


Klaus Liegener is a quantum physicist at the Walther Meissner Institute in Munich. **Oliver Morsch** is a research director at the National Institute of Optics in Florence and a professor of quantum technologies at the University of Pisa, both in Italy. **Guido Pupillo** is a professor of quantum physics at the University of Strasbourg in France.




Solving quantum chemistry problems on quantum computers

Klaus Liegener, Oliver Morsch, and Guido Pupillo

One of the earliest applications that the new era of computing may be used for is the simulation of the quantum effects that drive chemical reactions.





FREDDIE PAGANI



quantum” may not be the first word that comes to mind when thinking about chemistry. But at the atomic level, the physical and chemical properties of molecules can be affected by quantum mechanical interactions. The spin of an electron, for example, is a property that is fundamentally quantum, and electron spin affects the reactivity of the molecule that the electron is part of. Every chemical reaction, therefore, is guided by the laws of quantum theory.

The modeling of quantum properties, however, is difficult, and thus the R&D costs for industries that need to understand the effects are extraordinary. A pharmaceutical company pays an average of about \$2.4 billion to develop a new drug. Much of that cost is in preclinical research because chemists usually rely on trial-and-error processes to test the suitability of roughly 1000 small molecules for clinical trials. Only about 10% of clinical trials result in successful products. Many chemical companies, therefore, are on a quest to improve the efficiency of modeling molecules.^{1,2}

One top priority in the field of computational chemistry is to replace the current trial-and-error approaches with direct modeling that has sufficient accuracy to resolve all relevant processes, such as reaction rates of molecules at room temperature. With such accuracy, a system’s behavior could be predicted, which would be a pivotal benchmark for quantum chemistry. Although many methods exist for modeling quantum effects on classical high-performance computers, they struggle to reach the required accuracy for certain types of chemical systems, particularly large systems of hundreds or thousands of molecules with either strong electronic correlations or coupled vibrational and electronic correlations.

Medium-sized strongly correlated systems of tens to hundreds of molecules, however, could soon be analyzed more efficiently with a quantum computer than with a classical computer (see figure 1). A small quantum computer with around 100 error-protected qubits, for example, may be capable of calculating the energy of around 100 spin orbitals, each of which is a wavefunction that specifies the probability of the position and spin of an electron.^{3,4} That capability would be useful for numerous applications, including the search for transition metals that serve as efficient catalysts in chemical processes.

Take the Haber–Bosch process, for example, which has been used for more than 100 years to fabricate ammonia for fertilizers. The process consumes about 2% of the world’s energy. In the long run, quantum chemistry could simulate how various catalysts could improve the nitrogen-reduction reaction and other reactions in the Haber–Bosch process to make it more energy efficient. For now, such simulations aren’t possible because of the large number of qubits required. An error-resilient simulation of the Haber–Bosch process would need around a million physical qubits. Many companies are working on building quantum computers with substantially larger qubit numbers. Some have boasted that the 1 million–qubit mark could be reached within this decade (see figure 2).

Quantum computing and quantum algorithms

The classical computing methods for chemistry often aim to investigate the energy landscape of a given system. A chemical system’s ground state—the lowest energy state of a given molecule—is a quantum wavefunction of the configurations of all nuclei and electrons. But to determine the ground state, researchers often investigate only a small subset of all possible configurations. The complete set is too large to be thoroughly analyzed on a classical computer within a reasonable time because of entanglement that’s created by the quantum superposition of some of the molecule’s states. To ignore many of those configurations and make the analysis more tractable, classical computing algorithms use several approximations, some of which are detailed in box 1.

In addition to a molecule’s ground state, dynamical processes that occur at finite temperature are important to describe. For strongly correlated systems and large systems




FIGURE 1. ELECTRONS IN A MOLECULE are described by quantum mechanics. A molecule such as benzene is relatively simple, and its various electron configurations can be represented in a quantum computer with qubits, which can take the usual classical values of not only 0 (spin up, in this example) and 1 (spin down) but also superpositions of those binary values (indicated in the illustration as tilted spins). The quantum computer itself is a controlled quantum system, and its qubits can be manipulated and the information extracted to produce simulations. In the future, the goal is to achieve simulations that are more accurate than those of a classical computer. (Image by Freddie Pagani.)

with several hundred atoms, the computations are difficult, and approximations will typically struggle to produce highly accurate results. The approximating tools are useful only in cases in which quantum effects play a negligible role—for example, in the calculation of diffusion and in the study of conformational processes, such as protein folding. Such calculations are not expected to be done on quantum computers anytime soon, given the large system sizes that need to be simulated.

The basic idea behind quantum computing is the replacement of the classical bit, which can take a value of 0 or 1, with the qubit, which can exist in a superposition of on and off states. Among all the possible states, some exhibit unique quantum properties: The measurement of one qubit's state may influence the probability of the measurement of another qubit's state. In that case, the qubits are entangled. So long as the qubits are resilient to noise and errors from external sources have been corrected, a quantum computer's storage capacity scales exponentially. Because two qubits can be found in a superposition of four possible states, three qubits can be found in a superposition of eight possible states, and so on, n qubits can be used to store 2^n bits of information. Once the quantum computer has around 60 qubits, its storage capacity quickly surpasses that of classical computers.

Some of the highly entangled states of the quantum computer can serve to efficiently describe chemical systems that are predominantly driven by quantum effects and that may lack a classical description. In fact, quantum computers may support fundamentally different approaches to computations and simulations. The so-called quantum advantage can be reached when quantum computers perform applications more quickly and efficiently than classical computers. (For

more on the design of quantum computers, see the articles by Anne Matsuura, Sonika Johri, and Justin Hogaboam, *PHYSICS TODAY*, March 2019, page 40; and by Harrison Ball, Michael Biercuk, and Michael Hush, *PHYSICS TODAY*, March 2021, page 28.)

For classical computers, a computation proceeds through logic gates, which use one or more binary inputs to produce a single binary output. A quantum computer works similarly but with quantum gates. In the future, the simulation of a chemical system could be accomplished by breaking down the individual operations in a simulation into, for example, elements of a universal quantum computing gate set. With that approach, many simulations could run on the same computer.

The approach, however, requires a substantial amount of qubits and other computational resources. And given the number of logical operations to be performed, the approach increases noise sensitivity and the risk of errors in the computation. Indeed, how faithfully quantum gates are executed on an actual device is still a critical challenge for quantum computation. Various factors, including unwanted interactions of qubits with the environment and imperfect qubit controls, can degrade a computation.

Preventive methods, such as the creation of multiple copies of stored information, have been specifically adapted to quantum computers and are called quantum error correction. Although the methods require large numbers of qubits—some early estimates have used 1000 qubits to obtain a single error-protected qubit—innovation in error correction is proceeding fast. Recent breakthroughs in error control^{5–7} have the potential to usher in the era of fault-tolerant quantum computing (FTQC). That's when quantum computers will be programmable and their results will be trustworthy, much like the

Box 1. Some classical computing methods

To simulate the vast number of electron configurations of a molecule, several approximations are valid in certain regimes. Take the Born–Oppenheimer approximation, for example. It treats the atomic nuclei and electrons in a molecule separately. Because the masses of nuclei and of electrons differ by three orders of magnitude, the nuclei can be analyzed as stationary points, which speeds up the molecular computation. Instead of looking at the positions of the electrons, the approximation looks at configurations in

which they can typically be found, in the so-called spin orbitals.

That method and others help researchers to find the ground state of a chemical system in the presence of weak interactions. But most chemically interesting systems—strongly correlated ones in which chemical bonds break or ones featuring dispersive electron interactions in solvation chemistry, for example—are described by wavefunctions that exhibit a high degree of entanglement. In those cases, the Born–Oppenheimer approxima-

tion and similar methods may break down.

An option in those situations is to compute the energies of the entire system in a procedure called exact diagonalization. Because of the rapid increase in computational power over the past several decades, tools that implement exact diagonalization can fully analyze small molecules, such as ammonia. But larger molecules, say with 10 electrons distributed over 50 orbitals or more, have too many configurations for the capabilities of modern computers.

classical computers of today. In fact, there are signs that the transition into the FTQC era is already happening.⁸ The most interesting time for all industries, especially the chemical industry, will be when the quantum advantage can be fully harnessed with large, fault-tolerant quantum computers.

Currently the industry is in the noisy intermediate-scale quantum (NISQ) era. NISQ devices can execute gate-based algorithms, although with limited accuracy. The gate-based approach could be skirted in favor of less demanding analog quantum simulations. Molecular dynamics simulations or other methods could bypass the gate decomposition and directly emulate the time evolution of a real system of interest. Whether the quantum advantage can be realized in the NISQ era with such an approach is currently being investigated.

No matter the chosen approach, a key element for advancing quantum computing specifically for chemistry is the ongoing development of quantum algorithms that bring together quantum information theory and classical techniques. Because an n -qubit quantum computer can manipulate and store 2^n bits of information at the same time, quantum algorithms can potentially work quite differently from classical computer algorithms. Such algorithms are, however, still in their infancy, and the field has not yet converged on a single approach that can guarantee quantum advantage. The current list of promising algorithmic methods for simulating quantum chemistry problems is long, as shown in figure 3, and highlights the amount of work that has already been done on the theoretical side to realize the next stage of quantum chemistry simulations. For more details about some specific approaches, see box 2.

Hardware

With many algorithms in place and in development, the at-

tention of the industry, eager to exploit the expected quantum advantage, has turned toward the hardware providers and the question of when sufficiently reliable quantum computing devices will become available. Several companies are relying on a variety of technical platforms, including spin qubits, topological systems, and photonic quantum computing, to develop quantum computers. The winning approach may involve the integration of several technologies. The market will likely follow a winner-takes-all mentality, but for now, all those technologies are under consideration.

The goal is to develop a universal device that is capable of executing all possible quantum algorithms. To prevail, any of the platforms being studied will need to have several issues addressed. According to a majority of the community, two of the technologies hold great promise: trapped atoms^{9,10} and superconducting qubits.¹¹ (See figure 4 for an overview.)

Among the first approaches that researchers used were laser cooling techniques, which individually trap atoms and cool them to temperatures less than 1 mK. When ionized, atoms can be trapped by electric fields; for neutral particles, optical tweezers can trap them. In both approaches, a qubit is made when two energy levels—usually the ground state and some excited state—are coupled by electromagnetic radiation, and the result is a single-qubit quantum gate. Two-qubit gates can be formed if the trapped atoms interact, either through electrostatic forces for trapped ions or through van der Waals or dipolar interactions for neutral atoms.

In recent years, 1D and 2D arrays of neutral atoms trapped in optical tweezers have demonstrated highly accurate multi-qubit operations and have become a powerful platform for quantum simulations and computations. In that approach, strong interactions between atoms in adjacent tweezers are

MILESTONES IN QUANTUM COMPUTING




FIGURE 2. GROWING QUBIT NUMBERS. Although the number of qubits alone is not sufficient to judge the capabilities of a quantum computer, many research developments have greatly increased how many qubits can be hosted in one. With enough high-quality qubits, researchers could work toward developing a universal quantum computer that could handle many more computations than a classical computer can. (Image by Freddie Pagani.)

controllably induced when the atoms are excited to Rydberg states, high-energy states with large principal quantum numbers and large spatial extents.

A key strength of Rydberg atoms as a platform is their demonstrated scalability. The number of qubits in such a system can range from hundreds to tens of thousands. By dynamically reconfiguring the Rydberg atom arrays, researchers can engineer arbitrary lattice geometries and connectivity for various types of quantum gates. The qubits can then be read out by fluorescence imaging of the single atoms with a highly sensitive CCD camera and an imaging system with micrometer resolution. The technology of reconfigurable atom arrays is poised to be instrumental in advancing the FTQC era in the next few years.

For algorithm-based applications, one main challenge has been the limited gate accuracy because of technical imperfections and physical effects that lead to decoherence. Recent breakthrough experiments, however, have demonstrated remarkable technological achievements, including the operation of two-qubit gates with accuracies of 99.5% and error suppression with logical operations made with tens of logical qubits,⁸ which are themselves made of many physical qubits.

Another promising technology platform, superconducting quantum circuits, takes a different approach than atoms to develop qubits. On a superconducting platform, similar to a classical computing structure, an electric circuit is nanofabricated on a silicon chip. Its components are classical elements, such as inductors and capacitors. But on top of those, the circuit integrates superconducting elements, so-called Josephson junctions. With those junctions, the circuit functions as a qubit that exhibits a transition frequency between the ground state and the excited state that's typically between 300 MHz and 5 GHz, depending on the precise design and architecture of the superconducting circuit. Since each superconducting quantum circuit is individually fabricated, they are often called artificial atoms. They offer a rich parameter space of possible qubit properties and operation regimes with predictable performance.

Many industrial players have already created large devices with hundreds of superconducting qubits and peak accuracies of 99.9% for two-qubit gates. The circuits, however, require expensive, complex nanofabrication technology for optimal results. Another limitation is that, like any other qubit, superconducting qubits are susceptible to noise: Temperature fluctuations, for example, may cause the qubits

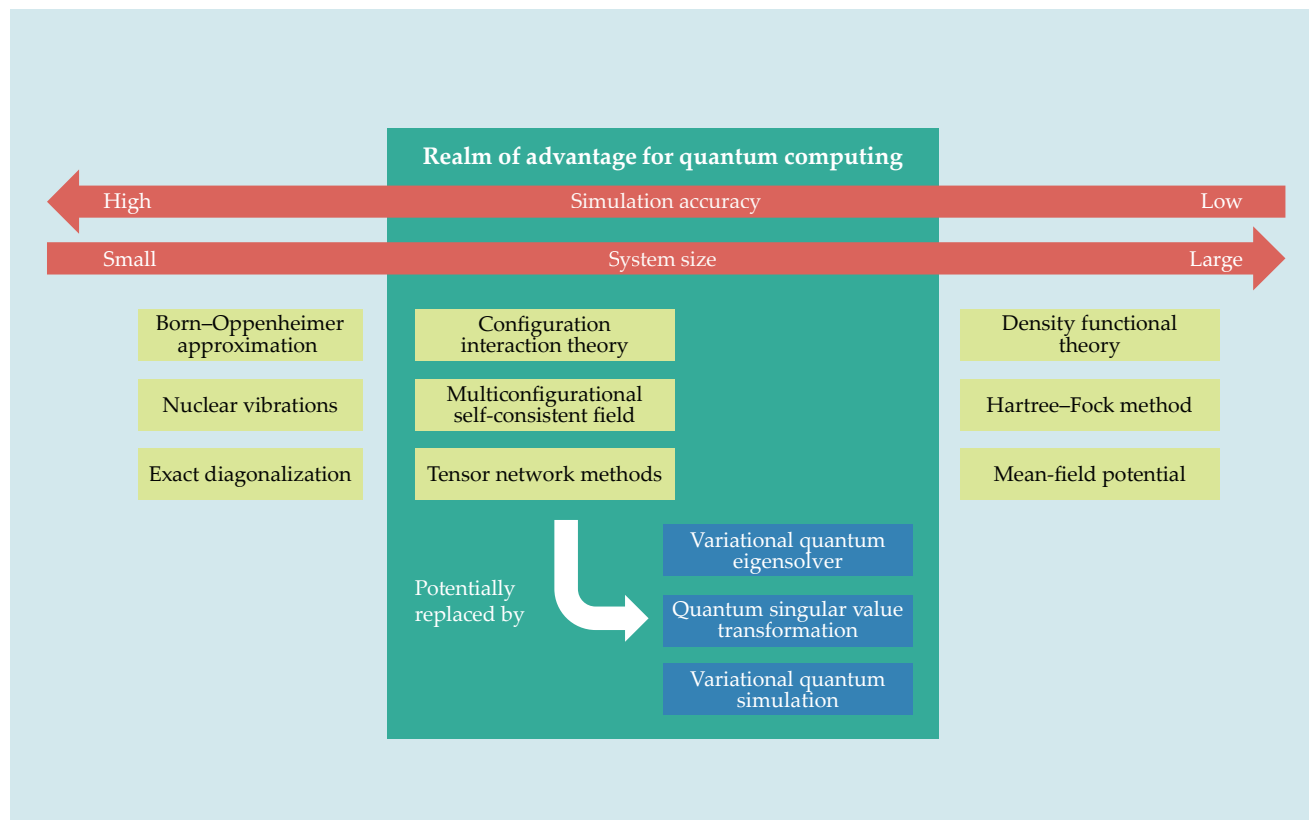


FIGURE 3. REALM OF ADVANTAGE. Several classical computation methods can simulate quantum systems, but the accuracy decreases as the system size increases (red arrows). For medium-sized chemical systems, with tens to hundreds of molecules, emerging quantum algorithms that could run on quantum computers could replace some classical methods and simulate certain chemical systems more efficiently. (Image by Freddie Pagani.)

to be excited to higher energy levels and introduce errors in computations. For that reason, any superconducting quantum processor unit needs to be operated at low temperatures, around 10 mK. (For more on error correction for superconducting qubits, see the article by José Aumentado, Gianluigi Catelani, and Kyle Serniak, *PHYSICS TODAY*, August 2023, page 34.)

With modern cryostats, it is possible to reach those low temperatures. Their maintenance, however, might be challenging if the system is scaled to a large number of qubits. Each qubit is typically controlled and read out via a dedicated wire, which connects the quantum processor with the control software outside the cryostat. To have thousands of cables to address each qubit individually is a major routing challenge because such a dense wiring arrangement causes unwanted heat radiation, which disturbs the qubits. Hence, ideas are emerging in the field for how to decrease the wiring complexity, such as using the same wire for multiple qubits, known as multiplexing, and using control electronics that work at cryogenic temperatures.

Although those and many more technical challenges apply to nearly all platforms, existing prototypes already

have demonstrated that, in principle, the cold-atom technology and the superconducting circuits work. Several companies are researching error-mitigation strategies and are encoding logical qubits on superconducting and trapped-atom platforms.¹² Even applications to chemical systems have been demonstrated on a superconducting quantum computer with 20 qubits.¹³ The achievement makes it the largest simulation of quantum chemistry on a quantum computer so far. With the proof of principle achieved, the central questions now are, What are the remaining steps needed to close the gap to real-life use cases, and Who is going to accomplish them?

The MOQS consortium

Quantum computing for quantum chemistry still faces many challenges, not least the high error rates and low qubit numbers of existing hardware. Considerable developments are needed to both build a powerful quantum processor unit and identify appropriate algorithms. The most prominent companies pushing the boundaries are tech giants in the US, such as Google, Microsoft, and IBM, and more than 100 quantum startups have emerged worldwide.

The quantum industry in the European Union is growing fast, predominantly in the academic sector. The EU trains the highest number of master's students in quantum tech-related fields: about 135 000, according to an April 2023 analysis by McKinsey & Company. To connect new knowledge leaders with the emerging quantum industry, the Horizon Europe program has funded the European Training Network for Molecular Quantum Simulations (MOQS).

The MOQS consortium includes several companies and research institutions from across the EU. With its expertise, the consortium trains about 15 PhD fellows annually in multiple disciplines, such as experimental physics and theoretical chemistry. In addition to searching for efficient quantum algorithms and tackling the technical challenges, the MOQS fellows work on a broad set of topics at the cutting edge of quantum computing research. In some cases they have already broken new conceptual ground.

At IBM Research in Zurich, Switzerland, for example, a MOQS fellow has investigated an alternative to the usual measurement in the computational basis of the logical values 0 and 1. Known as the positive operator-valued measurement, it has advantages for calculating the expectation values of operators, which are important in many quantum algorithms. The measurements had been assumed to need extra qubits for their implementation, which makes their realization harder and more resource demanding. The recent MOQS results show that positive operator-valued measurements can be efficiently implemented in existing superconducting quantum computing platforms without the use of additional qubits.¹⁴

At Eindhoven University of Technology in the Netherlands and the University of Strasbourg in France, MOQS fellows have looked into quantum logic gates based on interacting Rydberg atoms. The optimized laser pulses that have been developed¹⁵ should make it possible to reach gate accuracy of more than 99.9%, with 99.5% just recently demonstrated.¹⁶ In addition, new results from Ljubljana, Slovenia, are helping to unravel the complicated dynamics of interacting quantum many-body systems.¹⁷

A close connection to industry further guarantees that the doctoral students involved with MOQS obtain the necessary skills to propel the field of quantum chemistry forward. Major companies, including BASF, HQS Quantum Simulations, and Bosch, are involved in the consortium and ensure that the research projects lead to applicable use cases for




FIGURE 4. TWO QUANTUM COMPUTERS. (a) Cold atoms, trapped in an optical tweezer array in a vacuum cell, can serve as qubits in a quantum computer. Beamlets and mirrors focus laser light on the qubits, and the resulting interactions function as quantum gates that can perform logic functions. (Adapted from ref. 9.) (b) An electric circuit can transform into a quantum computer with the inclusion of superconducting materials, which can form qubits (highlighted in orange). Electromagnetic signals are sent through the circuit wires to implement quantum gates and to read the qubit states. (Courtesy of the Walther Meissner Institute.)

businesses. Only such close alignment can shorten the expected time frame to the quantum advantage.

Early adoption

The chemical industry's interest in quantum computing stems from the hope that quantum approaches might help in the development of new and better products at lower costs and in less time. Nonetheless, businesses will likely have to wait several years before quantum computers will

Box 2. Quantum computing algorithms

A few quantum computing algorithms have been developed, often as extensions of existing classical ones. For a long time, high expectations rested on an algorithm known as the variational quantum eigensolver (VQE) to address such problems as the search for the ground state of a chemical system.¹⁸ The VQE's design combines both classical and quantum computing methods. The hybrid approach is more practical because it avoids long run times, which increase error rates on today's error-prone quantum computers, and leverages the power of today's high-performance classical computers.

By switching between classical and quantum processors, the hybrid algorithm eventually trains the quantum computer in the following way: The VQE introduces free control parameters to a specified quantum algorithm. Then the algorithm is executed to prepare some desired ground state. The results from reading out the

final state and measuring its energy are used to update the control parameters. After several iterations, the control parameters are expected to converge on an optimal set that accurately prepares the ground state. Because of the optimization procedure, the VQE can work with simple algorithms and thus achieve higher accuracy than other quantum methods available today. The optimization procedure, however, scales exponentially with system size, so the VQE may not be suitable for quantum computers containing more than 100 physical qubits.

But quantum chemistry algorithms need to compute more than just system energies. The search for catalysts for various reactions, including the Haber–Bosch process for producing ammonia, for example, requires a careful study of how molecules react with each other over time, which cannot be easily implemented with standard VQE algorithms. Many methods have been

developed to implement the time simulation of a given system of interest.³ Two prominent examples are Trotterization and variational quantum simulation.

Trotterization breaks the time evolution of a chemical reaction into many tiny steps, such that the interactions of the system can be expressed by many quantum gate operations. Today the procedure is limited because quantum gates have imperfect accuracy. Variational quantum simulation, in contrast, relies on control parameters, like the VQE. The optimization of an algorithm that mimics the short time evolution of test states can converge to a suitable representation of the chemical system's total energy structure. The optimized algorithm can then be used to simulate long time evolution on complicated states in the Hilbert space and ultimately yield useful dynamical quantities, such as particle correlation functions, which reveal what phases of matter the system exists in.

create real benefit. As mentioned, it may be possible to simulate molecules with 100 spin orbitals on a 100-qubit quantum computer, but even that is still far from the long-term goal of the pharmaceutical companies. When they design new medicines and materials via quantum simulation, 95% of approved drug molecules are larger than 200 spin-orbital systems.³

Another factor to consider is error correction, in which up to 1000 physical qubits are combined to create a single error-resilient logical qubit. Only with error-resilient qubits can algorithms function properly and provide a quantum speedup over classical algorithms. Many recent advances have shown how to realize error correction on quantum devices with several hundred qubits.⁸ So the field may transcend the NISQ era sooner than anticipated.

The three of us are hopeful that the chemical industry can see a quantum advantage in computing toward the end of this decade—or maybe sooner if progress in technology proceeds at the currently staggering pace. In addition, molecular modeling is not the only area where quantum effects can be useful. Although chemical companies may see the benefits first, quantum computing stands poised to affect many other industries, including those that design and develop catalysts and new materials, batteries, semiconductors, magnets, and high-temperature superconductors.

We thank Daniel Bultrini and Christoph Hohmann for their help with figures and graphical representations. We also thank the Molecular Quantum Simulations consortium of the European Union's Horizon 2020 research and innovation program (grant number 955479) for its help and valuable comments.

REFERENCES

1. F. Budde, D. Volz, "The Next Big Thing? Quantum Computing's Potential Impact on Chemicals," McKinsey & Company (12 July 2019).
2. P. Gerbert, F. Ruess, "The Next Decade in Quantum Computing—and How to Play," Boston Consulting Group (15 November 2018).
3. S. McArdle et al., *Rev. Mod. Phys.* **92**, 015003 (2020).
4. Y. Cao et al., *Chem. Rev.* **119**, 10856 (2019).
5. U. Réglade et al., *Nature* **629**, 778 (2024).
6. S. Ma et al., *Nature* **622**, 279 (2023).
7. S. Bravyi et al., *Nature* **627**, 778 (2024).
8. D. Bluvstein et al., *Nature* **626**, 58 (2024).
9. M. Morgado, S. Whitlock, *AVS Quantum Sci.* **3**, 023501 (2021).
10. J. I. Cirac, P. Zoller, *Phys. Rev. Lett.* **74**, 4091 (1995).
11. P. Krantz et al., *Appl. Phys. Rev.* **6**, 021318 (2019).
12. Y. Kim et al., *Nat. Phys.* **19**, 752 (2023).
13. Google Quantum AI, *Nature* **614**, 676 (2023); T. E. O'Brien et al., *Nat. Phys.* **19**, 1787 (2023).
14. L. E. Fischer et al., *Phys. Rev. Res.* **4**, 033027 (2022).
15. S. Jandura, G. Pupillo, *Quantum* **6**, 712 (2022); M. Mohan, R. de Keijzer, S. Kokkelsmans, *Phys. Rev. Res.* **5**, 033052 (2023).
16. S. J. Evered et al., *Nature* **622**, 268 (2023).
17. Y. Kasim, T. Prosen, *J. Phys. A* **56**, 025003 (2023).
18. N. Moll et al., *Quantum Sci. Technol.* **3**, 030503 (2018).

PT