

Opinions Libres

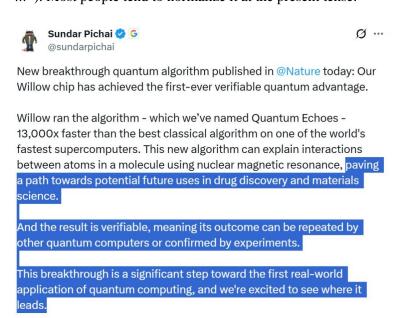
le blog d'Olivier Ezratty

Understanding Google Echoes

In October 2025, Google announced the publication of a paper in Nature on a "verifiable quantum advantage" with "Quantum Echoes". I like the name. It reminded me of one of my **preferred Pink Floyd song**, from the Meddle album, released in 1971!

Google's announcement was widely reported in social networks and the media as a breakthrough event. This was completed by the publication of an arXiv preprint with a use case to determine the structure of two simple molecules. But here, we are not at all in a quantum advantage regime.

As usual with Google, their first visible quantum advantage is their ability to obtain huge visibility, whatever they say, and whenever people understand what it is all about or not. Folk's understanding is inversely proportional to their propension to dispatch the news on social networks, particularly on LinkedIn. The scientific complexity behind these announcements and paper is, as usual, staggering. It obfuscates our grasp of the matter, whatever the pedagogy. It also amplifies the famous Dunning-Kruger syndrome. And we are exposed to such quantum computing breakthroughs every other month and year. See this **X** thread rewinding history back over 20 years on these multiple breakthroughs! These declarations manipulate in subtle ways the future tense ("path toward potential future...") and the present tense ("result is verifiable, can be repeated, ..."). Most people tend to normalize it at the present tense.



Like the title of my book, I try here to understand what was achieved. But whatever the effort, understanding the whole story remains a huge challenge. So we are like detectives, finding clues here and there about the achievement and looking at their consequences, without necessarily understanding the whole thing. I still had some discussions with various quantum physicists, superconducting hardware vendors and classical simulation specialists, and our LLM chatbot friends.

Here is the content of this paper:

- Algorithm: what problem was solved by the so-called "Quantum Echoes" algorithm in the Nature paper?
- Quantum advantage: what advantage was demonstrated in the Nature paper?
- Verifiability: what is this about? Is it different from past quantum advantage experiments?
- Used hardware: what chip was used, with which gates, and how many qubits?
- Lieb Robinson & Light Cone: why is this new? What are the consequences?
- Classical simulability: is Google Echoes / OTOC that difficult to simulate classically?
- OTOC NMR spectroscopy arXiv paper: was it delivering some quantum advantage in chemistry? Was it related to the Nature paper? (double nope!).

TL;DR version: Google Echoes experiment is a sort of extension of their past supremacy experiment, with a circuit containing random gates but producing some useful data. The arXiv paper on NMR spectroscopy looks potentially useful but is using a different algorithm than Google Echoes. The two Google experiments reach the limits of NISQ quantum computing with their its Willow processor. It could bring some quantum advantage on real-life use cases only with future FTQC systems.

Sources

The papers and posts we are dealing with are the following:

- June 2025 arXiv paper: Constructive interference at the edge of quantum ergodic dynamics by Dmitry A. Abanin, Rajeev Acharya et al, arXiv, June 2025 (77 pages).
- October 2025 Nature paper: Observation of constructive interference at the edge of quantum ergodicity by Google Quantum AI and Collaborators, Nature, Nature, October 2025 (8 pages and 68 supplemental material pages).
- Google blog post: A verifiable quantum advantage by Xiao Mi and Kostyantyn Kechedzhi, Google AI,
 Octobre 2025.
- October arXiv paper: Quantum computation of molecular geometry via many-body nuclear spin echoes by C. Zhang, R. G. Cortiñas et al, arXiv, October 2025 (76 pages).
- Nature comments: Google claims 'quantum advantage' again but researchers are sceptical by Elizabeth Gibney, Nature, Nature, October 2025.

Algorithm

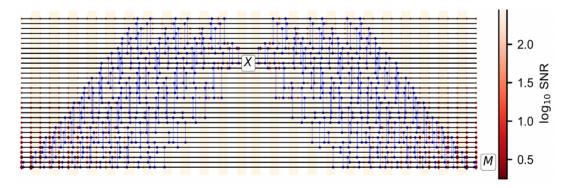
The "Quantum Echoes" algorithm presented in the Nature paper measures Out-of-Time-Order Correlators (OTOC). It is based on the creation of quantum interference between two evolutionary paths, one with a local perturbation, the other one without it. They extract signatures of quantum chaos and information propagation.

By exploiting the structure of the light cone induced by the Lieb-Robinson limit, the experiment identifies a region of constructive interference where measurements are verifiable and noise-resistant. That experiment was generic and didn't correspond to a real-life molecular system.

Google used a random circuit to generate a compute instance that seems difficult to simulate conventionally. The quantum circuit is fixed and is created only once but with using random gates, followed by echo generation operations to detect interference. It is therefore not a variational algorithm like what is usually used in the NISQ regime where you iterate multiple executions of a circuit containing an "ansatz" (a circuit with single qubit gate implementing rotations around the X, and Z axis in the Bloch sphere, and CNOT gates), and then a form of gradient descent using some classical optimizer to tune the values of each of the rotations.

Its randomness exhibits some similarities with the random circuits of the XEB benchmarks of their supremacy over Sycamore and Willow. But it generates some output data, not just random numbers like with the XEB. The algorithm simulates quantum dynamics, a use case known to be a bit more accessible than ground state estimation in quantum chemistry.

Below is what the version of the circuit made for 31 qubits looks like.



The total computing paper time in the Nature paper is 2.1 hours. This includes a large number of executions of the quantum circuit which has a depth of 29 quantum gate cycles, as well as some classical post-processing. The quantum circuit was probably executed millions of times, but without modifications. This number is not indicated in the paper. A potential 95-qubit version would require 180 million circuit shots as indicated in the paper.

Is this algorithm useful? There is some doubts around this claim, particularly given that it was not used in the arXiv paper published simulateously with the Nature paper as we will see later.

Quantum advantage

The "quantum advantage" presented in the Nature paper relates to running some OTOC circuit and its pre-post classical processing in 2.1 hours, vs 3.2 years on the DoE Aurora supercomputer, with each classically simulated circuit lasting 5 minutes. The difference in computing time is 13,000 in favor of the quantum computing solution.

This is impressive, but less than the 10^25 years in classical equivalent of Willow's quantum supremacy in 2024, which compared with less than 3 minutes of quantum computing time, running their cross-entropy benchmark (XEB) randomized circuit. The difference is that we are dealing here with a "real" problem instead of running a completely random circuit. Also, the number of qubits in the experiment was smaller, with 65 qubits in the OTOC experiment instead of 105 in the 2024 Willow supremacy experiment.

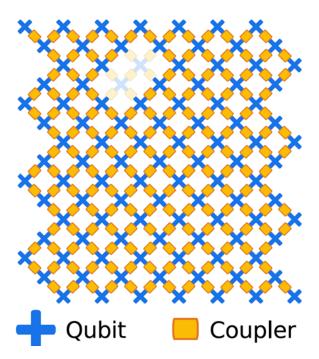
Verifiability

Google insists that their supremacy is "verifiable". Why is that so?

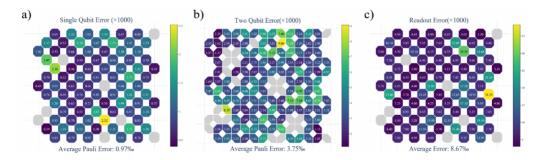
- Because the circuit used generates "expectation values" and not random data. But it is the general case of NISQ quantum algorithms.
- They simulated (emulated) their algorithm conventionally up to 40 qubits. Fine.
- The results remain the same when run on different quantum computers. But they haven't verified it, for example with an IBM Heron QPU. Heron's heavy-hex physical layout may be a limitation.
- They indicate that their results can be verified with real MNR spectroscopy, which is related to the arXiv paper. But the arXiv didn't use Echoes! And it was not tested with the 65-qubit version, we'll explain why later.

Hardware

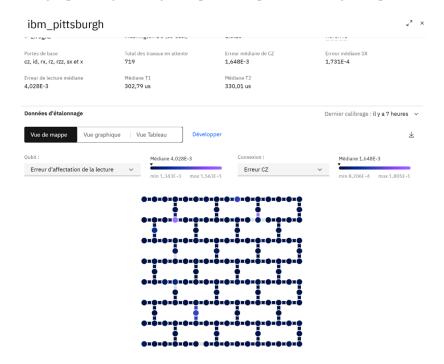
The OTOC experiment uses a version of Willow with 103 qubits. Why 103 and not 105? Because two qubits were dysfunctional (*below*). Like all players in this field, Google regularly iterates the production of new chips to improve their characteristics. But this chip seems of a lesser quality than the one tested in the 2024 Nature paper.



Exactly like for Sycamore in 2019 where they used 53 qubits of a processor that was supposed to have 54, one of which didn't work. This story of dysfunctional qubits is troubling and very common. With China's superconducting qubit QPUs, the story is the same, and even worse: Zuchongzhi 2.1 with 72 qubits only has 66 operating qubits while Zuchongzhi 3.0 with its 105 qubits has only 83 operational ones (*below*, **source**).



With IBM, it's the same story. In their very recent 156-qubit "Pittsburgh" Heron processor, 5 two-qubit gates are inoperative (*below*, **source**). This raises questions about what this will look like on a larger scale, with Google planning to integrate up to 10K qubits on a single chip, and John Martinis, up to 20K at Qolab (**source**).



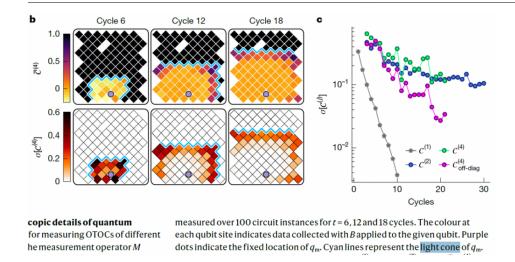
Google Quantum Echoes algorithm uses only 65 qubits of these 103 qubits. The reason is related to the quantum usable volume given the fidelity of the chip quantum gates. They cannot go beyond that. The circuit included 23 to 29 cycles, and about 1,000 two-qubit gates, and they exploited error mitigation technique, a method adapted to noisy qubits computing (NISQ). It uses zero-noise extrapolation which is machine-learning based and resource intensive.

They used Willow's "chip 2" variant which includes two-qubit iSWAP-like gates. The other chip, "chip 1" includes CZ gates. In the 2024 Willow paper, chip 1 was used for the creation of the logical qubit and chip 2 for the cross-entropy benchmarking of quantum supremacy. iSWAP gates have better fidelity than CZs, but CZ gates are useful for error correction. We can therefore consider that this experiment reaches the limits of the feasibility of NISQ algorithms on this chip. To go further, they will either have to drastically improve their gate fidelities, or switch to FTQC.

Lieb Robinson & Light Cone

One interesting item is Google's paper referring to the notion of "light cone".

This is related to the Lieb Robinson limit of the propagation of entanglement, correlations and information in "many-body" quantum systems. Google's experiment verifies this by observing this "light cone" of propagation. This is illustrated in Figure 2 of Nature (*below*). The speed of this propagation is not indicated in relation to the gate speeds.



This notion of light cone has important consequences in the scalability of FTQC QPUs of large physical size. In a chip like Willow, it's not problematic. But with several chips distant of over a meter, and connected by flexible microwave cables, the correlations propagation speed and time may largely exceed the two-qubit gate times. Therefore, it might create some constraints on circuit execution orchestration. I created a ten-page subsection on this matter in **Understanding Quantum Technologies 2025** to describe its implications on FTQC system architecture.

Classical simulability

Google tested several classical methods to simulate their circuits, using a mix of tensor networks, neural quantum states and Monte Carlo methods. It looked like state of the art classical equivalents. These classical simulations struggled beyond 40 qubits, due to the volume-law entanglement in the circuits. The **referees comments** of the Nature paper are rather positive.

Still, referee 4 wrote in his evaluation: "In this context a very promising alternative method, which for dynamics in 2D has already demonstrated to be often superior compared to tensor networks, is variational Monte Carlo, or in more general terms neural quantum states. Because of the recent advances of these Monte-Carlo inspired algorithms, I would suggest to revise the manuscript to properly account for this class of methods, when claiming quantum advantage". He added: "In order to actually settle whether quantum interference is of key importance, it would be necessary to study this systematically as a function of system size. Overall, it is not clear to me of how fundamental importance quantum interference actually is here, which, however, might also be rooted in not properly appreciating the random injections due to an unclear presentation"... "In the outlook the manuscript speculates that the hardware development "will rapidly increase the margin of quantum advantage, outrunning any gains that may result from improving classical simulation algorithms". This speculative comment might be correct when taking as a reference tensor networks, but it might not be for alternative numerical methods, which don't suffer from entanglement growth or matrix contraction complexity. I would suggest to adapt this speculation accordingly".

Referee 4 considers that the comparison with classical methods is incomplete: "The authors still seem to not properly present the existing literature and knowledge gained in recent years on neural quantum states, which is important insofar that any claim about advantages over classical algorithms is directly affected by this". The authors pushed-back on this by indicating that they had not found an effective NQS method to simulate their system. It is quite possible that a paper will soon be published showing that all this is classically simulable.

The reason for the referee's doubt is that theoretical work confirms that the computational complexity class of Out-of-Time-Order Correlators (OTOCs) is not fully understood. While OTOCs are widely used to probe quantum chaos and scrambling, their formal classification within complexity theory, meaning whether

computing them is BQP-hard, #P-hard, or classically simulable under certain conditions remains an open question. See for example **Out-of-time ordered correlators, complexity, and entropy in bipartite systems** by Pablo D. Bergamasco, Gabriel G. Carlo and Alejandro M. F. Rivas, Physical Review Research, October 2019 (5 pages).

It has also been proven in the past that noisy NISQ circuits can be classically simulated, like with the following papers:

- Classically Sampling Noisy Quantum Circuits in Quasi-Polynomial Time under Approximate Markovianity by Yifan F. Zhang, Su-un Lee, Liang Jiang, and Sarang Gopalakrishnan, arXiv, October 2025 (32 pages) which however doesn't indicate classical simulation run-times.
- Achieving Energetic Superiority Through System-Level Quantum Circuit Simulation by Rong Fu, Jian-Wei Pan, et al, arXiv, June 2024 (11 pages).
- Leapfrogging Sycamore: Harnessing 1432 GPUs for 7 times Faster Quantum Random Circuit Sampling by Xian-He Zhao, Jian-Wei Pan, Ming-Cheng Chen et al, arXiv, June 2024 (8 pages).
- Phase transition in Random Circuit Sampling by A. Morvan et al, Google AI, April-December 2023 (39 pages) where Google itself could ironically run its 2019 supremacy experiment in 6 seconds on a classical systems using tensor networks.
- What Limits the Simulation of Quantum Computers? by Yiqing Zhou, Miles Stoudenmire and Xavier Waintal, Physical Review X, November 2020 (15 pages) which was first to simulate Google supremacy experiment with using MPS to compress noise.

Similarly, bosonic boson sampling experiments can also be classically simulated, and ironically, by Google researchers:

- Efficient approximation of experimental Gaussian boson sampling by Benjamin Villalonga, Murphy Yuezhen Niu, Li Li, Hartmut Neven, John C. Platt, Vadim N. Smelyanskiy, and Sergio Boixo, arXiv, February 2022 (17 pages).
- The boundary for quantum advantage in Gaussian boson sampling, by Jacob F.F. Bulmer, 2022 (8 pages).

We can suspect that similar papers will be published in relation to the Google Echoes experiment.

OTOC NMR spectroscopy

In its blog post, Google stated that "This work demonstrates a verifiable quantum advantage and paves the way for solving real-world problems like Hamiltonian learning in Nuclear Magnetic Resonance (NMR)". It was associated to the arXiv published the same day, with 341 authors!

This second paper corresponds to running some MNR spectrography on toluene and dimethylbiphenyl, two rather simple organic molecules which have well-characterized spin networks. Using spin echo dynamics, the goal was to reconstruct their molecular geometry by simulating dipolar couplings via spin echo dynamics. It is a smaller problem instance than the famous protein folding problem. They could compare it with actual physical spectroscopy of these molecules, describing very well the related experimental protocole, which was

implemented at UC Berkeley.

It uses AlphaEvolve, "an evolutionary coding agent that leverages large language models (LLMs) to discover novel and efficient solutions for a variety of scientific problems. In this work, AlphaEvolve is used to generate a set of Python functions that produce quantum circuits approximating Hamiltonian evolution across an entire parameter landscape of time and Hamiltonian inputs, as opposed to optimizing individual quantum circuits or circuit parameters directly. The process was seeded with a human-written program (a first-order Trotter formula) that served as the initial solution to a given task". In other words, AlphaEvolve optimizes a variational circuit.

Google carried out this test with 9 and 15 qubits respectively, so in a regime that can be emulated on a simple laptop. The reason is that for this use case, the number of gate cycles is such that the depth of the circuit multiplied by the number of qubits reaches the limit of a thousand gates supported by their NISQ+QEM mode. Given the number of physical qubits used here, there is no quantum advantage at all. A laptop will always be cheaper and faster than any quantum computer for the execution of such circuit. The used qubits were arranged in 1D lines. They also used a different set of quantum gates than with the Nature paper, like fSim gates which are two-qubit quantum gates used to simulate fermionic interactions. They combined partial SWAP and controlled phase operations with arbitrary angles. They are native to many superconducting qubit platforms like Google's Willow. In this paper, Google used the "chip 1" Willow version, with native CZ gates and support for fSim gates. They didn't provide any indications of the number of circuit shots nor of computing time, or scaling laws for AlphaEvolve when used with a larger number of orbitals. This looks like a technical shell game!

I wondered whether it could be possible to use 65 qubits to run some MNR spectroscopy for larger molecules. The circuit depth required to run the equivalent OTOC experiment for a larger molecule using 65 qubits can be estimated by extrapolating the required complexity for simulating all-to-all coupled dipolar spin Hamiltonians using a Trotterized SWAP network. This task would demand 3.7×10^5 gates, requiring a qubit error rate of 2.7×10^6 . We are not far from the MegaQop.

This could only be achieved with an FTQC quantum computer. Google's 2024 Willow paper which deals with error correction estimated that obtaining an error rate of 10^-6 would require 1,457 physical qubits per logical qubit. So, we are up to 94K physical qubits. Google plans to integrate up to 10K physical qubits in 10×10 cm chips, which is already a huge manufacturing challenge. It is very hard to maintain low qubit variability (Qolab's team published a blueprint dealing with this), and to avoid defective qubits. Google had 2 defective qubits in the Willow chip that they used in the 2025 Nature experiment. Here, they would also need to interconnect at least 10 of these chips with flexible cables, adding losses, gate infidelities, and some circuit partitioning overhead. And the potential effects of the Lieb-Robinson velocity of entanglement/correlation/information propagation. There are significant challenges. Not there yet, for sure!

Here's a comparison between these various experiments, Sycamore supremacy in 2019, Willow supremacy in 2024, the Willow Nature experiment and the arXiv NMR geometry experiment, both in 2025:

Experiment	Year	Qubits Used	Task	Claimed Speedup	Verifiability	Practical Use
Sycamore (XEB)	2019	53	Random circuit sampling	~10^9	×	×
Willow (XEB)	2024	105	Random circuit sampling	~10^25	×	×
Willow (OTOC)	2025	65	OTOC measurement	13,000×		 Limited
NMR Geometry	2025	9-15	Molecular structure	None		×

Google simulated spin Hamiltonians and extracted long-range correlations, which are useful for determining bond angles and distances, mapping spin networks in small molecules and learning effective Hamiltonians from experimental data. This makes it a promising tool for quantum-enhanced molecular geometry inference, like in small organic molecules for finding drug candidates and functional materials. With FTQC hardware and hybrid algorithms, they could help with quantum-assisted NMR interpretation for protein fragments, spin-based mapping of local environments in folded structures, and Hamiltonian learning for biomolecular dynamics. But all of these scenarios require FTQC QPUs.

With regards to the experiment utility, Daniel Volz from Kipu Quantum is less optimistic on these use cases on LinkedIn: « Quantum advantage in chemistry achieved? Billions in business value unlocked? Unfortunately, not yet. Chemists already have many reliable tools for determining molecular structure. The molecules studied, Toluene, a common solvent found in paint thinner, and 3',5'-dimethylbiphenyl, an industrially irrelevant niche compound are far from challenging targets. The industrial questions explored by Google's team are, frankly, scientifically trivial and can be answered easily without a quantum computer. As a former research chemist, I struggle to see how determining distances between non-neighbouring atoms or dihedral angles is industrially relevant. Let's be honest, 99.9% of chemists wouldn't use a quantum computer for this".

As a result, all Google's statements about their "breakthough" have to be taken with a grain of salt. What they achieved was to create some quantum advantage experiment at the limit of what could be done with their NISQ QPU. It didn't pave the way for some immediate practical advantage. The experiment shows that quantum advantage is creeping toward utility, but we're still in the "proof-of-concept" phase. The real bottleneck is not algorithms, but hardware scalability and error correction. Practical advantage will only come when FTQC systems are available with at least 100 logical qubits and the support of one million operations. This is planned for the end of the decade by many hardware vendors like IBM, Alice&Bob, IonQ, Pasqal, QuEra and also by Google.

Google stated that the OTOC-based Quantum Echoes experiment "paves the way for solving real-world problems like Hamiltonian learning in Nuclear Magnetic Resonance (NMR)". This gives the impression that the Google Echoes OTOC algorithm could be used for practical NMR spectroscopy. While the Nature paper demonstrated some pending quantum advantage using Echoes OTOC interference circuits running on 65 qubits, it did not perform any chemistry or NMR task. Meanwhile, the arXiv paper, released the same day, did perform NMR spectroscopy on small molecules but did not use any Echoes OTOC circuit. Instead, it used AlphaEvolve-generated circuits based on Trotterized Hamiltonian evolution, and iteratively. Without any quantum advantage, since being tested on 9 and 15 qubits, and with relatively bad scaling prospects, at least in NISQ mode. It may generate some confusion on the actual work done by Google, despite the excellent presentation and scientific work achieved by these teams.



Google's Quantum Computer Makes a Big Technical Leap

Designed to accelerate advances in medicine and other fields, the tech giant's quantum algorithm runs 13,000 times as fast as software written for a traditional supercomputer.



It misled Elizabeth Gibney in her Nature piece when she wrote: "Google researchers and their collaborators fleshed out how they could apply the algorithm to simple molecules in a preprint study that they have submitted to arXiv". Same story in the New York Times with Cade Metz writing likewise: "Google's new algorithm is a step in that direction. In another paper published on Wednesday on the research site arXiv, the company showed that its algorithm could help improve what is called nuclear magnetic resonance, or N.M.R., which is a technique used to understand the structure of tiny molecules and how they interact with one another". It was not that algorithm!

Google's NMR spectroscopy experiments did not employ the full Quantum Echoes/OTOC algorithm from their Nature paper because current their superconducting hardware lacks the fidelity and scale needed to apply such interference-based protocols to complex molecular systems. They focused on small molecules like toluene, where simplified quantum simulations could be cross-validated with classical NMR data, avoiding the need for deep circuits and error-sensitive reversals, even though this method can't scale as well for larger molecules. So, both the OTOC framework which showcases a powerful route for Hamiltonian learning and verifiable quantum advantage and the AlphaEvolve-based MNR spectroscopy require FTQC QPUs.

There is also not much connection between Google's Echoes / OTOC and many other chemical simulations challenges like evaluating the ground state of many-body systems with high chemical accuracy, excited states, molecular dynamics, molecular docking and the like. These will require even larger FTQC QPUs, supporting thousands of logical qubits and billions of gates.

As Jens Eisert and John Preskill explained in their recent Mind the gaps: The fraught road to quantum advantage, arXiv, October 2025 (16 pages), there are lots of challenges to deliver on these promises. And their paper was just skimming the surface of these challenges, focusing more on algorithms challenges than physics and hardware engineering challenge to build these FTQC QPUs (summary *below!*).



Google's other work

All this story points to the need to create FTQC QPUs. Hopefully, Google is very active and has continuously provided multiple advances: its first memory logical qubit below the threshold with Willow in 2024, many quantum error correction innovations like with the yokes surface codes introduced in 2023, magic state cultivation also introduced in 2024 by Craig Gidney, quantum arithmetics improvements, and new resource estimates downgrading the RSA-2028 factorization need from 20 to 1 million physical qubits, also by Craig Gidney. Ryan Babbush et al are regularly proposing new chemical simulation algorithms with reduced resource requirements, mostly for FTQC QPUs. Now, they have to align these advances with hardware progress, which seems a bit slower.

As David Gilmour and Richard Wright sang in Echoes, "And no one showed us to the land, and no one knows the where's or why's, but something stirs and something tries, and starts to climb toward the light". That's our long-lasting quantum computing quest in a nutshell.

See also Quantum Echoes and the Future of Molecular Simulation: A Realistic Look at Google's Breakthrough by Phalgun Lolur, from Capgemini.

Cet article a été publié le 27 octobre 2025 et édité en PDF le 27 octobre 2025. (cc) Olivier Ezratty – "Opinions Libres" – https://www.oezratty.net