Quantum information theory (20110401)
Lecturer: Jens Eisert
Chapter 12: Non-universal quantum computers

Contents

Chapter 12

Non-universal quantum computers

12.1 Quantum simulators

So far, we have discussed universal quantum computers. But then, quantum computers as we have them are not quite universal. They are smallish (the largest qubit device available embodies 433 qubits), they have limited control, and are still rather noisy. One speaks of *noisy intermediate scale quantum* (NISQ) devices. What can such devices do after all? The answer to this is not quite known, and is for good reason a topic of intense study and research. But we have a look at some aspects thereof. The first application is that of a *quantum simulation*. Simulation is a core task, and a substantial fraction of modern supercomputers is dedicated to simulating quantum materials. It makes a lot of sense, therefore, to ask the question whether quantum systems can simulate other quantum systems with smaller resource requirements. This turns out to be true, at least in principle. The idea of a quantum simulator is indeed a compelling one.

Quantum simulators come in several flavours. One distinguishes *static quantum simulators* probing static such as ground state properties from *dynamical quantum simulators*. The latter probe dynamical properties that can be probed in a time evolution of the type

$$\langle O(t)\rangle = \mathrm{tr}(e^{-itH}\rho(0)e^{itH}O), \tag{12.1}$$

where

$$H = \sum_{j=1}^{N} h_j \tag{12.2}$$

is typically a *local Hamiltonian* acting on a lattice involving n sites of local dimension d, equipped with a Hilbert space $\mathcal{H} = (\mathbb{C}^d)^{\otimes N}$. Local here means that each of the terms $\{h_j\}$ acts on a small number of sites only, usually nearest neighbours. The observable O is commonly local as well, such as a local Pauli matrix, say, $O = Z$ acting on a single site. We have seen that local Hamiltonian evolution is in principle BQP complete, so

one can think of quantum computation in this fashion. But this is not how one usually thinks about such quantum simulators. Instead, one is interested in learning properties of strongly correlated quantum systems, beyond classical capabilities.

12.1.1 Analog quantum simulators

In *analog quantum simulation*, one recreates, mimicks, the exact Hamiltonian of the original system, but under precisely controlled conditions. This is much less of a ridiculous idea than it may first appear. In fact, in systems of cold atoms in optical lattices or with trapped ions one can recreate interacting systems very precisely, and also probe them in the laboratory. One can also probe properties under enormous precision that would be out of reach for the real quantum material. To elaborate the specific architectures that allow for analog quantum simulation is interesting in its own right and could be the topic of a course on its own.

The fact that time evolution is BQP complete here comes in as an advantage: One cannot devise universal classical simulation algorithms of quantum dynamics. Therefore, classical simulation algorithms such as quantum Monte Carlo Methods, density functional theory or tensor network methods soon reach their limitations. But analog quantum simulators do not face such restrictions: In this way, one can argue that analog quantum simulators allow to probe questions, e.g., in quantum statistical physics, already as of today beyond the means of classical supercomputers, in the sense that with the best known classical algorithms, one cannot reliably keep track of quantum dynamics. This is an enormously interesting field of research.

12.1.2 Digital quantum simulators

Digital quantum simulators require basically quantum computers. In the dynamical reading, one slices time evolution into stroboscopic time slices and approximates the continuous time evolution by a number of small time evolution steps approximated by a gate-based quantum computer. The formula at the heart of the matter is here the *Trotter formula*.

Trotter formula: For any two Hermitian operators A and B, one has

$$e^{A+B} = \lim_{n\to\infty} \left(e^{A/n} e^{B/n}\right)^n . \tag{12.3}$$

This can be made a functioning quantum algorithm. Here, time is chopped into n small time steps, and one keeps track of time evolution in a stroboscopic fashion. For two Hamiltonian terms H_1 and H_2, with $\|H_1\| \leq K$ and $\|H_2\| \leq K$ (this upper bound of the operator norm captures the coupling strength and can be set to unity to give time a unit), one has

$$e^{-iH_1} e^{-iH_2)} = e^{-i(H_1+H_2)} + \mathcal{O}(K^2). \tag{12.4}$$

Pushing this idea further, for local terms $h_1, \ldots, h_n$ of a local Hamiltonian

$$H = \sum_{j=1}^{N} h_j \tag{12.5}$$

one finds

$$e^{-ih_1} \ldots e^{-ih_N} = e^{-i(h_1+h_2+\cdots+h_N} + \mathcal{O}(N^3K^2), \tag{12.6}$$

if for all j, $\|h_j\| \leq K$. For a constant $C > 0$ such that the number of steps n is lower bounded as

$$n > CN^3(Kt)^2/\epsilon, \tag{12.7}$$

one has

$$\|(e^{-ih_1t/n} \ldots e^{-ih_Nt/n})^2 - e^{-iHt}\| \leq \epsilon. \tag{12.8}$$

This sense of approximation is the right one: If two unitaries are close in operator norm, the respective quantum states will be close in trace norm. Hence, in this fashion, one can approximate time evolution under local Hamiltonians to arbitrary precision. This gives rise to an algorithm to approximate continuous time evolution of local Hamiltonians with an effort of $\mathcal{O}(\mathrm{poly}(N)(\|H\|t)^2/\epsilon)$: The time evolution is slices into pieces that can be gate decomposed on a quantum computer. This is the most basic of all digital quantum simulation schemes. It has been further developed into higher order schemes that have a favourable error scaling, linear combination of unitaries approaches, an idea called qubitization and randomized schemes. It is a flourishing topic of research still.

12.2 Variational quantum computers

12.2.1 Variational quantum eigensolvers

Quantum computers exist, with system sizes up to $N = 433$ (as of 2023) qubits. To add insult to injury, these are still comparably noisy and not quantum error corrected. This is too little to perform the Shor algorithm for reasonable system sizes including quantum error correction. But is still good enough to tackle interesting problem. One such approach is to use quantum computers in *hybrid algorithms* in which a quantum computer is only a part of a larger classical algorithm. This classical algorithm takes data from measurements from the quantum circuits, alters the control parameters and accordingly prepares states and alters the variational quantum circuits along the way.

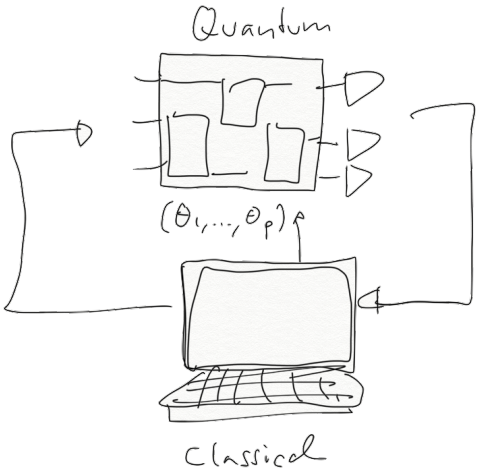

The quantum part at the heart of the algorithm is a short quantum circuit

$$\theta := (\theta_1, \ldots, \theta_p) \mapsto U(\theta_1, \ldots, \theta_p) \tag{12.9}$$

defined by real variational parameters $(\theta_1, \ldots, \theta_p)$. These can be seen as "knobs" in an experiment, as parameters of quantum gates that can be freely chosen by the experimenter. In variational quantum eigensolvers, one hence basically solves the following problem.

Variational quantum eigensolvers: They aim at approximately finding a solution to

$$E_{\min} = \min\langle\psi(\theta)|H|\psi(\theta)\rangle, \tag{12.10}$$

where

$$|\psi(\theta)\rangle := U(\theta_1, \ldots, \theta_p)|0, \ldots, 0\rangle. \tag{12.11}$$

This may look like a simpler an enterprise that it actually is in practice. After all, one has to find good variational sets so that one can expect a good approximation in the first place. This is not obvious, and one has to find a good compromise between expressivity and depth of the involved circuits. This applies, e.g., to problems in quantum chemistry where one has to find good representations to start with. Then, one has to find strategies of finding good updates

$$(\theta_1, \ldots, \theta_p) \mapsto (\theta'_1, \ldots, \theta'_p). \tag{12.12}$$

How to optimally do this is an interesting question in its own right.

- Indeed, one can *estimate gradients* by taking measurements and performing updates based on that. One can also acquire higher order information, such as elements of Hessians. Natively, this approach requires many measurements.

- Using ideas such as *parameter shift rules* help reducing the number of expectation values one needs to measure to find good updates. These are rules based on instances of symmetries that exploit that certain functional values can be evaluated by taking measurements at shifted values. To find good strategies of classical control is still a subject of intense research.

12.2.2 Quantum approximate optimization algorithm

Quantum approximate optimization goes a step further. It aims at finding good approximations to combinatorical optimization problems, such as `MaxCUT`. Such problems are often `NP`-hard in worst case complexity. Still, one hopes to find good approximations with quantum algorithms. The goal is to find an approximate solution to a cost function

$$f : \{0,1\}^n \to \mathbb{R}_0^+. \tag{12.13}$$

Such a cost function can be encoded in a Hamiltonian H_f , so that finding the maximum of f amounts to finding the ground state energy

$$E_{\min} := \min\langle\psi|H_f|\psi\rangle, \tag{12.14}$$

of H_f over all quantum states. One aims at finding a binary string that achieves an approximation ratio r for

$$\frac{f(z)}{f_{\max}} \geq r \tag{12.15}$$

for all $z \in \{0,1\}^n$, for the optimal solution being

$$f_{\max} := \max_{z\in\{0,1\}^n} f(z). \tag{12.16}$$

One does not expect to get an exact solution: The problem being `NP`-hard in worst case complexity means that an exact solution is also out of reach for a quantum computer. But one aims at realizing an approximate solution. This is achieved by initially preparing a product state described by a state vector $|+\rangle^{\otimes N}$ and by applying a sequence of steps

$$e^{-i\gamma_j H_f} \tag{12.17}$$

for $j = 1, \ldots, p$, implementing the encoding Hamiltonian, followed by products of local Pauli X rotations

$$(e^{-i\beta_j X})^{\otimes N}, \tag{12.18}$$

that "kick the state out of the Pauli-Z basis". At the end of the day, after p steps, this sequence is followed by an estimation of the expectation value $\langle H_f\rangle$ by virtue of repeated measurements. The variational parameters are again chosen using similar strategies as above. Similar issues of expressivity and classical control appear as well.

Quantum approximate optimization algorithm (QAOA): The entire state vector before the measurement in the quantum approximate optimization algorithm involving N qubits is given by

$$(e^{-i\beta_p X})^{\otimes N} e^{-i\gamma_p H_f} \dots (e^{-i\beta_1 X})^{\otimes N} e^{-i\gamma_1 H_f} |+\rangle^{\otimes N}. \tag{12.19}$$

The variational parameters $(\gamma_1, \dots, \gamma_p, \beta_1, \dots, \beta_p)$ are optimized in order to minimize $\langle H_f \rangle$, the expectation value of a fictitious Hamiltonian.

12.2.3 Further thoughts on quantum approximate optimization

It is important to note that the quantum approximate optimization algorithm can be seen as a particular instance of a variational quantum eigensolver, just that a specific choice has been made for the variational quantum circuit. At the same time, one can argue that the algorithm is inspired by discrete time adiabatic quantum computations. This algorithm is much studied, for good reason. At the same time, there is evidence that with too short circuits, no quantum speedup can be attained. For example, it has been shown that for $p = 1$, for a single step version, classical algorithms can achieve the same performance as the single step QAOA on the combinatorial optimization problem `MAX-3-LIN-2`. Since one cannot expect an exact solution, people are also deeply concerned with what improvements one can precisely expect over classical algorithms. One has to be aware of the fact that one can also find good classical efficient approximations, for example, by so-called *convex relaxations* of the original combinatorial problem (leading to semi-definite problems discussed above). It is an interesting and intense question for ongoing research in what precise sense one can hope for quantum advantages. Again, since the problems considered are usually `NP`-hard not only in worst case complexity, but also in approximation up to a constant ratio, one cannot expect a quantum computer to solve all instances in polynomial time.

12.3 Final thoughts

12.3.1 Random circuit sampling

The presumably simplest near-term algorithm is that of *random circuit sampling*. Here, one does not alter the circuit in any variational way on the fly. Instead, one considers the native output distribution of a random quantum circuit. A moment of thought reveals that at the end of the day, every experiment is quantum physics is a random sampling scheme: One just has to take the raw native data that follow Born's rule. Random quantum sampling is inspired by this kind of thinking. For N qubits, one starts off with $|0\rangle^{\otimes N}$ and implements a random circuit, by implementing a collection of random two-qubit and single qubit gates that are taken from a *universal gate set* as discussed above. This gives rise to a state vector $|\psi\rangle \in (\mathbb{C}^2)^{\otimes N}$. Then one performs a measurement in

the computational basis (the Pauli-Z-basis), to get outcomes $x = (x_1, \ldots, x_N)$ with probability

$$p(x) = |\langle\psi|x_1, \ldots, x_N\rangle|^2. \tag{12.20}$$

This distribution is pretty structureless, in fact, it is close to being uniform. Yet, the tails of this distribution are intricate. They are so intricate that one cannot sample from a distribution close to this one efficiently on a quantum computer.

Random circuit sampling: An efficient classical sampling from a distribution that is close up to a constant error in the $\|.\|_{l_1}$ distance for random circuit sampling with size and depth of the circuit appropriately chosen leads to a collapse of the polynomial hierarchy to the third level.

The mentioned $\|.\|_{l_1}$ distance is nothing but the sum of absolute values of a vector and is the same as the trace norm applied to diagonal quantum states. It also has basically the same statistical interpretation. The "collapse of the polynomial hierarchy to the third level" is an elaborate concept in theoretical computer science. It is sufficient to say that proving $\mathtt{P} = \mathtt{NP}$ would lead to a collapse of the polynomial hierarchy altogether, so the mentioned collapse is a similarly implausible (but strictly speaking unproven) collapse of complexity classes.

Such random sampling schemes have been experimentally implemented, to great attention, first by the Google AI Quantum team using a superconducting device involving 53 qubits. In the meantime, it has been argued that one can classically sample from the given distribution to a small error also on classical supercomputers. The question of how to fairly compare quantum and classical schemes is an exciting and important topic of ongoing research. Also, it is interesting to note that one cannot *black-box verify* such devices, and more indirect methods of verifying the output distributions such as *linear cross entropy benchmarking* must be resorted to. Also, in the meantime, larger scale experiments have been performed with 66 qubits (by a Chinese team) which is to date unchallenged by classical computers. These experiments show "quantum advantages" or what is sometimes called "quantum supremacy" – relating to the situation that quantum computers can outperform classical computers for the same well defined computational problem, underpinned by a language of computational complexity. But they are not practically minded algorithms in their own right, and applications are scarse.

12.3.2 Closing remarks

To precisely find out what near-term quantum computers can do is a highly interesting and exploratory field of research that receives a lot of attention. People are much concerned with the impact of noise that basically limits variational quantum circuits to logarithmic depth. Researchers are asking questions of good control in quantum variational algorithms. Importantly, it is much studied what applications in *combinatorial optimization* and *quantum-assisted machine learning* are conceivable. Indeed, there is evidence that quantum computers can help in both families of problems. How this is precisely done is presumably one of the most exciting questions of present-day physics and computer science.