–5 orders of magnitude required to reach useful systems. At that scale, the constraints that dominate system design are no longer qubit coherence in isolation, but wiring, control, connectivity, and execution time.

An honest evaluation of what an at-scale system looks like reveals fundamentally different technology from that of a lab experiment. C12's roadmap reflects this, with each system introducing the capabilities required for the next stage of scaling.

Introduction

The business world is rapidly becoming dominated by AI, which has led to a race to hyperscale data centers providing more and more compute. But AI remains limited by the tools it has. To reach its true potential, an essential tool it will need is the ability to evaluate enormous numbers of possibilities at once to select the optimal one. Only a quantum computer can achieve this.

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 scalability, speed, and real-world viability.

Scalability

Quantum computers are useful when they solve real-world problems, at scale. Yet the approach companies in the field tend to take is to incrementally scale up lab experiments.[1]

One can't achieve 4-5 orders of magnitude scaling in linear increments. So that is not C12's approach. An honest evaluation of what an at-scale system looks like reveals fundamentally different technology from that of a lab experiment. 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 are used in Styx, and 12 Styx modules are used in Panopeia.

Semiconductor fabricated circuits

C12 is able to integrate better with CMOS than most quantum technologies because we leverage modern semiconductor fabrication techniques and introduce carbon nanotubes only as the last step of the integration process. This is made possible by our patented nanoassembly process. High-throughput pick-and-place technology[2] is already able to incorporate surface-mount devices into circuits at a rate of tens of thousands per hour, and C12 is confident its own pick-and-place process can scale accordingly.

Efficient error correction

C12's architecture supports efficient quantum error correction via code switching using transversal gates, but also supports more standard techniques such as magic state cultivation, remaining flexible for inevitable coming advances in error correction. C12's qubit connectivity can support error correction codes with a 10x lower physical qubit requirement than the surface code.[3]

Uniform, all-digital qubit control

Experimental validation has shown that carbon nanotube devices allow very fine electrical control of the electron environment. C12's qubits can consequently be tuned at a millisecond scale to match each other, rather than impose a huge burden of custom timing and microwave frequencies specific to every qubit for every gate.

This is a mere nice-to-have at small scale, but critical at large scale. The requirement for a unique microwave signal to be generated and delivered to every qubit, as is common on quantum computing platforms, would be a deal-breaker at scale due to the bandwidth and heating challenges imposed. Only simple control signals are required to execute our high-speed gates, since the qubits can be treated as identical. This permits C12 to rely solely on efficient digital control signals entering the cryostat, much more in line with how classical computing has been successful.

1 It's understandable; there's pressure for near-term hardware sales. Likewise, quantum supply chain vendors lack motivation to provide meaningfully scalable components, since merely incremental technology can be sold to dozens of companies, today.

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

3 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)

High fidelity reduces qubits needed

Carbon nanotube based qubits maximize fidelity by offering the best noise isolation of any solid-state qubit, especially from nuclear spin noise thanks their purified carbon-12 isotope. 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. Since the number of qubits to correct errors drops as qubit fidelity increases, this allows C12 to execute algorithms using fewer physical qubits.

Carbon nanotube based qubits are not the only ones that offer high fidelity; trapped ion and neutral atom qubits do as well, because the qubit is isolated far from noise sources. But this comes at a huge cost to speed: the qubits must then be moved around for operations.

Fast electrical connections are what made computers possible and this will be the same for qubits: no atom movement necessary. The optimal way to support this with a minimum of local noise is to host the qubit in arguably the closest thing available to an ideal, one-dimensional pathway for electrical signals: a carbon nanotube.

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 speed of tens of ns for physical single-qubit gates and hundreds of ns for physical two-qubit gates, similar to superconducting qubits, the industry speed standard.

Negligible atom or electron movement

C12's two-qubit gates don't rely on physical atom or electron movement. The need to move atoms to entangle them slows trapped ion and neutral atom systems to such a degree as to make them largely unviable for utility-scale algorithms. While these technologies report "fast" gates, they tend to omit from the numbers the lengthy process of moving two qubits of interest next to each other.

Similarly, the electrons used for C12's qubits remain where they are, with entanglement occuring through fast virtual photon exchange with a microwave resonator. This is more efficient than the electron-shuttling approach relied on by other spin qubits. Shuttling is a slow process requiring a lot of physical gates that create a major challenge in wiring, timing, and coordination. Shuttling electrons also tends to lead them through areas of low valley splitting, where the qubit decoheres.

All-to-all connectivity zones

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.

This avoids the heavy burden faced by nearest-neighbor connectivity designs is which qubits can only interact with distant ones through extensive message-passing through a chain of intermediary qubits. Every attempt to pass a message in this chain is an opportunity for a new error. Meanwhile, the qubits along the way need to be free for use, creating immense compiler complexity when scaling up.

Connectivity and parallelization

The option to either entangle qubits within a nanotube, or with up to 400 qubits across a long-distance quantum bus, allows C12 to optimize the balance between connectivity and parallel operations. Above all, this provides the flexibility necessary to support not just today's best error correction algorithms but those yet to be developed. All-to-all connectivity zones are also common for atom-based qubits, but

face roadblocks in terms of speed as well as technology barriers in terms of photonic chips for neutral atoms[4] and photonic interconnects for trapped ions.[5]

4 Neutral atom systems seem ideal for all-to-all connectivity, since atoms can in principle be rearranged in 3D and entangled in parallel. In practice, the spatial light modulators generating the traps are vastly too slow for real-time rearrangement. Current systems instead move atoms individually using single lasers and acousto-optic deflectors. Scaling will require a fundamentally different approach such as integrated photonic chips, but this remains early-stage research. A challenge will remain that determining optimal 3D routing would itself require a quantum computer.

5 Trapped ions support all-to-all connectivity zones, but limited to short chains of ions. Scaling requires QCCD architectures that shuttle ions between zones, but shuttling is sequential, slow, and introduces motional heating, especially at junctions. Photonic interconnects between zones offer an alternative but entanglement success rates per attempt remain low, making them a parallelization bottleneck.

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. A utility-scale carbon nanotube quantum computer could fit into a volume 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.

Small system footprint

More efficient control electronics and much smaller size than superconducting qubits allow C12 to target 100,000 physical qubits with only a single dilution refrigerator and a few standard server racks. The total system size for 792 logical qubits is expected to be around 17 m². Systems of similar capabilities relying on distributed quantum processors anticipate the need for dozens of dilution refrigerators occupying vastly more space. C12's design does not require lasers, optics, single photon generators, photon transducers or photon detectors, avoiding an enormous equipment and cryogenics burden that even many technologies not based on photonic qubits face in order to scale.

Reduced cooling requirement

Compared with superconducting qubits, C12's spin qubits can be operated at temperatures an order of magnitude higher (low hundreds of mK), which actually supports two orders of magnitude more cryoelectronics.

Reduced maintenance challenge

By relying on one cryostat, long-term operation of a C12 system is straightforward. Quantum computer builders targeting "distributed" computing by scaling up the number of cryostats, on the other hand, face a substantial maintenance challenge. Dilution refrigerators are not maintenance-free and taking one offline for repairs raises extreme challenges for maintaining system functionality and the crucial quantum links among systems.

Low energy consumption

Carbon nanotube based qubits lead to vastly more affordable system running costs. Higher operating temperature, more efficient control electronics, 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.

Helium-3 efficient

C12's single-cryostat approach requires only a few tens of liters of Helium-3, unlike approaches that distribute processing over many dilution refrigerators and may require a significant fraction of the world's Helium-3 reserves. 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" (https://doi.org/10.48550/arXiv.2512.15001).

One objective

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

C12's approach starts from the constraints of large-scale systems and designs the architecture accordingly. By prioritizing scalability, inherent speed, and real-world viability from the outset, it defines an ambitious but realistic pathway to a genuinely useful quantum computer.

Conceptual rendering