Agent teams and Lean proofs for quantum algorithms: what DeepMind and Google Quantum AI report

25 November 2026 · 8 min read · blogs.iictp.org/gemini-agent-teams-lean-verified-quantum-circuits-explained

Short answer: on 23 November 2026 Google DeepMind and Google Quantum AI described teams of Gemini 4 agents that proposed improvements to known quantum algorithms, with every result checked by the Lean theorem prover before it was accepted. DeepMind reports an 11.8x cost reduction for a fusion-energy circuit, a 290x speedup for an arcsin routine and a 2.57x speedup for a Grover implementation, and has open-sourced the Lean library, problem specifications and verified solutions. The technical point is the method: search by AI agents, filtered by machine-checked proof.

TL;DR The short version

- Google DeepMind and Google Quantum AI used teams of Gemini 4 agents to find improvements to known quantum algorithms, accepting only results proven in Lean.
- They report large cost reductions, such as 11.8x for a fusion simulation, and have open-sourced the Lean library and verified solutions.
- Lean checks mathematical correctness, not whether hardware can run the circuit; the lesson is to generate widely and verify strictly.

Read online, listen to it, or change the text size: https://blogs.iictp.org/gemini-agent-teams-lean-verified-quantum-circuits-explained

What did the researchers do?

According to the DeepMind post by Sergey Shirobokov, Ben Jaderberg and Ryan Babbush, separate agent teams were assigned to each problem and asked to find algorithmic improvements. Quantum circuits cannot be tested at full size on classical computers or on today's small quantum processors, so correctness had to be shown mathematically. Each idea was therefore proved in Lean before entering a shared notebook.

How it works

- **Independent teams per problem.** Each team searches for better circuits on its own.
- **Formal verification gate.** Lean, a formal theorem prover, checks each claimed result, so an idea that sounds plausible but is wrong does not get through.
- **Shared notebook and a recurring meeting.** Every four hours, agents from each team met at a virtual conference to trade, debate and compare ideas, and the computational cost gradually fell until the experiment

ended, the post says.

- **Open Lean library.** DeepMind built a new open-source Lean library for quantum algorithms so teams of agents could build on verified building blocks.

What is new compared with earlier approaches

Earlier quantum-algorithm improvements were mostly found by specialists by hand. DeepMind's claim is that proof-checked AI search compresses work that took years of manual design, while the proof requirement keeps the output trustworthy. It links this to its other work, such as AlphaEvolve and AlphaQubit.

Results reported by the source

- **Fusion (inertial confinement):** computing cost, a metric combining qubits and operations, cut by 11.8x, with required memory down from 5,650 to 4,338 error-corrected qubits.
- **Light-sensitive molecule:** quantum simulation 8x cheaper.
- **Cytochrome P450:** a 2.5x improvement, with lower memory.
- **Building blocks:** a 2.57x speedup for a specific Grover implementation and a 290x speedup for arcsin, a routine used to load numbers into quantum states.
- All figures are DeepMind's and Google Quantum AI's own, from the blog post.

Limitations and open questions

- The post states that running these algorithms in real scientific simulations would still need more resources than the first error-corrected quantum computers are likely to provide.
- Lean checks mathematical correctness of a circuit, not whether hardware can run it or how noise affects it.
- The post summarises a paper; read the paper for the exact problem definitions and baselines. Sources were read on 25 November 2026. Products, rules and standards change, so check the linked originals before you rely on any detail.

How to apply the pattern

1. Choose a problem with a precise formal specification, so a proof checker can say yes or no.
2. Let agents propose candidate improvements in parallel, but accept only those that pass the checker.
3. Keep a shared log of accepted ideas so later attempts build on verified work.
4. Measure the improvement against a fixed baseline cost metric, as the post does with qubits and operations.

What it means for engineers and researchers

The transferable lesson is "generate widely, verify strictly". The same shape appears in software verification and in other formal domains. If you work with agents, the checker is as important as the model. See AlphaProof Code: formal verification of software and Antigravity /teamwork.

Frequently asked questions

- **What did DeepMind and Google Quantum AI announce?** Agent teams built on Gemini 4 that improved known quantum algorithms, with each result proved in Lean.
- **What is Lean?** A formal theorem prover that checks mathematical proofs by machine.
- **Which results are reported?** An 11.8x cost cut for a fusion circuit, 8x for a photochemistry simulation, 2.5x for Cytochrome P450, 2.57x for a Grover implementation and 290x for arcsin.
- **Is the code available?** DeepMind says it is open-sourcing the Lean library, problem specifications and verified solutions on GitHub.

Related reading

- AlphaProof Code: formal verification of software
- Antigravity /teamwork: multi-agent verification
- AI Practitioner certification

> Related credential: AI Practitioner. Explore the exam facts and syllabus.

Sources and further reading

1. Shirobokov, Jaderberg, Babbush, Google DeepMind and Google Quantum AI. Finding faster quantum algorithms with agents and formal verification. 23 November 2026. deepmind.google/blog/finding-faster-quantum-algorithms-with-agents-and-formal-verification

Published by IICTP, an independent technical credential body, on 25 November 2026. The current version of this post is at https://blogs.iictp.org/gemini-agent-teams-lean-verified-quantum-circuits-explained. More guides: blogs.iictp.org.