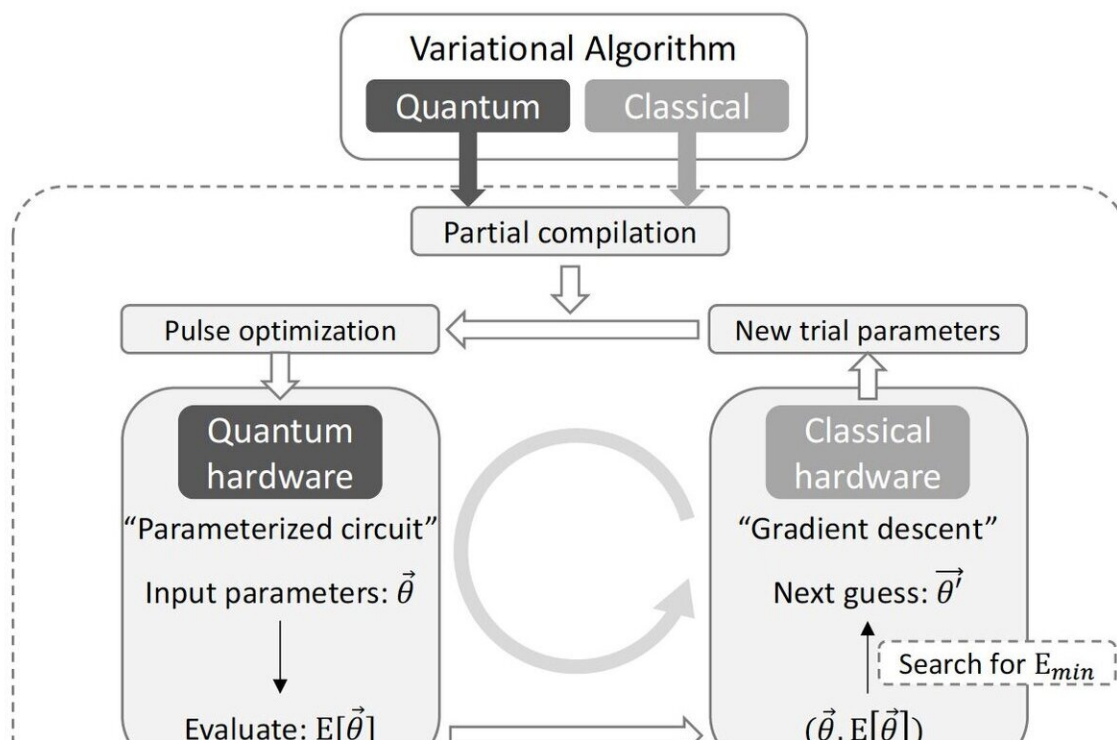


New compiler makes quantum computers two times faster

October 11 2019, by Epiqc



A flow chart describing the compiling of variational algorithms to speed up quantum computations. Credit: EPiQC/University of Chicago

A new paper from researchers at the University of Chicago introduces a technique for compiling highly optimized quantum instructions that can be executed on near-term hardware. This technique is particularly well suited to a new class of variational quantum algorithms, which are

promising candidates for demonstrating useful quantum speedups. The new work was enabled by uniting ideas across the stack, spanning quantum algorithms, machine learning, compilers, and device physics. The interdisciplinary research was carried out by members of the EPIQC (Enabling Practical-scale Quantum Computation) collaboration, an NSF Expedition in Computing.

Adapting to a New Paradigm for Quantum Algorithms

The original vision for [quantum algorithms](#) dates to the early 1980s, when physicist Richard Feynman proposed performing molecular simulations using just thousands of noise-less qubits (quantum bits), a practically impossible task for traditional computers. Other algorithms developed in the 1990s and 2000s demonstrated that thousands of noise-less qubits would also offer dramatic speedups for problems such as database search, integer factoring, and matrix algebra. However, despite recent advances in quantum hardware, these algorithms are still decades away from scalable realizations, because current hardware features noisy qubits.

To match the constraints of current and near-term quantum computers, a new paradigm for variational quantum algorithms has recently emerged. These algorithms tackle similar computational challenges as the originally envisioned quantum algorithms, but build resilience to noise by leaving certain internal program parameters unspecified. Instead, these internal parameters are learned by variation over repeated trials, guided by an optimizer. With a robust optimizer, a variational [algorithm](#) can tolerate moderate levels of noise.

While the noise resilience of variational algorithms is appealing, it poses a challenge for compilation, the process of translating a mathematical

algorithm into the physical instructions ultimately executed by hardware.

"The trade-off between variational and traditional quantum algorithms is that while variational approaches are cheap in the number of gates, they are expensive in the number of repetitions needed," said [Fred Chong](#), the Seymour Goodman Professor of Computer Science at UChicago and lead PI for EPIQC. "Whereas traditional quantum algorithms are fully specified at execution time and thereby fully optimizable pre-execution, variational programs are only partially specified at execution time."

Partial Compilation

The researchers address the issue of partially specified programs with a parallel technique called partial compilation. Pranav Gokhale, a UChicago PhD student explains, "Although we can't fully compile a variational algorithm before execution, we can at least pre-compile the parts that are specified." For typical variational algorithms, this simple heuristic alone is sufficient, delivering 2x speedups in quantum runtime relative to standard gate-based compilation techniques. Since qubits decay exponentially with time, this runtime speedup also leads to reductions in error rates.

For more complicated algorithms, the researchers apply a second layer of optimizations that numerically characterize variations due to the unspecified parameters, through a process called hyperparameter optimization. "Spending a few minutes on hyperparameter tuning and partial compilation leads to hours of savings in execution time", summarizes Gokhale. Professor Chong notes that this theme of realizing cost savings by shifting resources—whether between traditional and quantum computing or between compilation and execution—echoes in several other EPIQC projects.

The researchers next aim to demonstrate their work experimentally. Such experimental validation has become possible only recently, with the release of cloud-accessible quantum computers that can be controlled at the level of analog pulses. This level of control is much closer to hardware than standard gate-based control, and the researchers expect to realize greater efficiency gains from this pulse interface.

The researchers' paper, "Partial Compilation of Variational Algorithms for Noisy Intermediate-Scale Quantum Machines" ([arXiv link](#)) will be presented at the MICRO computer architecture conference in Columbus, Ohio on October 14. Gokhale and Chong's co-authors include Yongshan Ding, Thomas Propson, Christopher Winkler, Nelson Leung, Yunong Shi, David I. Schuster, and Henry Hoffmann, all also from the University of Chicago.

More information: Partial Compilation of Variational Algorithms for Noisy Intermediate-Scale Quantum Machines, arXiv:1909.07522 [quant-ph] arxiv.org/abs/1909.07522 . DOI: [10.1145/3352460.3358313](https://doi.org/10.1145/3352460.3358313)

Provided by University of Chicago

Citation: New compiler makes quantum computers two times faster (2019, October 11) retrieved 10 April 2024 from <https://phys.org/news/2019-10-quantum-faster.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.