Research provides speed boost to quantum computers
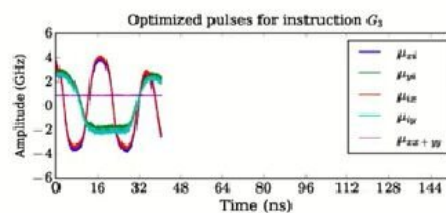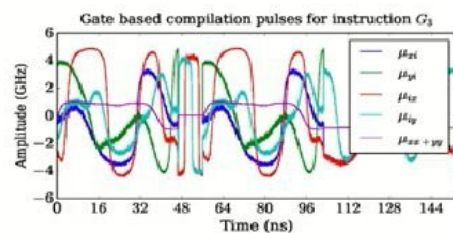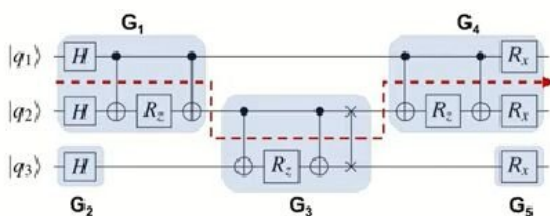
April 12 2019

A new finding by researchers at the University of Chicago promises to improve the speed and reliability of current and next generation quantum computers by as much as ten times. By combining principles from physics and computer science, the researchers developed a new scalable compiler that makes software aware of the underlying quantum

hardware, offering significant performance benefits as scientists race to build the first practical quantum computers.

The UChicago research group comprises computer scientists and physicists from the [EPiQC (Enabling Practical-scale Quantum Computation) collaboration](#), an NSF Expedition in Computing that [kicked off in 2018](#). EPiQC aims to bridge the gap from existing theoretical algorithms to practical [quantum](#) computing architectures on near-term devices.

Merging Approaches from Computer Science and Physics

The core technique behind the EPiQC team's paper adapts quantum optimal control, an approach developed by physicists long before [quantum computing](#) was possible. Quantum optimal control fine-tunes the control knobs of quantum systems in order to continuously drive particles to desired quantum states—or in a computing context, implement a desired program.

If successfully adapted, quantum optimal control would allow quantum computers to execute programs at the highest possible efficiency...but that comes with a performance tradeoff.

"Physicists have actually been using quantum optimal control to manipulate small systems for many years, but the issue is that their approach doesn't scale," said researcher Yunong Shi.

Even with cutting-edge hardware, it takes several hours to run quantum optimal control targeted to a machine with just 10 quantum bits (qubits). Moreover, this running time scales exponentially, which makes quantum optimal control untenable for the 20-100 qubit machines expected in the

coming year.

Meanwhile, [computer scientists](#) have developed their own methods for compiling quantum programs down to the control knobs of quantum hardware. The [computer science](#) approach has the advantage of scalability—compilers can easily compile programs for machines with thousands of qubits. However, these compilers are largely unaware of the underlying quantum hardware. Often, there is a severe mismatch between the quantum operations that the software deals with versus the ones that the hardware executes. As a result, the compiled programs are inefficient.

The EPiQC team's work merges the computer science and physics approaches by intelligently splitting large quantum programs into subprograms. Each subprogram is small enough that it can be handled by the physics approach of quantum optimal control, without running into performance issues. This approach realizes both the program-level scalability of traditional compilers from the computer science world and the subprogram-level efficiency gains of quantum optimal control.

The intelligent generation of subprograms is driven by an algorithm for exploiting commutativity—a phenomenon in which quantum operations can be rearranged in any order. Across a wide range of quantum algorithms, relevant both in the near-term and long-term, the EPiQC team's compiler achieves two to ten times execution speedups over the baseline. But due to the fragility of qubits, the speedups in quantum program execution translate to exponentially higher success rates for the ultimate computation. As Shi emphasizes, "on quantum computers, speeding up your execution time is do-or-die."

Breaking Abstraction Barriers

This new compiler technique is a significant departure from previous

work. "Past compilers for quantum programs have been modeled after compilers for modern conventional computers," said Fred Chong, Seymour Goodman Professor of Computer Science at UChicago and lead PI for EPiQC. But unlike conventional computers, quantum computers are notoriously fragile and noisy, so techniques optimized for conventional computers don't port well to quantum computers. "Our new compiler is unlike the previous set of classically-inspired compilers because it breaks the abstraction barrier between quantum algorithms and quantum hardware, which leads to greater efficiency at the cost of having a more complex compiler."

While the team's research revolves around making the compiler software aware of the underlying hardware, it is agnostic to the specific type of underlying hardware. This is important since there are several different types of quantum computers currently under development, such as ones with superconducting qubits and trapped ion qubits.

The team expects to see experimental realizations of their approach within the coming months, particularly now that an open industry standard, OpenPulse, has been defined. This standard will enable operation of quantum computers at the lowest possible level, as needed for quantum optimal control techniques. IBM's quantum roadmap highlights OpenPulse support as a key objective for 2019, and other companies are expected to announce similar plans as well.

The team's full paper, "Optimized Compilation of Aggregated Instructions for Realistic Quantum Computers" is now published on arXiv and will be presented at the ASPLOS computer architecture conference in Rhode Island on April 17. In addition to Shi and Chong, co-authors include Nelson Leung, Pranav Gokhale, Zane Rossi, David I. Schuster, and Henry Hoffman, all at the University of Chicago.

More information: Yunong Shi et al. Optimized Compilation of

Aggregated Instructions for Realistic Quantum Computers, *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems - ASPLOS '19* (2019). DOI: 10.1145/3297858.3304018

Provided by University of Chicago