Leonie Mück

How I will learn to love organic chemistry through quantum computers

When the experiment failed the sixth time, my patience started to fade. I was an undergraduate chemistry student and had been assigned a small research project performing an early step in a natural compound synthesis. It should have been a bog-standard substitution reaction. But something unexpected about the structure of this particular molecule made it impossible. We assumed steric effects were at play, but who could really tell? After weeks of work experimenting with temperature, reaction time, reactants, solvents, a tiny bit of the desired product appeared. How I wished I could simply zoom in with a giant magnifying glass to see what was happening during that reaction!

As I learned later, such a giant magnifying glass exists – computers. Electrons, the particles that determine chemical reactivity and behaviour, are quantum mechanical so the task for the computer is to approximately solve the Schroedinger equation. Armed with clever approximations, an understanding of the underlying physics and a powerful computer, we can simulate what goes on in our flask. And not only that. We can predict chemical phenomena like the geometry of a molecule, the transition states during a reaction, spectra or crystal structures.

In the last couple of decades, computational chemistry has made giant leaps fuelled by an explosion in computing power and new methods. Alas, progress is grinding to a halt. A phenomenon known as Moore's law stipulated that computing power increases by a factor of two or so every year – an exponential increase. But this curve is flattening. On top of that, we seem to have exhausted the arsenal of mathematical tricks that could improve our approximations. As it stands today, computers are an imperfect magnifying glass of chemistry: They can either give a crisp view of a very small part of a reaction – for example, how one solvent molecule behaves. Or they can zoom out to show the full picture, simulating all solvent molecules and all reactants – but everything is a bit blurry and out of focus.

Enter quantum computers. These new-fangled apparatuses are promising a step change in computational chemistry. Being quantum systems themselves, they can represent electrons in molecules with much greater ease than your regular bit-based machine. On classical computers, calculating the exact

solution to the Schroedinger equation scales exponentially with system size. On a quantum computer this scaling is linear. The challenge is to build a quantum computer that is powerful enough to handle the complexity of chemistry. Recent developments are heralding rapid progress: In 2019, the world witnessed how a quantum computer comprised of only 53 qubits, the quantum version of the classical bits, built by Google outperformed the world's fastest classical computer, Summit at Oak Ridge National Lab, at a specific task. [1]

While this task, random sampling from a circuit, wasn't particularly interesting for chemistry, it gave me glimmers of hope that, by 2050, we will have built sufficiently sizeable quantum computers, so that no-one will ever have to leave as much blood, sweat and tears at the organic chemistry bench as I did when I was an undergrad: We would use quantum computers as the ultimate magnifying glass to see everything that happens in our flasks with clarity.

The inner workings of quantum computers are fundamentally different from your regular laptop. In your laptop, the Central Processing Unit, CPU, is your workhorse executing all the programmes that you, the user, tells it to run. Nowadays, the building block of a CPU is the transistor, a semiconductor device that can be in an "on" or "off" state corresponding to the bit we are so familiar with. These are then used to construct NOT-AND (or "NAND") gates which form the basis of everything that happens in your CPU. This is all quite straightforward.

Quantum computers, on the other hand, are based on qubits. Qubits are constructed from two-level quantum systems corresponding to the bits of a regular computer. Using advanced experimental techniques, scientists can produce linear combinations of these states. The coefficients of this linear combinations describe the probability of finding the qubit in either a 0 or a 1 state upon measurement. Now imagine you have a system of many qubits, let's say n. Each qubit has a coefficient corresponding to the probability of finding it in the 0 state and a coefficient that describes the probability of finding it in a 1 state. This means there are 2^n possibilities for different combinations of zeros and ones that you need to consider! Fundamentally, this is what gives quantum computers their power. You can consider 2^n possibilities at once.

Hold on, you may say. What about this measurement that you just mentioned in passing? From your quantum mechanics course, you may still remember Schroedinger's poor cat that was dead and alive at the same time before you opened the box

Dr. Leonie Mück
Chief Product Officer
St Andrew's House
59 St Andrew's Street, Cambridge CB2 3BZ
leonie.mueck@riverlane.com

and "measured". That measurement determined the definitive fate of the feline. You are quite right that the same happens in quantum computers – we are in a state of sheer endless possibilities before we measure. But as soon as we open the box, figuratively speaking, the state collapses into just one of the many possibilities. This is the fundamental reason that harnessing the power of quantum computers is a hard nut. We must figure out how to combine information from the different combinations before we measure to extract some global information about them.

Very clever mathematicians and physicists have succeeded in designing algorithms that give quantum computers a decisive leg up. More precisely, there are around 400 algorithms that display a so-called "quantum speed-up".[2] The most famous is Shor's algorithm with which you can perform prime number factorisation exponentially faster than on a classical computer. Algorithms that solve quantum chemical equations are in the same class of exponential quantum speed-ups. Other algorithms, for example Grover's algorithm which inverts a function, yield a square-root quantum speed-up. [3]

Various different physical implementations of qubits are in the race to become the transistor of quantum computers. In atomic qubits, the internal states of atoms and ions are used as two-level systems. Superconducting qubits, on the other hand, are mesoscopic systems in which the two-level quantum system is formed by a Bose-Einstein condensate of Cooper pairs, which are pairs of electrons that are bound together in superconductors. Superconducting and atomic qubits are considered the most advanced at present, but spin systems in silicon and diamond as well as photons are quickly catching up.

No matter the specific implementation, qubits are fickle objects. The slightest contact with the classical world destroys their quantum natures, a process dubbed decoherence. Complicated control systems, consisting of lasers, waveform generators and fast programmable elements which can react quickly to changes, are necessary to keep the qubits alive and manipulate them. If we ever want to construct big and powerful quantum computers, one of the most important tasks of this control system is working against the decoherence, extending the lifetime of the qubits so that we can run long calculations on them. In a process called error correction, the same information gets encoded in the qubits multiple times, creating redundancies such that thousands of physical qubits represent one logical qubit. Through a clever layout, we can check whether one of the qubits has decohered – and if it has, we can perform the necessary operation to correct the error. All of this has to happen incredibly fast, in microseconds for atomic qubits and in nanoseconds for superconducting qubits.

Your heart must be sinking right now. If we need thousands of qubits to enable error correction and have 53 error-prone qubits today, using a quantum computer as a true magnifying glass for chemistry sounds like a pipe dream! Constructing a fully error-corrected quantum computer is certainly a herculean task. However, much like Moore's law in classical computation, we have seen a yearly doubling of quantum computing power, as measured by how many qubits we can build combined with
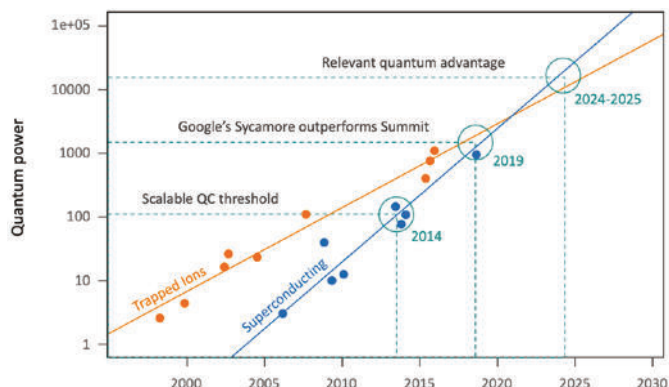
Fig. 1: The development of quantum power, a measure for how many gates can be run on a quantum computer, over time. If this trend continues, quantum computers will be sufficiently powerful in 2024-2025 to perform computations of academic or commercial relevance ("Relevant quantum advantage"), for example computing small molecules with a complicated electronic structure with high accuracy.

how noisy they are (see figure 1), in the past decade. And there is another piece of good news: We won't necessarily need a fully error corrected quantum computer to do useful things, we can probably use the "noisy intermediate scale quantum" or NISQ devices that will be constructed in the next couple of years.

Take chemistry as an example. On a quantum computer, there are two main algorithms with which we can compute chemical properties. Just like conventional quantum chemistry methods, both of them can be used to solve the Schroedinger equation to calculate the energy of a system. Quantum phase estimation was invented first. In this algorithm, we directly encode the physical description of the system, the Hamiltonian, into the quantum computer and perform a series of manipulations to extract information about the energy before we measure. Since the series of necessary manipulations is rather long, we will probably need an error-corrected quantum computer to run it.

But there is another algorithm on whose shoulders the hopes and dreams of the NISQ era lie: The Variational Quantum Eigensolver. In this algorithm we take an educated guess about how the solution to the Schroedinger equation would look, referred to as an "ansatz", and map it to a quantum computer. After performing only a few quick manipulations, we immediately measure the energy. To get an accurate estimate of the energy, we must prepare and measure multiple times. This solution gets passed onto a classical computer which performs an optimisation step and then feeds the results back into our "ansatz". The process repeats itself until we find no difference between the ansatz that we mapped to the quantum computer and the solution that we get after measurement. If you know a bit about how we use classical computers to approximate the Schroedinger equation, this iterative process of finding the lowest energy will look familiar to you. Indeed, the processes are very similar – in the Variational Quantum Eigensolver you just outsource the hardest bit of this algorithm to the quantum computer. [4, 5]

In the Variational Quantum Eigensolver, we have substituted the long series of manipulations from Quantum Phase Estimation for a shorter series plus measurement performed many times, and – voilà – we can use a device with a relatively short qubit lifetime to perform a chemistry calculation. In theory,

the Variational Quantum Eigensolver and Quantum Phase Estimation are both claimed to deliver an exponential speed-up compared to conventional quantum chemistry. However, predicting the exact time that the Variational Quantum Eigensolver will need to run in practice for a realistic system is a bit tricky. We need to consider many different elements: What are our errors in the system? How many times do we have to prepare and measure to get a certain accuracy? How far off the actual solution do we start and how will our classical optimiser perform? Various results on how to combine the merits of the two algorithms have been published so that we can make the series of manipulations as long as permitted by qubit lifetimes. This will mitigate some of the drawbacks of the Variational Quantum Eigensolver that may make quantum chemical calculations impractically long. [6]

Whichever algorithm we use on quantum computers, it is up against almost 100 years of method development in conventional quantum chemistry. In practice, hardly any of the methods that a computational chemist would use in their day-to-day life solves the Schroedinger equation exactly. Over the decades, approximations have been developed that bring the unfavourable scaling down so we can calculate chemical systems of interest. The biggest success story in quantum chemical method development is probably density functional theory, or DFT, which happens to hit the sweet spot between computational scaling and accuracy. In density functional theory we don't solve the Schroedinger equation directly. Instead, we express our equations in terms of the electron density and fit certain parameters using empirical parameters from a benchmarking set of chemicals.

With my failing organic reaction, I would have most likely turned to DFT since it provides sufficiently decent geometries for organic molecules to get some insight on steric effects and participating solvent molecules. In the magnifying glass analogy, DFT is a method with a medium-wide view and just about acceptable focus. While getting transition states and geometries qualitatively right, it struggles to achieve sufficient accuracy to decide which one of several competing reaction mechanisms is more likely to occur, particularly if these competing mechanisms are close in energy. DFT also usually fails at describing the bond making and bond breaking mechanism during a reaction. For that purpose, coupled cluster theory, would have been my method of choice. Coupled cluster theory is more accurate and more systematic than DFT and would have given me more control over the errors that I make. But coupled cluster theory is also less computationally efficient and it would have been nigh impossible to calculate more than a small part of the system. Coupled cluster theory is usually used to obtain quantitatively accurate results for geometries or reaction energies in small molecules, up to 30 atoms or so.

Initially, in the NISQ era, small quantum computers will be built on which we will only be able to calculate relatively small molecules. These small quantum computers are hence poised to compete with highly accurate methods such as coupled cluster theory. The first hurdle to clear is the accurate description of small molecules and reactions for which even coupled cluster theory fails. We are fairly confident that a couple of hundred

qubits with decent lifetimes could beat a conventional quantum chemistry computation in terms of obtaining chemical insight using the Variational Quantum Eigensolver as an algorithm. There are still some open questions regarding which noise level would be acceptable, what the actual runtime of such a calculation would be and how it would behave on different qubit types. But then consider what happens if the quantum version of Moore's law holds: Within 3 years, we will have increased our computing power 16-fold – we will be able to zoom out and compute the behaviour of large systems at no loss in accuracy! It won't take long until we have enough qubits to properly perform error correction and run long calculations on the quantum computer.

The increase in quantum hardware resources will not happen in a vacuum. Through better software and algorithms, we can understand how we can use the available hardware more efficiently. We can optimise the layout and geometry of the qubits for specific applications [7], design error mitigation techniques that increase the qubits' lifetimes [8] and optimally map an application to the complicated control system that sits on top of the qubits. [9]

And, so, I am fairly confident that by 2050 an organic chemistry student will be able to use a quantum computer to get more than just blurry insights into why their reaction isn't working. In fact, I think that the advent of quantum computers could make computational chemistry so central to gaining chemical understanding that it will be the first topic chemistry students learn about when they enter university. Of course, chemical rules and knowledge will be by no means obsolete. We still must know how to ask the right questions and devise sensible hypotheses – but testing them will be much less laborious and painful.

From today's perspective, the opportunities that quantum computers will afford are almost dizzying. Quantum computational methods will shake up the rigid process of small-molecule drug discovery, making it possible to accurately calculate enzyme inhibitor interactions for the most difficult metalloenzymes at scale. In materials science, we will be able to compute strongly correlated systems and surface reactions to design better solar cells or battery cathode materials. For catalysis and chemical processes, the kinetics and mechanisms of competing reactions can be disentangled, saving resources and mitigating pollution.

Ten years ago, when I personally first heard of quantum computing, the field was academic and the challenges enormous. But the trajectory and achievements since then make me confident that with grit, clever engineering and a fair bit of investment, quantum computers will have changed the world by 2050.

References

[1] Arute et al. *Nature* **574**, 505–510 (2019). https://doi.org/10.1038/s41586-019-1666-5

[2] The Quantum Algorithm Zoo https://quantumalgorithmzoo.org/

[3] Montanaro, A. *npj Quantum Inf* **2**, 15023 (2016). https://doi.org/10.1038/npjqi.2015.23

[4] Bauer, B et al. *Chem Rev* (2020), 0.1021/acs.chemrev.9b00829

[5] Cao Y et al. *Chem Rev* **119** (2019), 19, 10856–10915, https://doi.org/10.1021/acs.chemrev.8b00803

[6] Wang D *et al. Phys. Rev. Lett.* **122** (2019), 140504

[7] Murali P arXiv:1901.11054

[8] McArdle S et al. *Phys Rev Lett* **122** (2019), 180501

[9] Cruise J *et al.*, Practical Quantum Computing: The value of local computation arXiv:2009.08513

Leonie Mück

When I left school in 2004, I knew that I wanted to study a science, but I wasn't sure which one. The choice fell on chemistry since I felt that I could branch out in different directions from there. Through various serendipitous and not so serendipitous events, I specialised in computational chemistry. For my PhD, I chose to go into the very depths of this subject and develop highly accurate quantum chemical methods. Under the tutelage of Juergen Gauss at the Johannes Gutenberg-University of Mainz, I learned how to understand and derive quantum chemical equations and translate them into computer programmes. This proved hard for me: I hadn't done any programming before. My supervisor's patience as well as the support of my peers were instrumental in keeping me going.

The environment at Mainz was also stimulating in other respects: Juergen Gauss suggested that I join the Gutenberg Academy, a multi-disciplinary group of selected PhD students who met once a month to discuss topics spanning from anthropology to particle physics. I became interested in the meta-aspects of science: How does the scientific process work? Why do we value certain scientific results more highly than others? With support of my graduate school, Materials Science in Mainz, my fellow PhD students and I founded the Journal of Unsolved Questions, JUnQ for short, a multidisciplinary outlet for negative and null results. Working on JUnQ was pure joy. When I learned through my co-supervisor that one could be a scientific editor for a living, I was thrilled.

At the end of my PhD, in 2012, I applied to be an editor at the journal *Nature Communications* and got the job. Based in London, I was now responsible for handling all manuscripts in condensed matter physics and theoretical physics that were submitted. This required learning the basics of a variety of subjects in a very short amount of time, always keeping the big picture in mind. But human nature was the biggest challenge: Academics are not known for being the most mellow of people, in particular when it comes to publishing their work, and I learned how to keep a cool head in tricky, emotional situations.

I later switched to the journal *Nature* becoming a senior editor and later team leader. In terms of subjects, I was responsible for fundamental physics including quantum computing. While I loved my job at *Nature*, there was one thing missing: When working to launch JUnQ, I had caught the bug for building something from scratch. And so, I joined the Public Library of Science to build a Physical Sciences and Engineering division for one of their journals *PLOS ONE*.

While I was taking a maternity break to take care of my baby daughter, an old acquaintance, Steve Brierley, called me up to chat about his new quantum computing start-up Riverlane. I met him for a coffee, baby in tow. As it turned out, he wanted to recruit me to help him build his company and set the right milestones. I took the fact that my daughter pooed all over me during that meeting as a sign of good luck and have not looked back. Life at Riverlane is exhilarating. I'm surrounded by some of the smartest people in the world and get to shape a completely new field together with them.

Stimmen aus der Wissenschaft

Computational Catalysis 2050

Im Jahre 2050 wird Data Science als dritte Säule des Erkenntnisgewinns in den Naturwissenschaften unverzichtbar sein, neben Computation/Simulation und Experimentation.

In der Quantenchemie werden wir wissen, ob die theorie-basierte Suche nach besseren quantenmechanischen Näherungsverfahren, die durch clevere Algorithmen die Hardware optimal nutzen, obsolet geworden und durch Machine Learning Methoden überflügelt worden ist.

Wenn ich jetzt als Doktorand oder Postdoc anfinge, würde ich mich für Quantenchemie auf (realen) Quantencomputern interessieren. Damit könnte die Quantenchemie vielleicht die harte Wand tunneln, die der Lösung von sogenannten Multireferenzproblemen entgegensteht. Diese Methoden werden benötigt, um Systeme mit nahezu entarteten elektronischen Zuständen zu beschreiben, also z.B. Biradikale oder Übergangsmetallverbindungen. Die Physiker nennen das strongly correlated systems, wozu magnetische Kopplungen in Übergangsmetalloxiden gehören.

Übergangsmetallverbindungen dominieren in der Katalyse, und wenn das Gebiet Computational Catalysis die Limitierungen der nicht prediktiven Dichtefunktionaltheorie überwinden will, sind chemisch genaue (± 4 kJ/mol) Multireferenzmethoden für Systeme mit zwei Übergangsmetallatomen und einer entsprechenden Zahl von Liganden unverzichtbar. Gegenwärtige Multireferenzmethoden sind weder auf Systeme dieser Größe anwendbar, noch geht die Genauigkeit für die behandelbaren Systeme über 0,2 eV hinaus. Werden sich 2050 die auf Quantencomputer gesetzten Hoffnungen erfüllt haben, oder wurde auch bei Multireferenzproblemen ein Durchbruch mit Machine Learning erzielt?

In der Katalyseforschung hat sich das Paradigma der aktiven Zentren (aktive Phasen) in den letzten 50 Jahren als wissenschaftlich äußerst produktiv erwiesen. Im Jahre 2050 wird Computational Catalysis (ebenso wie das Experiment) dieses Paradigma hinter sich gelassen haben und Katalysator und katalytische Reaktion als eine sich dynamisch ändernde Einheit betrachten.

Joachim Sauer
Bunsen-Denkmünze 2020
1. Vorsitzender der DBG von 2015–2016

Stimmen aus der Wissenschaft

Quantencomputer und co. – Welche Lösungen hält die Grundlagenforschung für 2050 bereit?

Betrachtet man den Trend des Moore'schen Gesetzes, werden bis 2050 die Grenzen der konventionellen Transistortechnik und damit der Entwicklung leistungsstärkerer Computer erreicht. [1] Abhilfe sollen hier Quantencomputer schaffen – eine Lösung, die Ihren Ursprung in der physikalischen Grundlagenforschung hat und bereits heute getestet wird. [2] Die Entwicklung neuer Materialien und Methoden hierfür erfordert große interdisziplinäre Zusammenarbeit, v.a. in den Bereichen Physik, Chemie und Ingenieurswissenschaften. [3] Es wird spannend sein zu verfolgen, wie die Entwicklungen voranschreiten und die (Quanten-)Computer in 30 Jahren aussehen, welche Anwendungen sie finden werden und welchen Einfluss sie auf unser Leben und unsere Forschung haben werden. [4] Doch noch viel spannender ist vielleicht die Frage, welche Grundlagenforschungen von heute die (technologischen) „Quantensprünge" von morgen mit sich bringen werden...

Emiel Dobbelaar
Bundessprecher des JungChemikerForums der Gesellschaft Deutscher Chemiker e.V.

[1]   G. E. Moore, *Proceedings of the International Electron Devices Meeting (IEDM '75)*, 1975, **21**, 11–13.
[2]   F. Arute et al, *Nature*, 2019, 505-510.
[3]   E. Grumbling, M. Horowitz, *Quantum Computing: Progress and Prospects*, National Academies Press, Washington, DC, 2018, p. 9.
[4]   X. Yuan, *Science*, 2020, **6507**, 1054–1055.