



ARE QUANTUM COMPUTERS REALLY ENERGY EFFICIENT?

Proponents often tout quantum computing as a more energy efficient alternative to classical computing methods. However, the extent to which it can reduce energy usage remains unclear, as experts have not yet agreed on metrics to determine its energy consumption. **By Sophia Chen**

We live in an era of unprecedented computational capability. In 2016, an artificial intelligence program, AlphaGo¹, beat a professional human player at the game Go. In 2020, the artificial intelligence (AI) software AlphaFold demonstrated that it could predict the three-dimensional structure of proteins from their amino acid sequence², a 50-year old grand

challenge in biology. In 2021, researchers at the University of California, San Diego, used a supercomputer in tandem with AI to model 1 billion atoms in the SARS-CoV-2 virus³. In 2022, OpenAI released ChatGPT, a large language model with an ability to mimic a human conversationalist. The list goes on and on.

However, this technology comes at a high cost. To reach the current level of accuracy and efficiency, these computational

approaches require staggering amounts of energy, resulting in substantial carbon emissions. For example, the world's fastest supercomputer, Frontier, draws 8 megawatts when it idles – a quantity that could simultaneously power thousands of homes – and training a large language model one time generates about the same carbon emissions as a passenger on a flight from New York City to San Francisco⁴.

Given current trends, the energy use of computation looks to continue increasing. Between 2012 and 2018, the amount of computing power required for cutting-edge large AI models doubled every 3.4 months, amounting to a more than **300,000-fold increase**. Another estimate⁵ found that to cut the error rate in half, an AI model would require more than 500 times the computational cost. Computing technology uses 5 percent of all energy consumption in the US⁶, and higher computational costs will drive up energy use. That increased computing will then contribute to carbon dioxide emissions worldwide, as **82 percent** of the world's primary energy consumption comes from fossil fuels. The **Inter-governmental Panel on Climate Change** has urged the world to halve carbon emissions by 2030 to limit global warming to 1.5 degrees Celsius. At this rate, our computing ambitions could threaten the sustainable future of the planet.

Consequently, experts are looking to new strategies that can rein in energy use while continuing to improve computing performance. One proposed solution: quantum computing. Proponents are already using prototype quantum computers to make the case that they can solve new types of problems with less energy. For instance, a 2020 demonstration showed that a quantum computer could perform a math puzzle using 50,000 times less energy than the world's most powerful supercomputer at the time⁷. "Quantum has real potential for an energy advantage," says Jerry Chow, who leads quantum computing hardware development at IBM.

A different form of logic

A quantum computer is a device that manipulates information using the mathematics of quantum mechanics, as opposed to binary logic. For example, a quantum computer doesn't represent information as 1s and 0s. Instead, its basic unit of information, known as a qubit, corresponds to the probability of being either 1 or 0. The qubit's state is like a coin flipping in the air: before landing, the coin's state is neither heads nor tails, but some probability of either. In quantum lingo, the coin is in a superposition of heads and tails. Similarly, a qubit represents a superposition of 1 and 0.

Erwin Schrödinger famously illustrated the concept of a superposition in a thought experiment involving a cat in a box with a vial of poison and a radioactive substance. When the radioactive substance decays, it releases a particle that triggers the release of the poison,

killing the cat. According to quantum mechanics, before anyone opens the box, the cat is technically in a superposition of being dead and alive at the same time.

By exploiting superposition and other quantum properties, such as entanglement, a quantum computer is capable of fundamentally different mathematical operations than classical computing. The potential applications have attracted billions of dollars of both public and private investment. Banks, for example, are keen to investigate quantum computers for optimization problems⁸, whereas pharmaceutical companies are studying **their capability** to simulate complex molecules and accelerate drug discovery. **Hyundai has partnered with US startup IonQ, and Mitsubishi has partnered with IBM** to use quantum computers to simulate battery chemistry, which could lead to electric cars with higher mileage per recharge.

Iordanis Kerenidis, a physicist at the startup **QC Ware**, develops algorithms for quantum computers. In a recent preprint, he and his colleagues used two of IBM's quantum computers to demonstrate how they could make machine learning more energy efficient⁹. In the study, a quantum computer was used to train a popular machine learning architecture, the transformer. Using an algorithm that requires both classical and quantum computers, they trained a machine learning model to accurately classify medical images. For instance, the model learned to diagnose the severity of a patient's diabetes based on images of their retina. The model also learned to diagnose whether a patient had pneumonia based on a chest X-ray. Its accuracy was competitive with several state-of-the-art classical machine learning algorithms, such as ResNet.

The model's potential energy savings come from the unique way that quantum computers encode information. During training, when mapping image patterns to vectors, a classical machine learning algorithm creates the vector space inefficiently, because it tends to extract similar pixel patterns, corresponding to vectors near each other in space. In contrast, quantum computers inherently extract dissimilar pixel patterns, corresponding to vectors that are orthogonal. "Every structure [that the quantum computer extracts] gives completely new information," says Kerenidis. This means that during training, the quantum computer needs to extract fewer pixel patterns than a classical computer to build out the same vector space.

Consequently, the quantum computer could train the transformer more efficiently.

Machine learning practitioners describe model complexity in terms of their number of parameters. Kerenidis and colleagues' neural network consisted of a few hundred parameters, compared to ResNet, which has millions of parameters. "We hope that quantum computers can use much fewer parameters, resulting in much smaller models with similar performance," says Kerenidis. "Smaller models are easier to train and cost less, both in terms of energy and time."

Although the results are promising, Kerenidis's team has not performed a formal comparison between their algorithm's energy use and that of a classical machine learning algorithm. In fact, they haven't tallied its energy use at all, because researchers haven't agreed on a metric to describe the resource requirements of a quantum computer.

A gauge for quantum computing energy use

The lack of agreement on an energy use metric is, in part, because the field is so new — researchers have only managed to build the first quantum computers of note arguably within the last decade. So far, many studies have focused on proving that quantum computers offer a speedup over classical computers, rather than understanding their energy use.

However, existing quantum computers have still not conclusively shown they are faster than classical computers at any specific task. Physicist John M. Martinis, working at Google at the time, and colleagues made this claim in 2019¹⁰, but experts dispute the significance of their achievement. In that demonstration, Google's team showed their quantum computer could randomly sample numbers from a specific statistical distribution in 200 seconds, a task that they claimed would take a state-of-the-art supercomputer 10,000 years. Since then, researchers have developed better classical algorithms for that task, with one group claiming that a supercomputer should be able to complete it in a few dozen seconds to beat Google's quantum computer¹¹.

Nevertheless, these experiments established that researchers have achieved a threshold level of control over their quantum computers to complete a relatively complex task. In their 2019 study, Google researchers also hinted that they thought the quantum computer offered energy savings over a supercomputer¹⁰. "[W]e estimate that performing the same task [...] would cost 50 trillion core-hours and consume one petawatt hour of energy. To put this in perspective,

[...] the net quantum processor time is only about 30 seconds.” Their machine runs on 26 kilowatts, three orders of magnitude less than a state-of-the-art supercomputer.

Still, some researchers say that the field hasn’t done enough homework to make claims about energy savings. “You will find lots of hand waving,” says Alexia Auffèves of the CNRS, “mostly based on this idea that a quantum computer, if it works properly, involves fewer physical operations than a classical computer, and thus will cost less energy.” Auffèves is working to come up with a more rigorous framework for evaluating quantum computing energy use. Some researchers assume that longer computations correspond to proportionally more energy use. However, Auffèves and her colleagues found that the relationship between energy consumption and computation time in a quantum computer is more complicated than the conventional wisdom. “There is no proportionality between time and energy,” she says.

Instead, quantum computers use disproportionately more energy as algorithms grow to involve more consecutive operations. This has to do with the challenge of preserving the information in a quantum computer. “The longer the algorithm is, the longer you need to preserve quantum information, and so the more you have to cool down your qubit,” says Marco Fellous-Asiani of the University of Warsaw, formerly Auffèves’s graduate student. “The power [draw] grows over the duration of the computation.”

Auffèves wants the community to come together to define a metric for energy efficiency for quantum computers, such as a ratio between a metric describing performance and the resource cost. For classical computing, the performance metric is typically measured by floating point operations per second (FLOPs), while the resource is in watts. Thus, the energy use of supercomputers is typically expressed in ‘gigaflops per watt’. However, the quantum computing community has not decided on how to define neither the numerator nor the denominator for this metric. They have some options for the numerator: IBM uses a performance metric known as ‘quantum volume’¹², which describes the complexity of the algorithm that the machine can run. The startup IonQ has defined a metric known as the ‘algorithmic qubit’, which rates a quantum computer’s ability to execute a set of benchmark algorithms. These values all depend on the quantum computer’s basic specifications, such as the number of qubits or the machine’s error rate. “There is no

agreement now on what should be the best way to evaluate the performance of a quantum computer,” says Auffèves.

“There is no agreement now on what should be the best way to evaluate the performance of a quantum computer.”

To get the conversation started, Auffèves, Fellous-Asiani, and their collaborators have presented a framework for optimizing the energy use of quantum computers¹³. In their proposed framework, they model the relationship between the quantum computer’s noise and its energy consumption, says Fellous-Asiani. The type of quantum computer they studied needs to be kept at a temperature close to absolute zero, where colder temperatures guarantee less noise, and thus, fewer errors. “You need to expend energy to guarantee success,” he says. Then, they determine the level of accuracy that the quantum computer needs to deliver for the task at hand and calculate the amount of acceptable noise to reach that accuracy threshold. This framework then allows them to optimize energy consumption based on the application.

Building a new community

In August 2022, Auffèves, and colleagues, co-founded a group called the [Quantum Energy Initiative](#), which declares the importance of researching the environmental footprint of quantum computing¹⁴. Their [manifesto](#) outlines the big questions in the field: “While [...] observations seem to point toward an energetic advantage of a quantum nature, the physical mechanisms behind them are barely understood: how does energy consumption scale with the processor size? How does it relate to the computational performance or the qubit technology? How does it compare to classical processors?”

The required expertise to answer these questions spans communities typically siloed from each other, says Auffèves. It requires quantum hardware experts who work directly with the qubits, but also engineers who design the cryogenics systems that cool the hardware, as well as the control systems that program the qubits, and researchers who study thermodynamics in quantum mechanical systems. “The big challenge is connecting different communities that are not talking to each other much right now,” she says.

Auffèves and her colleagues circulated the manifesto among the quantum computing community, and they have since gathered more than 300 people from 48 countries to form a community to study these topics. They plan to organize a [workshop](#) around quantum computing and energy use in Singapore in November 2023.

Although current quantum computing research mainly focuses on gaining computational speed, Auffèves says it is essential to center energy use in the design of quantum computers to move toward a more sustainable future. “You have to make a choice at some point if you want to optimize time, or if you want to optimize energy,” says Auffèves. Auffèves is partly motivated to study energy use in response to anti-technology sentiments she has observed in recent years. “We are at a critical point in our societies, where you have people who believe in innovation and science, and people who think that we should stop everything because science and innovation are responsible for climate change,” she says. She’s experienced anti-technology protesters at a conference firsthand. “The atmosphere is not healthy,” she says.

A probabilistic future

To develop new technology, researchers will need to think more explicitly about the environmental costs. “If you deploy your technology, you have to acknowledge the fact that we live in a finite world,” says Auffèves. As Auffèves and her colleagues work to quantify quantum computing energy use, the technology continues to evolve. For one thing, their analysis focuses on a specific type of quantum computing hardware whose qubits are made from superconducting circuits. However, researchers are contending with several other qubit types as well, including ones made of trapped ions, neutral atoms, and photons. These qubit types have fundamentally different engineering requirements; neutral atom quantum computers, for example, require fewer cryogenics than superconducting quantum computers. “It is a very complicated thing to research the energy savings of quantum, because it really depends on the hardware,” says Kerenidis. Auffèves is working to develop cross-disciplinary collaborations that could help researchers develop a consistent framework to account for energy on different types of hardware.

It is also unclear how much quantum computers will cut down on the energy consumption of supercomputers, as quantum computers will not replace classical computing. Instead, experts think that it’s more likely that a quantum

computing chip will reside within a supercomputer, which accesses it to solve specific problems. Their anticipated role in future computing is in flux, which means it is hard to estimate their overall energy savings, even if they do prove to be energy efficient.

These early studies show promise, but it is too soon to think of quantum computing as any sort of environmental silver bullet. Researchers still have to prove that quantum computers can actually do something useful. “One big challenge in the next three to five years is to figure out how to use the quantum hardware to improve a task in practice,” says Kerenidis.

Sophia Chen ✉

Freelance science journalist, Columbus, OH, USA.

✉ e-mail: schen311@gmail.com

Published online: 26 June 2023

References

1. Silver, D. et al. *Nature* **529**, 484–489 (2016).
2. Senior, A. W. et al. *Nature* **577**, 706–710 (2020).
3. Dommer, A. et al. *Int. J. High Perform. Comput. Appl.* **37**, 28–44 (2022).
4. Strubell, E., Ganesh, A. & McCallum, A. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* 3645–3650 (ACL, 2019).
5. Thompson, N. C., Greenewald, K., Lee, K. & Manso, G. F. Deep learning’s diminishing returns. *IEEE Spectrum* (24 September 2021); <https://spectrum.ieee.org/deep-learning-computational-cost>
6. Wolpert, D. Why do computers use so much energy? *Scientific American* (4 October 2018); <https://blogs.scientificamerican.com/observations/why-do-computers-use-so-much-energy/>
7. Villalonga, B. et al. *Quantum Sci. Technol.* **5**, 034003 (2020).
8. Chen, S. Banks are betting that quantum computing can find them an investment edge. *Protocol* <https://www.protocol.com/manuals/quantum-computing/finance-banks-investing-investment-edge> (4 May 2020).
9. Cherratt, E. A. et al. Preprint at <https://arxiv.org/abs/2209.08167> (2022).
10. Arute, F. et al. *Nature* **574**, 505–510 (2019).
11. Pan, F., Chen, K. & Zhang, P. *Phys. Rev. Lett.* **129**, 090502 (2022).
12. Cross, A. W. et al. *Phys. Rev. A* **100**, 032328 (2019).
13. Fellous-Asiani, M. et al. Preprint at <https://arxiv.org/abs/2209.05469> (2022).
14. Auffèves, A. *PRX Quantum* **3**, 020101 (2022).