Science | DOI:10.1145/3623641 Don Monroe

Quantum Speedup for the Fast Fourier Transform?

It is difficult to improve on the widely used, already-efficient algorithm.

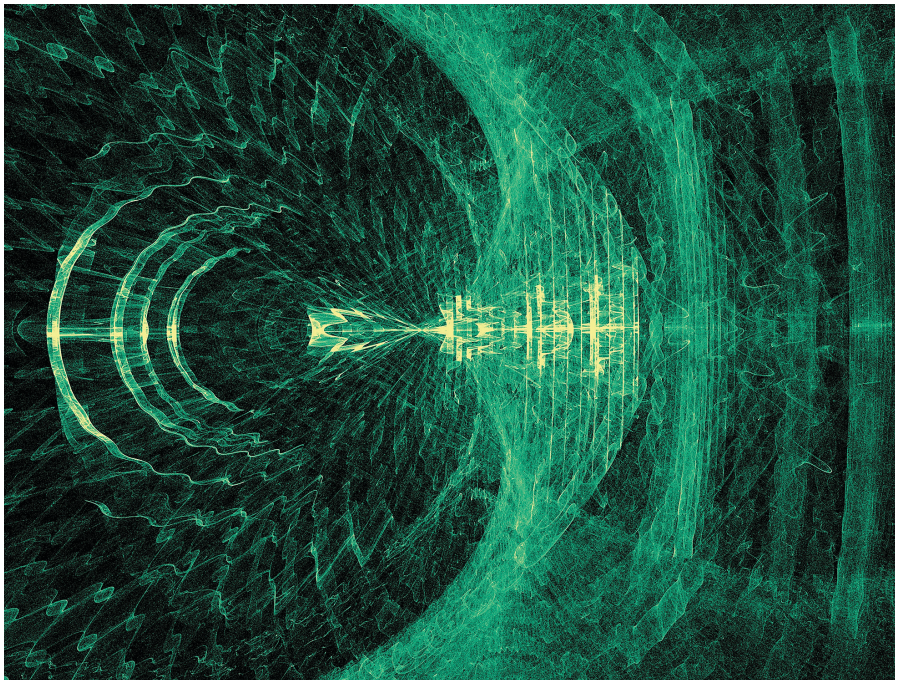

IMAGE BY MAXIM STUDIO

RESEARCHERS HAVE BEEN exploring the theoretical potential of quantum computers for decades. Only in recent years, however, have real systems begun to achieve an accuracy and scale that makes their eventual impact seem plausible.

A major impetus for the attention and investment is the revolutionary introduction of methods by Peter Shor, now at the Massachusetts Institute of Technology, that could allow quantum computers to break current public-key encryption schemes. The resources, such as time, required by classical techniques for factoring numbers, grow exponentially with the number of digits in the input. The demands of Shor's quantum algorithms grow only as a polynomial function, potentially rendering encrypted data vulnerable to future quantum computers.

A key ingredient of this method is the quantum implementation of a Fourier transform to rapidly determine the dominant repetition frequency present in a data sequence, which points to the prime factors of a large number. Indeed, the quantum Fourier transform, or QFT, is central to many of the exponentially powerful quantum algorithms that have been proposed to date. (Another class of algorithms developed by Lov Grover speeds unstructured quantum searches, but only as a power of the problem size.)

Fourier transforms already are widely used for tasks such as image processing, but QFT does not directly address these important applications. Researchers aiming to correct this situation have therefore developed a very different quantum approach, which could be part of a toolkit for more widespread quantum computing. So

far, however, the quantum advantage seems less compelling.

FFT: Fast Fourier Transform

Fourier transforms convert input data, such as a sequence of values over time, into the amplitudes (which are complex numbers) of a frequency spectrum. In most digital applications, the input comprises some number N of discrete samples, and the output spectrum contains N spatial or temporal frequencies. The spectrum is completely equivalent to the original data, which can be reconstituted using a similar process.

The transform also provides an easy way to compress the data, for example by representing the amplitudes with lower resolution or by discarding high-frequency components. JPEG (Joint Photographic Experts Group) formats, for example, let users shrink a file in size while preserving the most important features. A variant of the Fourier transform, the "discrete cosine transform," computes spatial frequency components from the two-dimensional pixel data.

Mathematically, the amplitude of each frequency is computed by summing all the inputs after a "rotation," multiplying the value of each input point by a factor that depends on the product of the frequency and the point's index. N points and N frequencies therefore seem to require N^2 products. However, as popularized by James Cooley and John Tukey in 1965, the computation can be vastly sped up through the use of a divide-and-conquer strategy called the fast Fourier transform, or FFT. For example, if N is a power of two, the full transform can be computed by repeatedly breaking the matrix into four submatrices. The resources required for this algorithm grow only slightly faster than linearly in N (as $N\log N$), rather than in proportion to N^2.

QFT: Quantum Fourier Transform

The full power of quantum computing was demonstrated with an even more powerful speedup of a Fourier transform, but its output is more limited than that of the FFT. Shor's factoring algorithm, for example, determines the dominant period (or frequency) of a sequence, but not the amplitudes of other frequencies.

The Fourier transform is a natural way to interrogate the relative information in the amplitudes, Abbot said, but "you have to, at the end, find a way to extract the information you want."

To determine the period, the states of quantum bits, or qubits, are transformed just as they are in a Fourier transform. After this manipulation, the qubits still contain all the input information. Typically, though, the manipulations involve "tiny, tiny rotations that could get very quickly swamped out by noise," warned Alastair Abbott of the Inria (French Institute for Research in Computer Science and Automation) center at France's Université Grenoble Alpes. "Finding a way to do these robustly is difficult."

Critically, extracting this information only produces one answer. When its value is measured, each qubit must take a value of 0 or 1, just like a classical bit. More subtle "quantum information," for example representing multiple possible outcomes, disappears at measurement.

The Fourier transform is a natural way to interrogate the relative information in the amplitudes, Abbot said, but "You have to, at the end, find a way to extract the information you want, which is the nontrivial part." The ingeniousness of the QFT is creating a state whose measurement almost certainly produces a 1 for the qubit representing the target periodicity, while others are measured as 0.

QFFT: Quantum Fast Fourier Transform

Despite its unquestionable importance, the QFT does not provide access to the full spectral information that is

ACM Member News

DIGITAL TECHNOLOGY AND SOCIAL CHANGE

Kentaro Toyama is W.K. Kellogg Professor of Community Information at the University of Michigan School of Information in Ann Arbor, MI.

Toyama earned his undergraduate degree in physics from Harvard University in Cambridge, MA, and his master's and Ph.D. degrees in computer science from Yale University in New Haven, CT.

On receiving his Ph.D. in 1997, Toyama went to work for Microsoft Research in Redmond, WA. In 2004, he relocated to India to co-found Microsoft Research India in Bangalore.

In 2010, he joined the faculty of the University of California, Berkeley. He moved to the University of Michigan in 2015, where he has remained.

Toyama's research focuses on digital technology and social change. He is the author of the book *Geek Heresy: Rescuing Social Change from the Cult of Technology*, which is largely based on his experiences applying digital technologies for social economic development while working in India. He explained, "My main conclusion is that technology does not add a net-positive value; instead, it amplifies whatever human force is already there—whether it's positive or negative."

Toyama adds that he believed technology could bring positive social change, but now he is much more skeptical that is the case.

"As a computer scientist, I don't have the right training for how to change the world, but given the prominence of computing in the world today, it is incumbent on those of us who are computer scientists to rein in the worst sides of technology, such as the negative effects of social media."

—John Delaney

widely used from the FFT algorithm. To address this issue, physicists at Japan's Tokyo University of Science (TUS) propose a quantum circuit to perform the FFT, which they call QFFT (although this name risks confusion with QFT).

The proposed circuit is not expected to show a quantum speedup over the classical FFT, at least for a single input (such as an image). However, the researchers argue their circuit could perform the calculation for many images in parallel without requiring additional resources.

To implement the Fourier transform, the TUS team uses a strategy called basis encoding, which they say is an "essentially different" strategy from the effectively analog premeasurement manipulations in QFT. Rather than directly representing each input datum as the amplitude of a qubit, the input data is into a binary representation. The binary representation is then transferred to the state of many qubits.

The quantum circuit is designed to perform calculations on these bits that resemble those of classical circuits. One difference is that additional input streams can be simultaneously embedded in the configuration as a quantum "superposition" and the FFT calculation done on all of them at once.

In the end, measuring the transformed output would only yield one result. However, if measurement is deferred, the scheme could be incorporated as a kind of subroutine into another quantum calculation that might take advantage of the full Fourier transform. "Arithmetic operations, such as adding and subtracting, can be applied to the output of QFFT, [but] not QFT," said TUS graduate student Ryo Asaka. "It's one of the advantages of QFFT." Most current applications have no need for comparisons between multiple images, however, which is presumably how retaining multiple quantum representations of the Fourier transform would be most powerful.

Importantly, the resources required should include those needed to prepare the input for the calculation, which is sometimes overlooked. "If the time for preparation of input increases in proportion to the number of [inputs], the advantage of QFFT disappears," Asaka noted. To avoid this outcome, he envisions preparation of a complex quantum configuration of the qubits and storing them in a quantum random-access memory. Such a "qRAM" was proposed in the 2000s, but has not so far been made practical.

"Some problems might be a bit easier to do, (while) other problems, without a fault-tolerant computation, you have really not much hope of doing."

Looking for the Quantum Advantage

Other researchers also are trying to imagine ways to use quantum computers for mainstream tasks. Despite massive investments, however, the qubits in existing quantum computer systems are too noisy and lose their delicate quantum state too quickly, and there are too few of them, although the systems are getting steadily bigger and better. In 2019, a team from Google announced their system performed a task that would take orders of magnitude longer for a classical system, and researchers in China later described similar demonstrations. The chosen tasks, however, do not demonstrate a practically important capability.

In the long run, most observers expect most interesting practical problems will depend on quantum error correction, a subtle technique that may require many more qubits. "Some problems might be a bit easier to do," Abbott said, but many "other problems, without a proper fault-tolerant computation, you have really not much hope of doing. Most quantum Fourier transform algorithms fall in that latter camp." He added, "You couldn't even do it badly at the moment."

Researchers already are exploring quantum algorithms, though, using systems available in the current "Noisy Intermediate-Scale Quantum" (NISQ) era. Some of the most promising candidate problems include quantum simulations of molecules and materials. Small quantum systems may also be recapitulated with a modest number of qubits, such as a recent claim to emulate the transport of information through a wormhole in spacetime using only seven qubits.

Another class of problem that has gotten high attention from even low-tech businesses is the ubiquitous challenge of optimizing a complex system, such as the allocation of manufacturing resources or finances. Finding the best solution might be possible despite significant noise, for example through hybrid classical-quantum algorithms. (D-Wave Systems also markets machines that target these problems with many more qubits, but many researchers question whether its operation is truly quantum mechanical, unlike more common gate-based architectures.)

Companies and governments worldwide continue to explore possible algorithms, architectures, and physical implementation of quantum. However, robust systems of thousands of qubits for robust, large-scale fault-tolerant computing are still probably years in the future, if ever. C

Further Reading

Asaka, R., Sakai, K., and Yahagi, R.
"Quantum circuit for the fast Fourier transform," ***Quantum Information Processing 19*****, 277 (2020). DOI: https://doi.org/10.1007/s11128-020-02776-5**

Arute, F. et al.
"Quantum supremacy using a programmable superconducting processor," ***Nature 574*****, 505 (2019) DOI: https://doi.org/10.1038/s41586-019-1666-5**

Hoefler, T., Thomas Häner, T, and Matthias Troyer, M.
"Disentangling hype from practicality: On realistically achieving quantum advantage," ***Communications 66*****, 82 (2023). DOI: https://doi.org/10.1145/3571725**

Don Monroe is a science and technology writer based in Middlebury, VT, USA.