Goals and opportunities in quantum simulation

J. Ignacio Cirac and Peter Zoller

The long-term promises of quantum simulators are far-reaching. The field, however, also needs clearly defined short-term goals.

Richard Feynman's presentation 'There's plenty of room at the bottom'[1] at the American Physical Society meeting in 1959 is widely acknowledged as one of the main inspirations for the development of nanotechnologies. In the same talk, Feynman also anticipated the possibilities that quantum mechanics can offer us in the microscopic world: "When we get to the very, very small world we have a lot of new things that would happen that represent completely new opportunities for design. Atoms on a small scale behave like nothing on a large scale, for they satisfy the laws of quantum mechanics. So, as we go down and fiddle around with the atoms down there, we are working with different laws, and we can expect to do different things."

The field of quantum information explores such 'different things' and tries to exploit the laws of quantum mechanics that involve the superposition principle to carry out computational tasks in a more efficient way than is possible with devices governed by classical physics. Experimental progress during the past decades has been extraordinary, and has enabled us to isolate single microscopic particles, to manipulate and control their internal quantum states, and to detect them with almost perfect fidelity. Those advances notwithstanding, a quantum computer is still a long-term goal, requiring the full control of a many-body system and, eventually, the implementation of sophisticated error-correction protocols to achieve fault tolerance. So, as we still have to wait for a fully fledged universal quantum-information processor, is it possible to build on the experimental advances made so far, and construct a device not quite at the level of complexity of a quantum computer, but one that still could perform some tasks that classical devices cannot?

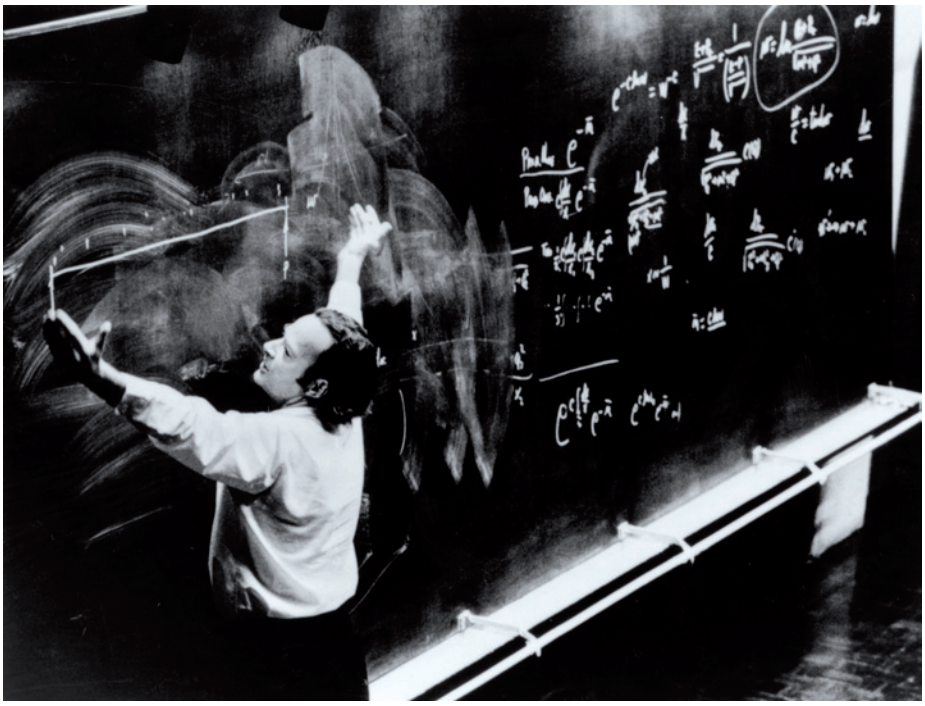

Vast opportunities. Richard Feynman — here during a lecture in 1970 — realized early on the possibilities offered by manipulating matter on an atomic level. In 1981 he formulated the idea to simulate one quantum system with another. This vision has now started to become an experimental reality.

A quantum for a quantum

The answer may be contained in another visionary lecture of Feynman. In his 1981 lecture 'Simulating physics with computers'[2], Feynman emphasized the complexity of simulating quantum systems using classical computers. The computational power required to even describe a quantum system scales exponentially with the number of its constituents. For instance, to describe the most general (pure) quantum state for N spin-½ particles, we have to store in our computer 2^N coefficients — a task that becomes practically impossible if N is greater than 50 or so. Moreover, to predict the value of a physical quantity, we will have to add, multiply or otherwise combine all of those coefficients, something that will require a time that scales exponentially with N as well. Feynman proposed that to overcome these problems, 'quantum simulators' that operate according to the laws of quantum mechanics should be used. If we had a system at our disposal, also composed of spin-½ particles, that we could manipulate at will, then we would be able to engineer the interaction between those particles according to the one we want to simulate, and thus predict the value of physical quantities by simply performing the appropriate measurements on our system.

To build a quantum simulator, Feynman proposed to discretize both space and

time. The problem at hand is specified in terms of a Hamiltonian, $H(t)$, which describes the (potentially time-dependent) interactions among the particles living on a lattice, an initial state $|\Psi(0)\rangle$ and a final time t. The goal is to determine certain physical properties of the final state $|\Psi(t)\rangle$. In 1996, Seth Lloyd proved[3] that evolving in small time steps would allow efficient simulation of any many-body quantum Hamiltonian containing few-particle interactions; that is, the overall time needed for the simulation would not grow exponentially with the number of particles, but polynomially. The main idea is to approximate each time step by a sequence of simpler operations according to a so-called Trotter decomposition, resulting in a sequence of quantum gates. Importantly, a quantum simulation might allow us to determine not only the dynamical behaviour of a many-body quantum system, but also (using adiabatic algorithms[4]) properties at zero temperature, corresponding to the ground state of a given Hamiltonian.

Focus on the real matter

Feynman envisioned a quantum simulator to be "a quantum machine that could imitate any quantum system, including the physical world"[2]. Such a general device would be as difficult to build as a quantum computer. However, if we are more modest and only demand our simulator to imitate certain physically interesting systems that cannot be simulated with classical computers, a quantum simulator may be easier to construct, but still would be an important device for the development of science and technology. Most of the time, we are not interested in the faithful representation of the complete many-body wavefunction, but rather in certain physical properties, typically related to intensive quantities — such as densities, magnetization per lattice site, or few-body correlations — which in turn determine the phase diagram. Moreover, a quantum simulator is expected to be more robust against imperfections than a quantum computer. Error correction, therefore, may not be needed. Roughly speaking, if one of the lattice sites undergoes an undesired perturbation, this may affect expectation values of intensive quantities by only a few per cent, as long as the lattice is sufficiently large. In fact, in many real materials that we would like to imitate with the simulator, those imperfections exist naturally, but nevertheless phases are typically robust against them.

Although Feynman's version is a digital quantum simulator with a stroboscopic time evolution, an analog simulation can also be considered, where the Hamiltonian $H(t)$ is built directly with a physical system that evolves continuously in time. A digital quantum simulator is more versatile, as it is able to simulate a wider range of Hamiltonians. For example, we can simulate directly N-body interactions; or, if a quantum simulator can simulate H_1 and H_2, it can also simulate $H_1 + H_2$. Analog simulators, in contrast, may be less prone to errors by avoiding those related to the Trotter expansion. Furthermore, simulations in continuous space, beyond those confined to a lattice, can also be considered.

Eyes on the prize

As we have emphasized, the main goal of quantum simulators is to solve problems that are not accessible to classical computers. Those will include many-body quantum systems, with many degrees of freedom, involving large-scale entanglement. However, among those problems there are still many that can be solved or simulated with classical computers. There are several powerful methods — Monte-Carlo and coupled-cluster methods, density functional theory, dynamical mean-field theory, density matrix renormalization group theory and others — that have been successfully employed to solve a wide range of problems. Nevertheless, there are still entire classes of many-body problems that cannot be tackled with those methods. The first class includes problems that deal with fermions or frustrated spins in two or higher spatial dimensions (where, for example, Monte-Carlo methods are limited by the so-called sign problem). The second class of problems involves dynamics, and includes most of the models that are not exactly solvable for sufficiently long times. This class of problems may involve closed or open systems; the latter refers to systems that are not isolated but interact with a reservoir with many degrees of freedom, resulting in dissipation.

It is hard, if not impossible, to prove that a specific problem cannot be simulated efficiently using classical computers — which may involve algorithms that have not been yet invented — but is amenable to a quantum simulator. Scientists working on computational complexity have long been searching for such kinds of proof, but without success. However, during the past year or so, work that builds on the synergy between quantum information science, condensed-matter physics and quantum field theory has suggested that problems that profit most from a quantum treatment are those involving a substantial amount of entanglement. Although in thermal equilibrium the entanglement cannot be arbitrarily high (owing to the so-called area law), in dynamical problems it can. Indeed, typically there will be a substantial amount of entanglement in out-of equilibrium scenarios, and a quantum simulator should prove particularly useful in those cases.

Problems of the kind mentioned above appear in different branches of science, including condensed-matter physics, quantum statistical mechanics, high-energy physics, atomic physics and quantum chemistry. A prominent example is related to high-temperature superconductivity: one of the important questions in this context is, which basic interaction between the electrons is responsible for the superconducting behaviour? In other words, which minimal Hamiltonian describes the phenomenon of high-temperature superconductivity? To answer this question, a quantum simulator could check the various candidate Hamiltonians for relevant phases. A quantum simulator may not only become a precious tool for tackling some specific problems in those fields, but should also prove a powerful instrument for developing, testing and benchmarking theoretical methods. For instance, very few theoretical techniques exist to describe the time evolution in a many-body quantum system after a sudden quench, even in one dimension. In fact, it is still an open question when such a system may thermalize, in the sense that the reduced state of a small connected sublattice may evolve to something that looks like the reduced state of a global thermal-equilibrium state. (The state of the whole lattice cannot evolve to the thermal-equilibrium state because it is a closed system). A quantum simulator may be essential in finding an answer to this question, but should also lead to the development of new numerical techniques to attack precisely this problem. Finally, by simulating we can also realize a system that may support exotic excitations that are useful in themselves, such as so-called anyons, which have a central role in topological quantum computing.

The road to beyond-classical simulations

We can try to be very explicit and distil a set of criteria that a quantum simulator must fulfil, akin to those for quantum computers[5]. Of course, as there are many different kinds of possible simulations[4], it is impossible to make a concise list including all of them. As we are interested in those that outperform classical devices,

we will concentrate on those capable of addressing questions in many-body quantum systems.

1. Quantum system. A quantum simulator should possess a system of bosons or/and fermions with or without internal degrees of freedom (pseudospin). The particles can be stored in a lattice or at least confined in some region of space. The system should contain a large number of degrees of freedom.

2. Initialization. A quantum simulator should be able to prepare (approximately) a known quantum state. Ideally, this should be a pure state, although in some situations it may be interesting to study the dynamics of a mixed state (for instance, a thermal state).

3. Hamiltonian engineering. It should be possible to engineer a set of interactions with external fields or between different particles, with adjustable values. These can be local (that is, acting among neighbouring particles) or have a longer range. They may involve a reservoir if simulation of open-system dynamics is required. Among the accessible Hamiltonians (or, in quantum open systems, Liouvillians) there should be some that cannot be efficiently simulated (at present) with classical techniques.

4. Detection. One should be able to perform measurements on the system. This could be individual (that is, addressing a few particular sites on the lattice) or collective (without the need of addressing any individual site). Ideally, one should be able to perform single-shot experiments that can be repeated several times. For instance, when measuring the collective magnetization in a set of spins along a given direction (corresponding to an observable S), one would be able to determine not only $\langle S \rangle$, but also $\langle f(S) \rangle$, where f is an arbitrary function.

5. Verification. Although by definition there should be no way of verifying if the result of the simulation is correct, as long as a problem that cannot be classically simulated is being solved, there should be a way of increasing the confidence in the result. For instance, the simulator could first be benchmarked with problems with known solutions. Or the evolution may be run forwards and backwards in time to check that it really ends up in the initial state. When an adiabatic algorithm is run, the technique may be used to adjust the time of the simulation by going back and forth with the parameters, each time getting closer to the target Hamiltonian. Alternatively, the results of simulations of different methods and systems could be compared.

There are several theoretical proposals to build quantum simulators with some specific purposes for different systems. We will not elaborate on any practical aspects here. Apart from the experimental challenges to build quantum simulators, this emerging field of research harbours a number of interesting theoretical questions. For example, we have argued that a quantum simulator should be more robust against decoherence than a quantum computer, at least for certain observables — or, even more crudely, when we are mainly interested in identifying the phases for different values of the external parameters governing the Hamiltonian.

Can we make a more rigorous statement of this sort? Seth Lloyd has already shown[3] how a quantum system may simulate the dynamics of another one efficiently by doing a stroboscopic change of the parameters. But how is this prone to errors? Classical chaotic systems are hard to simulate even classically owing to exponential sensitivity. We can ask an analogous question for certain non-integrable many-body quantum systems. Another important question is how to reach the ground state of the problem Hamiltonian or a thermal state. For the former, one can use adiabatic protocols. These methods may be useful in some situations, but new ideas of how to ensure and identify adiabaticity are needed. Above that, how can we determine for which kind of Hamiltonians the ground state can be found efficiently (in a time that scales 'mildly' with the number of particles or lattice sites)? In fact, this is exactly the problem computer scientists are interested in when they study complexity of quantum problems.

Quantum simulation is a field of great excitement, and one bearing great promises for the future. But the short-term goals should be clearly defined. As one of those goals, we propose that systems fulfilling the criteria laid out above are found, and, in particular, that simulation of a many-body quantum system involving (proven) large-scale entanglement that cannot be represented classically is demonstrated in the laboratory. This would make Feynman's vision a reality. ❐

J. Ignacio Cirac is at the Max-Planck-Institut für Quantenoptik, Hans-Kopfermann-Str. 1, Garching, D-85748, Germany. Peter Zoller is at the Institut für Theoretische Physik, University of Innsbruck and Institut for Quantum Optics and Quantum Information, Technikerstr. 25, A-6020 Innsbruck, Austria.
e-mail: Peter.Zoller@uibk.ac.at; ignacio.cirac@mpq.mpg.de

References

1. Feynman, R. P. *Engineering and Science* 22–36 (February 1960).
2. Feynman, R. P. *Int. J. Theor. Phys.* **21,** 467–488 (1982).
3. Lloyd, S. *Science* **273,** 1073–1078 (1996).
4. Buluta, I. & Nori, F. *Science* **326,** 108–111 (2009).
5. DiVincenzo, D. P. *Fortschr. Phys.* **48**, 771–783 (2000).