Digital Discovery

PERSPECTIVE

Check for updates

Extending quantum computing through subspace, embedding and classical molecular dynamics techniques

Thomas M. Bickley, [iD] †[a] Angus Mingare, [iD] †[a] Tim Weaving, [iD] †[a] Michael Williams de la Bastida, [iD] †[a] Shunzhou Wan, [iD] [a] Martina Nibbi, [iD] [b] Philipp Seitz, [iD] [b] Alexis Ralli, [iD] [ac] Peter J. Love, [iD] [cd] Minh Chung, [iD] [e] Mario Hernández Vera, [iD] [e] Laura Schulz [iD] ‡[e] and Peter V. Coveney [iD] ⋆[afg]

The advent of hybrid computing platforms consisting of quantum processing units integrated with conventional high-performance computing brings new opportunities for algorithm design. By strategically offloading select portions of the workload to classical hardware where tractable, we may broaden the applicability of quantum computation in the near term. In this perspective, we review techniques that facilitate the study of subdomains of chemical systems with quantum computers and present a proof-of-concept demonstration of quantum-selected configuration interaction deployed within a multiscale/multiphysics simulation workflow leveraging classical molecular dynamics, projection-based embedding and qubit subspace tools. This allows the technology to be utilised for simulating systems of real scientific and industrial interest, which not only brings true quantum utility closer to realisation but is also relevant as we look forward to the fault-tolerant regime.

1 Introduction

Quantum computers are a natural platform for simulating chemical systems as they can encode quantum states in linear space, rather than the exponential space required by classical devices. So far, they have been successful in realising small-scale demonstrations of electronic structure calculations. The variational quantum eigensolver (VQE) has facilitated simulations of up to 12 qubits,[2–35] while more recent developments in quantum-selected configuration interaction (QSCI) have unlocked scales up to the 77-qubit level.[36–44] For the first time in the chemical sciences, simulations in excess of 100 qubits are within sight. However, all the above works have been limited to gas-phase calculations, either to study small (often diatomic) systems in minimal atomic orbital basis sets, or modest active spaces of larger molecules and/or basis sets. In order to reach the milestone of quantum utility in this field, in which quantum computers can produce viable solutions to problems beyond the reach of exact solutions, there must be development in quantum-enhanced simulation of typical chemical workloads which include modelling the effects of large-scale environment regions such as solvents, biomolecules and surfaces.

The hybrid quantum mechanics/molecular mechanics (QM/MM) method is widespread, and allows one to situate a quantum mechanical calculation within a classical medium of point-charges, resolved using molecular mechanics (MM). The possibility of integrating quantum computational resources in QM/MM has been suggested,[45–50] but implementations on real quantum devices are scarce and do not make appreciable use of the hardware, nor provide any scalability guarantees.[51–54] It is becoming increasingly common to see quantum processing units (QPUs) being integrated with high-performance computing (HPC) platforms, bringing with it a need for computational workloads that challenge both quantum *and* classical resources. The QM/MM framework provides a realistic route towards achieving large-scale simulations that fully utilise the augmented capabilities of heterogeneous HPC + QPU systems.[55–58]

aCentre for Computational Science, Department of Chemistry, University College London, London WC1H 0AJ, UK. E-mail: p.v.coveney@ucl.ac.uk

bTechnical University of Munich, School of Computation, Information and Technology, Department of Computer Science, Boltzmannstraße 3, 85748 Garching, Germany

cDepartment of Physics and Astronomy, Tufts University 574 Boston Avenue, Medford, MA 02155, USA

dComputational Science Initiative, Brookhaven National Laboratory, Upton, NY 11973, USA

eLeibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities, Boltzmannstraße 1, 85748 Garching, Germany

fUCL Centre for Advanced Research Computing, Gower Street, London WC1E 6BT, UK

gInformatics Institute, University of Amsterdam, Amsterdam 1098 XH, Netherlands

† These authors contributed equally to this work.

‡ Current address: Argonne National Laboratory, 9700 S Cass Ave, Lemont, IL 60439, USA.

The motivation behind attempting such a computation lies in the exploration of the abilities of quantum computing in chemistry and material science, particularly in assisting calculations on systems which have complex electronic interactions within a large system. Prominent examples include catalysis and biomolecular systems. While current evidence for quantum advantage in stand-alone (i.e. full system) quantum chemistry simulations requires further work;[59] we believe that a reasonable application of near-term quantum hardware is the deployment of quantum algorithms in a small, highly-correlated region of a much larger classical QM/MM routine.

Sections 2, 3 and 4 are dedicated to reviewing techniques that can be used to reduce the quantum resources required to simulate a chemical system, many of which we utilise in a proof-of-concept demonstration presented in Section 5. Combined application of the methods discussed within this article allows quantum computation to be deployed within large-scale chemical simulation workflows, providing a practical route to scientific and industrial utility for the technology.

Section 2 discusses classical techniques for treating large-scale environmental effects, including QM/MM for general environments and continuum models designed specifically for solvents. In Section 3 we review two embedding techniques; projection-based embedding (PBE) is chemically-motivated and allows a QM calculation to be conducted at two different levels of chemical theory, while density matrix embedding theory (DMET) leverages the Schmidt decomposition in a similar vein as tensor networks to embed a subsystem within a surrounding bath. In Section 4 we review qubit subspace techniques that exploit (approximate) symmetries for additional resource reduction, such as qubit tapering and the contextual subspace method.

As a proof-of-concept for this approach, we present a workflow in Section 5 (also visually outlined in Fig. 1) where a QSCI simulation of the proton transfer mechanism in water, related to the structural debate over the aqueous form of the hydronium ion $[H_3O]^+$,[60] is performed. The results were obtained from the IQM QExa20 20-qubit superconducting device, integrated with the HPC cluster SuperMUC-NG at the Leibniz Supercomputing Centre (LRZ). While the QM/MM framework allows for the study of large chemical systems, realistic candidate structures for the QM subdomain are typically too large to be directly solved using quantum computers. One must therefore layer additional quantum embedding and qubit subspace methodologies to further distil the problem for feasibility on near-term quantum hardware.

2 Classical chemical environments

Methods to embed a more accurate quantum model within an extended classical region are motivated by prohibitive computational cost for full-system ab initio treatment. These methods allow for important classical effects to be included within quantum chemical calculations, such as the presence of complex molecular structures and surfaces, as well as solvation baths. Here we discuss two modes for treating these external environments; an explicit description via quantum mechanics/

molecular mechanics (QM/MM), and an implicit description via continuum models.

2.1 Quantum mechanics/molecular mechanics

The QM/MM method has emerged as a popular technique for studying many chemical systems where full treatment at the two individual levels of theory is either intractable or insufficient for modelling the desired properties. For example, this will be the case in biological systems where metal–ion coordination complexes contain regions of high electron correlation.[61,62] The method has also been used extensively throughout other fields in chemical/materials simulation, including photochemistry,[63] surface chemistry,[64] and condensed matter physics.[65] Even the use of favourably scaling quantum chemistry methods like DFT on systems like these can be challenging, and classical force-field approaches cannot capture all the relevant electronic effects.

Introduced in 1976 by Warshel and Levitt,[66] the idea of QM/MM is simple: treat a region of chemical interest with a high-accuracy, computationally-expensive quantum chemistry method, and the rest of the system is modelled with the less accurate but computationally cheaper MM. This partitioning clearly relies on the assumption that the targeted electronic effects within the QM region are mainly local and do not depend on long-range interactions with parts of the system within the MM region.

It is easy to find the individual energies of the QM and MM regions applied at their respective levels of theory, but including the interactions between the two regions is where the crux of implementing QM/MM lies. A simple approach is known as subtractive coupling, where the total energy is given by

$$E_{\text{full}}^{\text{QM/MM(sub)}} := E_{\text{QM}}^{\text{QM}} + E_{\text{full}}^{\text{MM}} - E_{\text{QM}}^{\text{MM}}, \tag{1}$$

where the subscripts indicate which region is included in the computation and the superscripts denote the level of theory applied to that region. As the only QM calculation in the subtractive scheme is applied to the QM region, the interactions between the QM and MM regions are treated at the MM theory level. It is easy to implement as the QM and MM codes do not need to communicate with each other, however there are some key disadvantages to this method including the inability to model the effect of polarisation of the QM region by the MM atoms.[61] The most common example of subtractive QM/MM is the ONIOM method.[67,68]

To capture some of these additional effects, the additive class of coupling methods can be used. In general, these methods consider the three types of interactions separately and sum them to get the total energy, for example

$$E_{\text{full}}^{\text{QM/MM(add)}} := E_{\text{QM}}^{\text{QM}} + E_{\text{MM}}^{\text{MM}} + E_{\text{full}}^{\text{QM/MM}}, \tag{2}$$

where the third term includes the elusive QM/MM couplings.[45] This coupling term can take several forms, most commonly:

- **Mechanical embedding:** the interactions are treated at the MM level, i.e. the usual MM force field parameters are used to model bonds, angles, torsions etc. between bonded QM and MM atoms and Lennard-Jones and Coulomb potential terms for non-bonded atoms.[69]
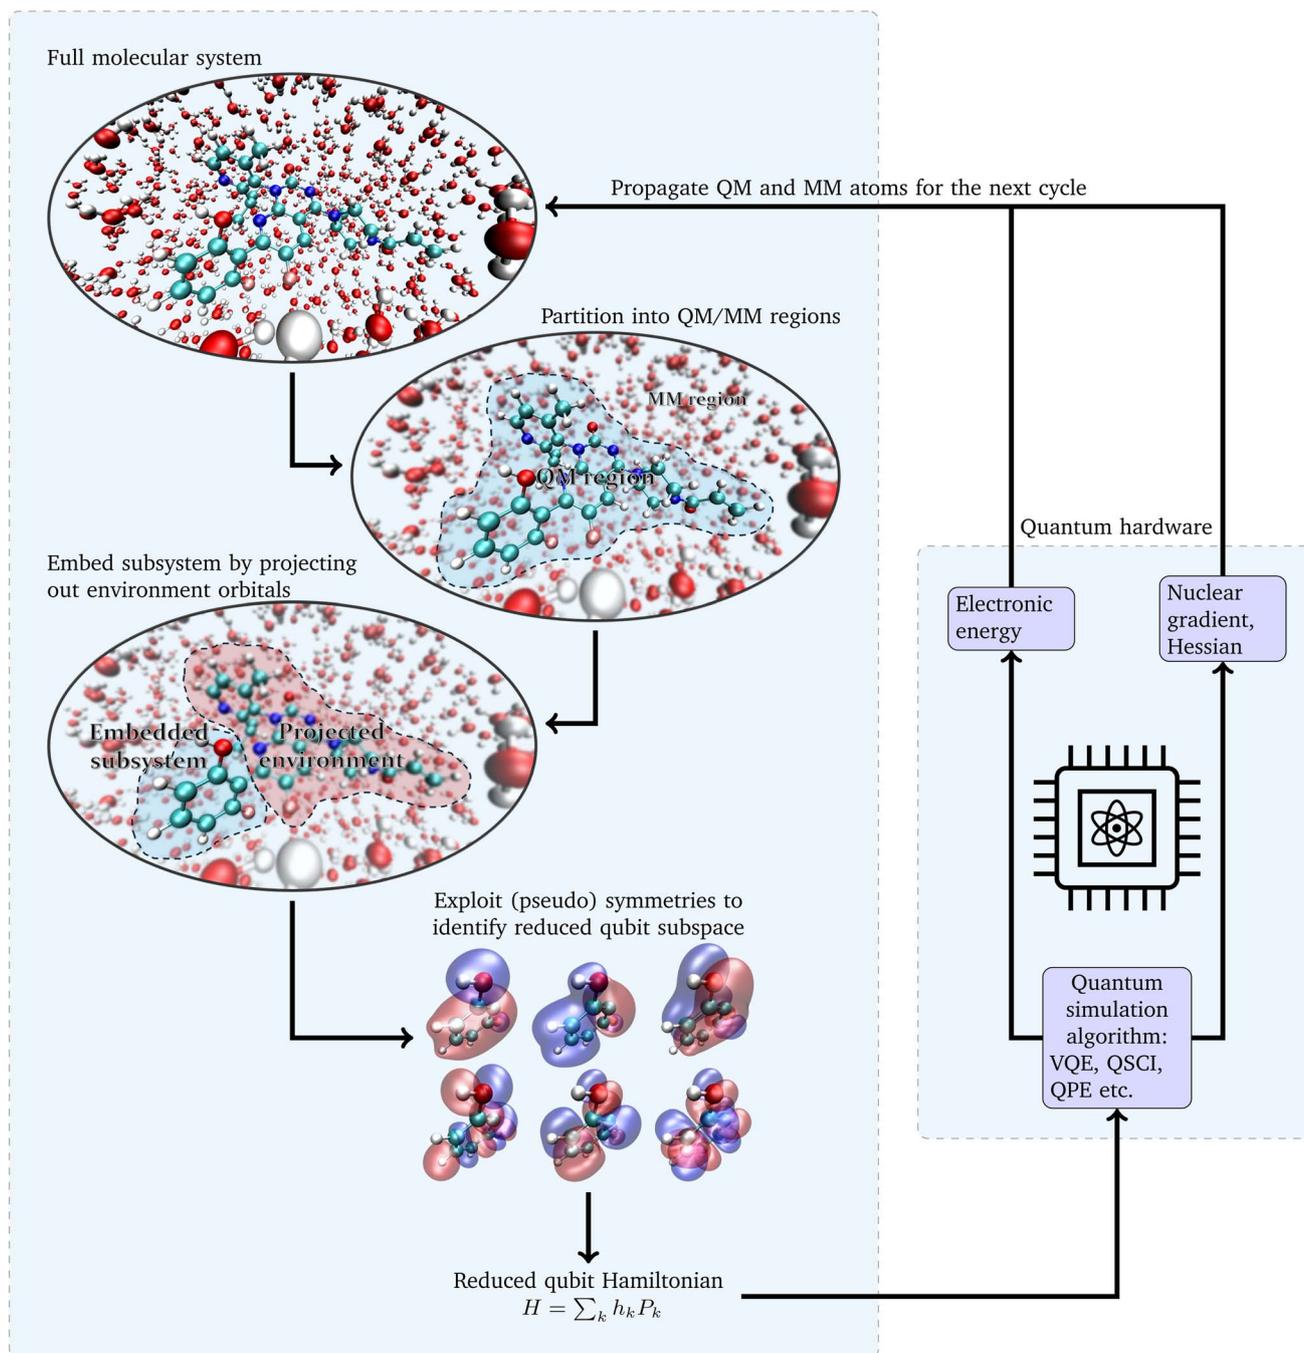
Fig. 1 Multiscale simulation workflow for embedding quantum computational capabilities within a surrounding classical molecular dynamics environment. The workflow consists of several nested layers of abstraction. At the highest level, we identify some molecular target entity within a larger system; while the former is resolved *via* quantum mechanics (QM) methods, the latter is treated at the classical molecular mechanics (MM) level for computational tractability. Within the QM region, the molecule is further partitioned into an active subsystem and surrounding environment *via* projection-based embedding (PBE), allowing a subdomain to be treated at a higher level of QM theory, while the environment is rendered at the density functional theory (DFT) level. Finally, within the embedded QM subsystem we may deploy qubit subspace techniques to further reduce the qubit overhead to utilise near-term quantum hardware in large-scale molecular simulation workflows. This allows us to leverage quantum processing units (QPUs) integrated with high-performance computing (HPC) platforms. Sotorasib molecule in water solvent drawn with VMD.[1]

- **Electrostatic embedding:** the QM Hamiltonian includes point charges from the MM environment as one-electron terms in the Hamiltonian, thus allowing one-way polarisability of the QM atoms by the MM atoms (see eqn (3)).
- **Polarisable embedding:** both the QM and MM regions are mutually polarisable and are solved in a self-consistent procedure.[70]

As an example, the Hamiltonian for the QM region under electrostatic QM/MM coupling (in atomic units) is

$$\hat{H}_{\text{QM/MM}} = \underbrace{-\frac{1}{2}\sum_{i}^{N_{\text{el}}} \nabla_i^2 + \sum_{i}^{N_{\text{el}}} v(\mathbf{r}_i) + \sum_{i<j}^{N_{\text{el}}} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|}}_{\mathbf{a}}$$

$$+ \underbrace{\sum_{A}^{N_{\text{MM}}} \sum_{B}^{N_{\text{QM}}} \frac{Q_A Q_B}{|\mathbf{R}_A - \mathbf{R}_B|}}_{\mathbf{b}}$$

$$- \underbrace{\sum_{A}^{N_{\text{MM}}} \sum_{i}^{N_{\text{el}}} \frac{Q_A}{|\mathbf{R}_A - \mathbf{r}_i|}}_{\mathbf{c}}, \tag{3}$$

where the three components **a**, **b**, **c**, correspond to the electronic Hamiltonian, QM-MM nuclear repulsion, and QM-MM electronic-nuclear attraction respectively. The other interaction energies related to the MM region are accounted for by the MM driver outside of this Hamiltonian.

In recent years, explorations of the potential use of quantum computing in QM/MM workflows have appeared in the literature. In a perspective piece by Blunt et al.,[45] the scaling of the fault-tolerant quantum phase estimation (QPE) algorithm was investigated within the context of large-scale biomolecular simulations. Whilst it is noted that algorithmic advances such as qubitization and Hamiltonian truncation have reduced predicted QPE runtimes by orders of magnitude, workflows relevant to pharmaceutical chemistry remain beyond the limits of reasonable predicted resource expectations. It is therefore noted that embedding approaches, specifically QM/MM and the similar QM-cluster approach will be essential for large scale chemical simulations even in the fault-tolerant regime. Similarly, the review by Capone et al.[47] offers some ideas on the future of multiscale chemical modelling with quantum computing, where embedding strategies including QM/MM are likely to have an essential role in enabling more chemically-relevant studies using quantum devices in the coming decades. Notwithstanding the many improvements needed in quantum hardware and algorithms, Santagati et al.[48] acknowledge the potential for quantum computing to enhance the precision of QM/MM like methods by targeting the quantum computer at the region of highest electron correlation in an extended system. Some simulated-VQE results were presented by Ma et al.[46] which made use of the many-body expansion method, which was chosen for its potential for integration into larger scale QM/MM workflows. In addition, polarisable embedding utilising a VQE subroutine to solve the wavefunction parameters and dipole moments self-consistently is introduced and performed on emulated hardware by Kjellgren et al.,[71] and subsequently emulated with GPU acceleration within CUDA-Q.[72]

There are also several examples of actual QM/MM workflows being performed on quantum hardware, where constrained active spaces are employed to make the calculations feasible on current devices. Izsák et al.[51] performed a 4 electron, 4 orbital (4e, 4o) active space energy evaluation of the enzyme ferredoxin hydrogenase and the photosensitizer temoporfin embedded within larger classical environments on Rigetti superconducting hardware, with VQE and iterative QPE algorithms respectively. Li et al.[53] studied a minimal (2e, 2o) active space of five atoms of a cancer target and drug interaction. This work embeds a VQE energy evaluation of this space on superconducting device within classical surroundings. Additionally, an embedded (6e, 6o) active space computation of the proton transfer step in carbonic anhydrase II was performed by Ettenhuber et al.[54] using both trapped-ion and superconducting devices.

Finally, recent work by Weisburn et al.[49] and Günther et al.[50] develops a three-level QM/QM/MM embedding scheme, where the two quantum regions are partitioned with bootstrap embedding. Quantum computations in these works are emulated, but they offer a potential path for multiscale chemical calculations with real quantum device assistance.

2.2 Continuum models

Explicit solvent models, which treat solvent molecules individually, can be computationally prohibitive at a quantum mechanical level, particularly for large systems. An efficient alternative is the use of continuum solvation models,[73] which describe the solvent as a continuous dielectric medium, thereby significantly reducing computational cost while capturing key solvation effects. A single solute molecule is immersed in an infinite solvent reservoir and treated at a homogeneous QM level. This solute can be a supermolecule composed of multiple molecules, including solvent molecules when appropriate. The use of continuum models enables efficient quantum mechanical calculations of the solute, allowing for the exploration of molecular properties in solution and providing valuable insights into solvent effects on structural stability, energetics and spectroscopy. Two widely used continuum solvation models are the polarisable continuum model (PCM)[74] and the conductor-like screening model (COSMO),[75] both of which are widely used in quantum chemistry and MD simulations for applications such as reaction mechanisms, spectroscopy, and drug design. This section discusses the principles, advantages and applications of these two continuum-based approaches.

The PCM describes the solvent as a polarisable dielectric continuum surrounding the solute.[74] The solute, typically treated at the quantum mechanical level, is embedded in a cavity constructed within the dielectric continuum. The solvent's response to the solute's charge distribution is modelled typically by solving the Poisson equation to compute the electrostatic potential at the cavity boundary. This potential induces a reaction field, which is incorporated into the solute's Hamiltonian, effectively accounting for solvation effects.

PCM methods are highly versatile and can be applied to a wide range of solvents and solutes. They are particularly

© 2025 The Author(s). Published by the Royal Society of Chemistry

effective for studying equilibrium properties, such as solvation free energies, solvent shifts in spectroscopy, electronic excitation energies, and reaction mechanisms in solution. For example, PCM has been employed within a time-dependent DFT to investigate the shifts in absorption and fluorescence energies in passing from apolar to polar solvent.[76] While the accuracy of the model depends on the choice of cavity construction and dielectric constant used to model the solvent environment, the flexibility of PCM allows for the inclusion of non-electrostatic contributions, such as dispersion and repulsion interactions, further improving the accuracy.[77] The computational cost of PCM scales with the complexity of the cavity and the dielectric response, which can be a limitation for large complex systems.

COSMO is another continuum solvation approach that approximates the solvent as a conductor.[75] The solvent's response is computed by assuming that the dielectric constant of the solvent is infinite. This approximation leads to a significant simplification of the electrostatic equations, improving numerical stability and convergence while keeping it computationally efficient. After the initial conductor-like screening, a scaling factor is applied to account for realistic dielectric effects. COSMO is particularly well-suited for rapid screening of solvation effects in larger molecular systems, such as drug-like molecules or materials. Its efficiency stems from the reduced complexity of the electrostatic problem, which avoids the need for iterative solutions of the Poisson equation.

An extension of COSMO, known as COSMO for real solvents (COSMO-RS),[78] integrates the continuum solvent approach and statistical thermodynamics to describe solvent effects beyond electrostatics, making it more predictive for thermodynamic properties. This has made COSMO-RS a particularly useful tool in applications such as solvation energy estimation, partition coefficients, and drug design.

Both PCM and COSMO provide efficient means to model solvation effects without the computational overhead of explicit solvent simulations. PCM is particularly well-suited for high-accuracy quantum chemical calculations due to its rigorous electrostatic treatment, while COSMO and COSMO-RS are advantageous in applications requiring stability and efficiency, such as large-scale screening studies. Recent developments have sought to combine the strengths of both methods, giving rise to hybrid approaches and extensions, such as the conductor-like modification of PCM (C-PCM).[79]

3 Quantum embedding methods

Quantum embedding methods enable the use of multiple independent quantum chemistry methods to directly solve a system.[80] Similarly to QM/MM, a system is partitioned into a region which requires a high level of theory and one which requires only a lower level. In doing so, it is possible to achieve results which are significantly more accurate than the lower level method alone, while being significantly less costly than a global application of the high-level method.[80,81] Implicitly, the Hamiltonian of the composite system is divided into multiple parts $H_{\text{system}} = H_A + H_B + H_{AB}$, the Hamiltonian of the individual systems and

also the non-additive part H_{AB}. Accurately representing this final term is challenging, and many methods have been developed to do so.[82,83] Importantly, there is a great degree of flexibility regarding both the partitioning of the system and the methods applied to each part. Within the context of quantum computing, embedding methods allow for quantum computing resources to be utilised in simulating parts of systems for which the whole would be impossibly large.[84] Clearly this is critical in the NISQ era, in which qubit counts and executable circuit depths are both very restrictive. However, as quantum hardware continues to develop, the flexibility of embedding methods will provide a straightforward path to fully utilise the resources which become available. Further, when fault-tolerant quantum computers are first realised it is likely that (at least initially) they will have few logical qubits and therefore suffer from the same restriction on admissible system size as current NISQ processors.[85] Quantum embedding methods may again be employed to fully utilise whatever fault-tolerant resources are available. We discuss two embedding methods: projection-based embedding (PBE) and density matrix embedding theory (DMET).

3.1 Projection-based embedding

Originated by Manby *et al.*,[86] projection-based embedding enables the use of a wavefunction method within DFT. Initially employed with classical methods,[87,88] quantum algorithms can straightforwardly be used.[84,89]

The method begins by selecting an active region and environment. In the original formulation, subsystems are manually selected at the level of atoms,[86] affording flexibility while requiring that chemical intuition is applied. In practice, it can be difficult to predict which selection is appropriate[82,90] and methods have been developed to perform this automatically.[90–92] Having defined a partition, whole-system DFT is performed to obtain a set of optimised molecular orbitals.

Electrons are then localised to the active and environment subsystems using any of a variety of standard localisation procedures such as IBO,[93] Pipek–Mezey[94] or SPADE.[95,96] Virtual orbitals can likewise be localised, for instance *via* VVO[97] or concentric localisation.[98] Where multiple system geometries are to be used, localising each individually may lead to changes in the number of active molecular orbitals and a resulting discontinuity in the potential energy surface.[90,92] Procedures have been developed to avoid this, although these naturally involve some compromise between geometries.[90–92,99]

Having localised electrons into the two regions, we may express the DFT energy in terms of the electron densities of each subsystem γ^{act} and γ^{env} as[95]

$$E[\gamma^{\text{act}}, \gamma^{\text{env}}] = \underbrace{\text{Tr}(\gamma^{\text{act}} h_{\text{core}}) + g(\gamma^{\text{act}})}_{\text{energy of isolated active system}}$$
$$+ \underbrace{\text{Tr}(\gamma^{\text{env}} h_{\text{core}}) + g(\gamma^{\text{env}})}_{\text{energy of isolated environment system}}$$
$$+ \underbrace{g(\gamma^{\text{act}}, \gamma^{\text{env}})}_{\text{non-additive two-electron energy}}, \qquad (4)$$

where g corresponds to the electron interaction terms, with $g(\gamma^{\mathrm{act}}, \gamma^{\mathrm{env}}) = g(\gamma^{\mathrm{act}} + \gamma^{\mathrm{env}}) - g(\gamma^{\mathrm{act}}) - g(\gamma^{\mathrm{env}})$.

Conspicuously absent from the above is the non-additive kinetic term. Having localised the electrons to subsystems, the environment orbitals are projected out of the Hamiltonian, thus suppressing transitions from the active region. The Fock operator of the active region is augmented with a projection term $P_{\mathrm{proj}}^{\mathrm{env}}$, and a term V_{emb} which includes the mean-field effect of the environment on the active region,

$$F_{\mathrm{emb}}^{\mathrm{act}} = h_{\mathrm{core}} + V_{\mathrm{emb}} + P_{\mathrm{proj}}^{\mathrm{env}} + g(\gamma_{\mathrm{emb}}^{\mathrm{act}})$$
$$= h_{\mathrm{emb}} + g(\gamma_{\mathrm{emb}}^{\mathrm{act}}). \tag{5}$$

Two forms of projector are typically used, the μ-shift $(P_\mu^{\mathrm{env}})_{ij} = \mu[S\gamma^{\mathrm{env}}S]_{ij}$[86] and Huzinaga $P_{\mathrm{huz}}^{\mathrm{env}} = -\frac{1}{2}(F\gamma^{\mathrm{env}}S + S\gamma^{\mathrm{env}}F)$.[100] Here, $S_{ij} = \langle \psi_i | \psi_j \rangle$ is the overlap of the atomic orbital basis. The effect of P_μ^{env} is to take the environment orbitals to a high constant energy (typically 10^6), while $P_{\mathrm{huz}}^{\mathrm{env}}$ sends the negative energy levels to their opposite value. Note that the Fermi-shifted Huzinaga projector can account for orbitals with initially positive values.[101] The Huzinaga projector is constructed such that it commutes with the Fock operator, as a result, it gives more precise energies.[102] Note that the PBE procedure is performed with only a single-shot embedding, with no need for computationally expensive feedback between classical and quantum methods.

With the environment frozen, the active region is self-consistently optimised under this new Fock operator, returning a wavefunction for the embedded active region $|\Psi_{\mathrm{emb}}^{\mathrm{act}}\rangle$. Using a selected wavefunction method, which may be run on a quantum device, the energy of the embedded region can be calculated. Taking $H_{\mathrm{emb}} = h_{\mathrm{emb}} + g(\Psi_{\mathrm{emb}}^{\mathrm{act}})$, where the second term is again a two-electron term but now acting upon the embedded active wavefunction. To correct for double counting of the Coulomb term resulting from the

Fig. 2 Bond dissociation of perfluoromethane in STO-3G basis. In blue and pink are the whole-system Hartree–Fock and density functional theory (B3LYP). Orange gives the whole system CCSD energy. μ-Shift embedded CCSD-in-DFT energy is given by yellow squares, while purple crosses show the embedded FCI-in-DFT energy.

environment's effect on the active region, and to negate the energetic effects of projection, a correction term must be included in the final total,[95]

$$E[\Psi_{\mathrm{emb}}^{\mathrm{act}}; \gamma^{\mathrm{act}}, \gamma^{\mathrm{env}}] = \underbrace{\langle \Psi_{\mathrm{emb}}^{\mathrm{act}} | H_{\mathrm{emb}} | \Psi_{\mathrm{emb}}^{\mathrm{act}} \rangle}_{\text{active region}}$$
$$+ \underbrace{E[\gamma_{\mathrm{env}}]}_{\text{environment}} + \underbrace{g(\gamma^{\mathrm{act}}, \gamma^{\mathrm{env}})}_{\text{non-additive}}$$
$$- \underbrace{\mathrm{Tr}\left(\gamma^{\mathrm{act}}\left(V_{\mathrm{emb}} + P_{\mathrm{proj}}^{\mathrm{env}}\right)\right)}_{\text{correction}}. \tag{6}$$

The active region wavefunction method may now be run completely independently. Any quantum simulation algorithm can be used, with the other terms derived from classical computation handled as an energy constant. We provide an example of the method in Fig. 2, showing the bond dissociation energy of perfluoromethane. In this example, the absolute value of the bond energy is marginally more accurately predicted by DFT alone; however, it serves to illustrate the reduction in problem size achieved by PBE. The CF_3 molecule has 50 spin-orbitals in the STO-3G basis, which is reduced to 28 by embedding.

With currently available quantum hardware, applications have so far been limited to simple demonstrations with only a few atoms.[84,89] Quantum advantage is still required to necessitate the use of a quantum processor for the embedded wavefunction method over existing classical methods. We therefore expect quantum-in-classical PBE to become a common technique in the future.

3.2 Density matrix embedding theory

Density functional methods such as PBE struggle to elucidate entanglement information between the system and its environment.[103] More sophisticated methods overcome this by replacing the single-particle density with a quantum variable that is better suited to capture entanglement. For example, in condensed matter, Green's function methods such as dynamical mean-field theory are popular.[104–107] However, Green's functions methods often require very large bath spaces to incorporate non-local interactions.[108] Together with the difficulty of dealing with time-dependent quantities such as the self-energy, this has limited the application of Green's function methods, particularly in quantum chemistry.[109] Density matrix embedding theory (DMET) as introduced by Knizia and Chan[110] was designed to overcome the challenges of Green's functions methods by only dealing with the single-particle density matrix. Furthermore, it was inspired by ideas from tensor networks to efficiently capture entanglement information.

DMET begins by computing an approximate ground state wavefunction for the full system, $|\Psi_0\rangle$, for example using a truncated configuration interaction theory[111] or anti-symmetrised geminal power wavefunctions.[112] For each subsystem that is to be treated at a higher level of theory a set of bath orbitals are computed from the low-level density matrix. A
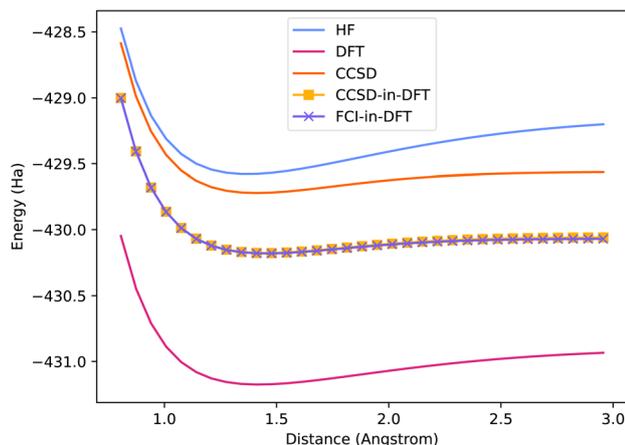
simple argument based on the Schmidt decomposition of the ground state solution $|\Psi\rangle$ shows that the size of the bath system need be no greater than the size of the subsystem under consideration. Letting $\{S_i\}_{i=1}^{d_S}$ be a basis for the subsystem and $\{B_i\}_{i=1}^{d_B}$ be a basis for the bath we can write

$$
\begin{aligned}
|\Psi\rangle &= \sum_{i=1}^{d_S}\sum_{j=1}^{d_B}\psi_{ij}|S_i\rangle\otimes|B_j\rangle \\
&= \sum_{i=1}^{d_S}\sum_{j=1}^{d_B}\sum_{k=1}^{\min(d_S,d_B)}U_{ik}\Sigma_{kk}V_{kj}^{\dagger}|S_i\rangle\otimes|B_j\rangle \qquad (7)\\
&= \sum_{k=1}^{\min(d_S,d_B)}\Sigma_{kk}|\tilde{S}_k\rangle\otimes|\tilde{B}_k\rangle.
\end{aligned}
$$

The set of subsystem and bath orbitals then defines an embedded Hamiltonian to which a high level of theory is applied to obtain the subsystem density matrix ρ_A and energy contribution E_A. The global density matrix $\rho = |\Psi_0\rangle\langle\Psi_0|$ is then optimised self-consistently with respect to some pre-defined cost function designed to match properties of the global state with properties obtained from the collection of subsystems. This process repeats until convergence is achieved. For a full introduction to DMET we refer the reader to Wouters et al.[109] A related but distinct class of methods exists, known as bootstrap embeddings, which instead enforces consistency by directly matching overlapping fragment density matrices.[113–117]

In the usual case of classical-in-classical embedding, standard high-level methods such as coupled-cluster theory or the density matrix renormalisation group algorithm can be applied to the embedded Hamiltonian. However, as the size of the active space increases these methods must trade off accuracy and computational tractability. One solution is to treat many subsystems at the higher level of theory although this increases the difficulty of the self-consistent optimisation. An emerging solution is to treat the subsystem on a quantum computer using a quantum algorithm for ground state computation which handles the embedded Hamiltonian.

As these nascent devices develop they may be able to extend the utility of DMET by allowing for larger active spaces. This idea has been proposed and numerically verified for small systems.[118–121] Additionally, several experiments run on real quantum hardware have yielded results matching classical benchmarks for DMET. Limited to small systems, the combination of DMET with VQE has facilitated simulations of a Hubbard lattice[122] and hydrogen rings.[16] Combining DMET with QSCI allowed a simulation of cyclohexane,[40] which used 32 qubits on a superconducting quantum chip. This proof-of-concept demonstration adds further support for the hope that QSCI may overcome the limitations of VQE; however, it remains clear that further hardware and algorithmic developments will be required before quantum computers are able replace classical methods within the DMET framework.

One key limitation for quantum computers is that DMET is iterative and requires the one-particle reduced density matrix of the active space to enforce self-consistency with the environment. While this explicitly captures entanglement information between the active space and environment, often leading to more accurate energy predictions than one-shot embedding methods, it requires computing many more expectation values and performing expensive quantum-classical feedback. This quantum overhead may be partially alleviated by replacing DMET with a bootstrapping embedding as suggested by Liu et al.,[123] however the suggested quantum subroutines preclude their implementation on near-term devices.

4 Qubit subspace methods

Dimensionality reduction is a common challenge in both conventional and quantum approaches to molecular electronic structure. Methods such as frozen core approximations,[124] active spaces[92] or virtual orbital truncation techniques like Frozen Natural Orbitals (FNO)[125] are commonplace. If used correctly, such methods can retain chemically relevant information at a reduction of computational overhead. These approximations typically rely on physically-motivated assumptions about electronic structure, allowing for systematic removal of orbitals or excitations that are expected to contribute minimally to the correlation energy.

Qubit subspace methods adopt a more abstract approach to yield hybrid orbitals that, while there might be a loss of physical motivation, come with the benefit of additional information being encoded in each qubit via the mixing of molecular orbitals to form an entangled basis in which to describe the molecular system. After application of this orbital-mixing unitary, we apply single-qubit projection operators to fix the state of desired qubits and trace them out of the system. This may be thought of as freezing the new hybrid orbitals, although it also has the flexibility to describe different qubit bases; more generally speaking, we project onto a stabiliser subspace of the qubits.

By contrast, in a method such as Complete Active Space Self-Consistent Field (CASSCF), we still apply a unitary to optimise the orbital basis at each step, but the form of this unitary is restricted to single-particle rotations and thus cannot reach as rich a class of molecular orbital bases. In CASSCF, the orbital transformation preserves its class structure, meaning we maintain a strict separation between inactive, active, and virtual orbital spaces, whereas this does not generally apply for qubit subspace methods.

The projection operator for a single qubit indexed $q\in\mathbb{N}$, stabilised by a Pauli operator $P\in\{X,Y,Z\}$, is $\mathbb{P}_{\pm}^{(q)} = \frac{1}{2}[I^{(q)}\pm P^{(q)}]$. For a subset of qubits with indexing set $\mathcal{I}\subset\mathbb{Z}_N$, where $N\in\mathbb{N}$ is the total number of qubits, the projection onto the corresponding stabiliser subspace takes the form

$$
\mathbb{P}_{\mathcal{I},\boldsymbol{\sigma}} := \bigotimes_{q\in\mathcal{I}}\mathbb{P}_{\sigma_q}^{(q)}, \qquad (8)
$$

where $\boldsymbol{\sigma}$ denotes the sector that defines the eigenspace.

Moreover, if one wishes to first apply a unitary prior to application of the single-qubit projectors, we may define the rotated projection

$$\mathbb{P}'_{\mathcal{I},\boldsymbol{\sigma}} := U^\dagger \mathbb{P}_{\mathcal{I},\boldsymbol{\sigma}} U = \frac{1}{2^{|\mathcal{I}|}} \bigotimes_{q\in\mathcal{I}} \left[I^{(q)} + \sigma_q U^\dagger P^{(q)} U \right]. \tag{9}$$

Finally, the reduced qubit subspace is obtained *via* the map

$$H \mapsto \mathrm{Tr}_{\mathcal{I}}\left(\mathbb{P}'_{\mathcal{I},\boldsymbol{\sigma}} H \mathbb{P}'_{\mathcal{I},\boldsymbol{\sigma}} \right), \tag{10}$$

where $\mathrm{Tr}_{\mathcal{I}}$ is the partial trace over the qubits indexed by $\mathcal{I}$. We note that, due to the cyclicity of trace operations, one may equivalently view this map as a rotation of the operator H, followed by single-qubit projections:

$$H \mapsto \mathrm{Tr}_{\mathcal{I}}(\mathbb{P}_{\mathcal{I},\boldsymbol{\sigma}} U H U^\dagger \mathbb{P}_{\mathcal{I},\boldsymbol{\sigma}}). \tag{11}$$

From an implementation point-of-view, it is typically beneficial to adopt this convention.

This framework is very general, for example if U diagonalises H, then this corresponds with solving the problem exactly. Instead, one must design U in such a way that it is classically efficient to realise the rotation $H \to UHU^\dagger$, for example by enforcing that U is (near) Clifford. Qubit subspace techniques are differentiated *via* the way in which we choose the rotation unitary U, qubit indices $\mathcal{I}$ and the eigenspace sector $\boldsymbol{\sigma}$. In the following subsections we explore several approaches.

4.1 Frozen core

In quantum chemistry, the canonical molecular orbitals (MO) and corresponding energies are the eigenvectors/values of the Fock matrix, optimised to self-consistency *via* Hartree–Fock. This is the standard approach to building the second-quantised electronic structure Hamiltonian

$$H = \sum_{p,q} h_{p,q} a_p^\dagger a_q + \sum_{p,q,r,s} h_{p,q,r,s} a_p^\dagger a_q^\dagger a_r a_s, \tag{12}$$

where $h_{p,q}, h_{p,q,r,s} \in \mathbb{R}$ are one and two electron integrals. The Hamiltonian is subsequently mapped onto qubits *via* a fermion encoding scheme such as Jordan–Wigner[126] or Bravyi–Kitaev.[127]

In Fig. 3 we plot the canonical MO energies for benzene, C_6H_6, in a minimal STO-3G atomic orbital basis set. One notes that the core orbitals lie deep with a potential energy well which is typical for most chemical systems. The implication of this is that it is energetically unfavourable for electrons lying deep
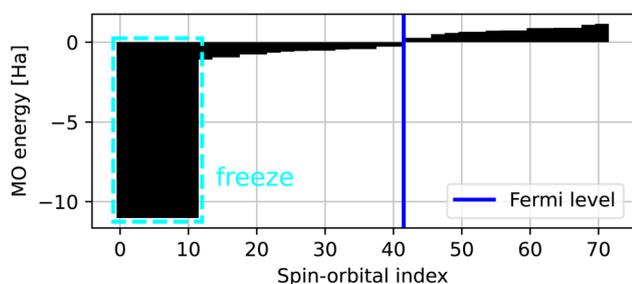

Fig. 3 Benzene (C_6H_6) STO-3G molecular orbital energies computed with the restricted Hartree–Fock method. The lowest twelve spin-orbitals may be frozen without dramatically affecting ground-state energy estimates.

within this well to be excited into the valence space. Instead, it is more likely that we will observe the greatest electronic activity around the Fermi level, where we find the gap between the highest occupied (HOMO) and lowest unoccupied (LUMO) molecular orbital. This motivates the notion of a frozen core approximation, in which electrons occupying the lowest-lying MOs are frozen in-place and not allowed to be excited into higher energy states.

Viewed as a qubit subspace technique, this is the simple case where $U = I$ is the identity, and $\mathcal{I} = \{q \in \mathbb{Z}_N : \mu_q < -\delta\}$ where $\boldsymbol{\mu}$ is the vector of canonical MO energies and $\delta > 0$ some energy threshold parameter to truncate the MO orbitals below the core potential. The sector is selected as $\sigma_q = -1 \quad \forall q \in \mathcal{I}$ to enforce that core orbitals are occupied. One may also freeze the valence space in a similar way by instead setting $\sigma_q = +1$ to project onto unoccupied orbitals and truncating the highest-energy MOs, rather than the lowest.

4.2 Qubit tapering

In physics, symmetries correspond to conserved quantities in a system of interest.[128] Given a Hamiltonian H, a symmetry is any operator S such that $[H, S] = 0$. For example, in quantum chemistry this may relate to particle number or spin symmetry. The presence of symmetry typically presents opportunities for simplification in some sense, from the exploitation of our knowledge of the problem structure that arises from that symmetry. In chemistry, molecular symmetries guide the construction of better ansatz circuits,[129] or may be used for the purposes of error mitigation through symmetry verification.[130,131] However, one may also exploit physical symmetries as a qubit subspace method.

A subset of symmetries that is of particular interest here are those of $\mathbb{Z}_2$-type, namely operators that describe a form of 2-fold symmetry. The $\mathbb{Z}_2$ symmetries possess a useful property such that, if $[H, S] = 0$ and S is $\mathbb{Z}_2$, then S commutes with every term of $H = \sum_k h_k P_k$ individually, *i.e.* $[P_k, S] = 0 \ \forall k$. This fact leads us to a mechanism for reducing the number of qubits in the Hamiltonian without sacrificing any accuracy, since the full and reduced Hamiltonians are isospectral up to a change in eigenvalue multiplicities.

From the theory of stabilisers,[132] given an independent set of N-qubit commuting Pauli operators $\mathcal{S}$, there exists a Clifford rotation C, a set of qubit indices $\mathcal{I} \subset \mathbb{Z}_N$ and bijective map $f : \mathcal{S} \mapsto \mathcal{I}$ such that, for each element $S \in \mathcal{S}$, we have $CSC^\dagger = P^{(f(S))}$ for a single-qubit Pauli operator $P \in \{X, Y, Z\}$. In other words, the unitary C maps elements of the set $\mathcal{S}$ onto distinct qubit positions. The positions must be distinct due to the requirement that $\mathcal{S}$ be independent, namely that no single element of the set is a product of other elements, and commuting, so that it is not possible to rotate X and Z onto the same qubit position, say. We may also assume without loss of generality that $P = Z$, since conjugation by Hadamard and phase gates relates the three choices.

Now, suppose we take $\mathcal{S}$ to be the set of $\mathbb{Z}_2$ symmetries of a Hamiltonian H. Then, since for each Hamiltonian Pauli term P_k we have $[P_k, S] = 0 \forall S \in \mathcal{S}$, it must be the case that $[CP_kC^\dagger, Z^{(q)}] = 0 \forall q \in \mathcal{I}$ as unitary rotations preserve

Fig. 4 Discrete geometrical symmetries are described by abelian subgroups of the molecular point-group, consisting of rotations, reflections and inversions. For example, benzene (C_6H_6) belongs to D_{6h}, which consists of group elements C_6 (60° rotations around the central axis perpendicular to the plane of the molecule), C_2/C_2' (180° rotations through axes parallel to the molecular plane), a reflection σ_h across the horizontal plane, two vertical reflections σ_v/σ_d, and finally the inversion symmetry i.

commutation relations. The implication of this is that the rotated Hamiltonian term $CP_kC^\dagger$ must consist of either I or Z in the qubit positions indexed by $\mathcal{I}$. These positions may subsequently be projected and dropped as in eqn (11), or in other words "tapered", from the Hamiltonian.[133]

When tapering, particular care must be taken to select the correct symmetry sector $\boldsymbol{\sigma}$, which can typically be assigned by some knowledge of the underlying problem. In electronic structure, $\mathbb{Z}_2$ symmetries can arise either from spin up/down parity or abelian subgroups of the geometrical point-group, describing discrete rotations, reflections or inversions of the molecule; an example is given in Fig. 4. We note that, while the full point-group might have a high dimensional irreducible representation, the limitation to abelian subgroups greatly restricts the degrees of freedom. One may refer to point-group tables to correctly choose the desired symmetry sector.[134,135]

4.3 Contextual subspace

In qubit tapering we may only remove as many qubits as there are Hamiltonian symmetries. The goal of the contextual subspace approach[136,137] is to relax this requirement to permit near-symmetry operators to define the subspace, chosen in such a way that we minimise the loss of information *via* the projection procedure. As such, this method will introduce some level of systematic error, but if the so-called "pseudo"-symmetries are selected carefully, this error may be controlled to allow high-accuracy simulations at a considerable saving of qubit resource.

Contextuality provides a broad conceptual picture of non-classical correlation;[138–141] the particular flavour that is exploited in the contextual subspace method is strong measurement contextuality.[142] In the setting of Pauli operators, this manifests as the presence of measurement contradictions in the same

Fig. 5 Molecular hydrogen, H_2 STO-3G, under the Jordan-Wigner transformation describes a noncontextual system with (a) 2-clique compatibility graph and (b) noncontextual energy spectrum, whose minimum coincides with the FCI energy.

vein as Peres-Mermin magic squares.[138,143] These contradictions arise from the violation of commutation transitivity among non-symmetry elements of a particular Pauli measurement set. Conversely, a noncontextual set is one whose compatibility graph consists of disjoint commuting cliques, all connected to a central symmetry set. An example of a noncontextual Hamiltonian is molecular hydrogen (H_2) in the minimal STO-3G basis set. Its Pauli terms consist of two disjoint cliques {*ZIII, IZII, IIZI, IIIZ*} and {*XXYY, XYYX, YXXY, YYXX*}, with the symmetry component {*IIII, IIZZ, IZIZ, IZZI, ZIIZ, ZIZI, ZZII*} as viewed graphically in Fig. 5a.

A noncontextual Hamiltonian has the structure

$$H_{NC} = \sum_{P \in \bar{\mathcal{S}}} \left(h_P + \sum_{k=1}^{K} h_{P,k} C_k \right) P. \tag{13}$$

where $\mathcal{S}$ is the symmetry generating set, $K \leq 2N + 1$ is the number of disjoint cliques, and C_k is a single representative from each clique, noting that $\{C_k, C_\ell\} = 0 \forall k \neq \ell$. For example, in Fig. 5a we have $K = 2$, $\mathcal{S} = \{IIZZ, IZIZ, ZIIZ\}$, $C_1 = ZIII$ and C_2

= $XXYY$. One obtains a classical objective function encoding the spectrum of H_{NC} *via* a phase-space description of the underlying hidden variable model:[144–146]

$$\eta(\mathbf{v}, \mathbf{r}) = \sum_{P \in \overline{S}} \left(h_P + \sum_{k=1}^{K} h_{P,k} r_k \right) \prod_{S \in \mathcal{S}_p} \nu_S. \quad (14)$$

where $\|\mathbf{r}\| = 1$ and $\mathcal{S}_P \subset \mathcal{S}$ satisfies $P = \prod_{S \in \mathcal{S}_P} S$. Optimising over the parameters $\mathbf{v}$, $\mathbf{r}$ yields the noncontextual ground state. In Fig. 5b we show the $\mathbf{r}$ optimisation landscape for H_2, noting that the minimum coincides exactly with the FCI energy.

Molecular hydrogen is a special case; in general, the electronic structure Hamiltonian is dominated by contextual interactions. In the contextual subspace approach, one projects symmetries $\mathcal{S}' \subset \mathcal{S}$ of a noncontextual model system over the full Hamiltonian, with the sector σ as in the introduction to Section 4 identified by optimising the noncontextual objective function of eqn (14), taking the elements ν_S that correspond with the chosen symmetries $S \in \mathcal{S}'$. The pairwise anticommuting clique representatives are moreover rotated onto a single Pauli operator prior to the stabiliser subspace projection, either *via* the sequence of rotations (R_{SeqRot}) or linear-combination of unitaries (R_{LCU}) construction.[147] Paired with the same Clifford rotations as in qubit tapering that map $\mathcal{S}'$ onto single-qubit Pauli operators, *i.e.* $CSC^\dagger = Z^{(f(S))}$ for each $S \in \mathcal{S}'$, we obtain the unitary $U = CR_{SeqRot/LCU}$ that one applies to the Hamiltonian prior to the single-qubit projections of eqn (11). We note that, due to the $R_{SeqRot/LCU}$ rotation, the unitary U is non-Clifford; however, in the *LCU* construction it is guaranteed that the increase in the number of Hamiltonian terms $L \in \mathbb{N}$ is L^K, and even this worst-case increase may only be encountered in highly contrived scenarios.[147]

The benefit of the contextual subspace method over those presented above, such as qubit tapering, is that we may project into an arbitrarily small qubit subspace, whereas tapering is limited by the number of $\mathbb{Z}_2$ symmetries present. This feature allows the contextual subspace approach to accommodate any given quantum device, regardless of size, and as such describes a highly flexible approach to quantum computing for many application domains, of which quantum chemistry has been investigated most extensively.[29,33,34,136,137,147,148] Of course, unlike tapering, we accumulate more error as we project into smaller subspaces, so this consideration needs to be navigated carefully. In Fig. 6 we demonstrate the decay of error as a contextual subspace is expanded from 1 to 23 qubits for sodium chloride (NaCl) STO-3G; while it begins as a 36-qubit problem in its entirety, application of the contextual subspace method allows us to achieve the target accuracy of 1.6 mHa for just 13 qubits, a considerable saving compared with the full system. These results were obtained using the Python packages PySCF[149] and Symmer.[150]

5 Proof-of-concept demonstration

In Fig. 1, we outlined a QM/MM simulation workflow for using quantum computational resources for a small region of a molecule, embedded within a larger framework comprising traditional quantum and classical computational chemistry methods. In this section we present results from an initial proof-of-concept example of this workflow where the proton transfer mechanism in water is considered.

Hydronium is the protonated cation of molecular water, *i.e.* $[H_3O]^+$. In solution, the simplest hydration structure of hydronium is the Zundel cation $[H_5O_2]^+$, where the hydronium ion is
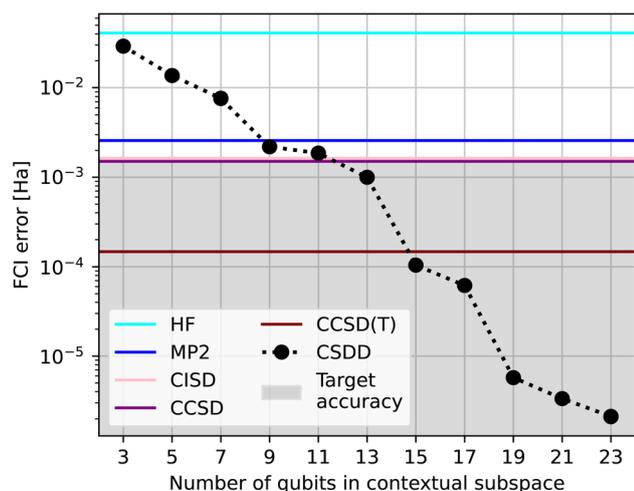

Fig. 6 NaCl STO-3G FCI errors at 3.5 Å (≈ 1.58 times the equilibrium bond length) against the number of qubits in an expanding contextual subspace. The reduced contextual subspace Hamiltonians were solved *via* direct diagonalisation (CSDD).
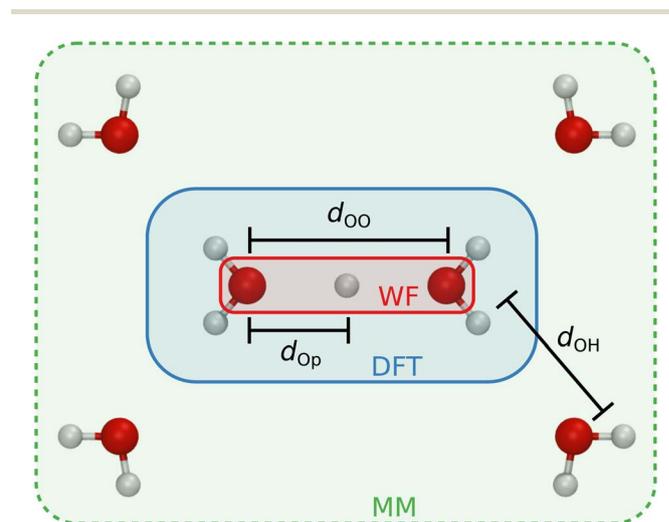

Fig. 7 The proton transfer geometry considered in the hardware proof-of-concept experiment. Two water molecules are separated by a distance d_{OO} Å, with a free proton placed in-between at a distance of d_{Op} Å from the left water. The proton ratio is defined by $r := d_{Op}/d_{OO} \in (0, 1)$. These atoms are treated at the QM level, where projection-based embedding is used to split the system into a density functional theory (DFT) and wavefunction (WF) subsystems. Further water atoms are placed around this system and treated at the molecular mechanics (MM) level, with the first solvation shell of four waters placed explicitly at d_{OH} Å along the OH bonds of the QM waters. The graphic was produced with VMD.[1]
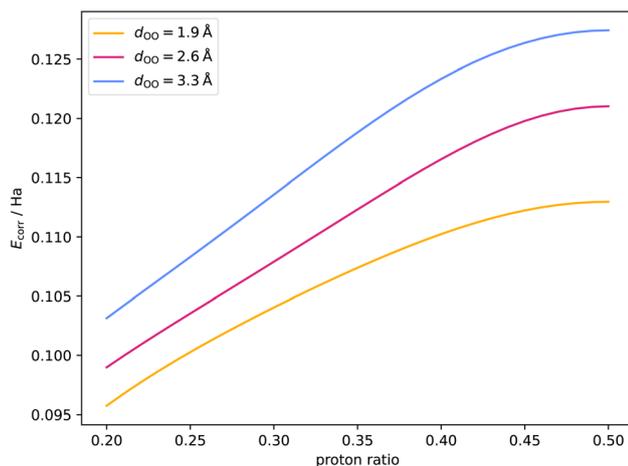
Fig. 8 The correlation energy $E_{corr} = E_{HF} - E_{FCI}$, as a function of proton ratio r. Restricted-HF and full-CI calculations are performed on a single Zundel cation (the atoms in blue region of Fig. 7) at three water separation distances with the STO-3G basis set. Across the three water separations, the mean correlation energy increase from $r = 0.2$ to 0.5 is 21%.

hydrogen-bonded to one other water molecule. For our testbed system we consider a planar Zundel cation, where the distance between the two waters is controlled by a parameter d_{OO}, and the free proton is placed somewhere along this length at a distance d_{Op} thereby defining a proton ratio $r := d_{Op}/d_{OO} \in (0, 1)$. Additional solvent water molecules are included around this Zundel system; initially four are placed at a distance d_{OH} from the four hydrogens of the Zundel cation, and more can be placed randomly in space beyond this. See Fig. 7 for a diagram of this geometry. As the free proton moves from being local to one water molecule ($r \approx 0.2$ or 0.8), to being shared between two ($r \approx 0.5$), the amount of correlation energy in the system increases by about 20% as can be seen in Fig. 8.

This system is then partitioned into subsystems which are treated at different levels of theory. The quantum region is defined as the Zundel cation, and the classical region comprises the remaining water molecules. The classical region is treated at the molecular mechanics level, whereas the quantum region is further partitioned into active and environment subregions *via* projection-based embedding (PBE), described in Section 3.1. The latter is treated with density functional theory, while the active region may be treated with conventional post-Hartree–Fock *ab initio* methods like coupled-cluster, or by constructing a qubit Hamiltonian and evaluating the energy with a quantum algorithm.

For the simulation we set up an electrostatically-coupled QM/MM simulation of the solvated Zundel system *via* the PySCF qmmm module,[149] where LAMMPS[151] drives the MM force evaluations and the MolSSI Driver Interface (MDI) package[152] facilitates communication between PySCF and LAMMPS. The PBE procedure is executed with Nbed[153] using the SPADE localisation method.[95,96] For the quantum chemistry calculations, the minimal STO-3G atomic basis set is employed, and Kohn–Sham DFT with the B3LYP exchange-correlation functional is used on

the environment region. We obtained a Pauli Hamiltonian with Symmer[150] *via* the Jordan-Wigner (JW) fermion encoding.[126]

The resultant qubit Hamiltonian of the core O–H–O atoms encodes the exact electronic energy embedded within the DFT environment (for the given basis set and exchange–correlation functional used), with additional polarisation effects caused by the MM atoms which are modelled as point charges (see Section 2.1 for further details). We refer to this quantity as the FCI-in-DFT-in-MM energy. It is useful to note that the Zundel cation has 15 molecular orbitals in the minimal STO-3G basis set, which would therefore require 30 qubits to describe the QM system on a quantum computer under the JW encoding. However, through the application of PBE (see Section 3.1), we may project out the four environment hydrogen atoms, each contributing a single occupied molecular orbital in this basis set and thus resulting in a reduction by 8 qubits in the resulting qubit Hamiltonian. The planar Zundel system has several planes of symmetry, allowing for an additional 4 qubits to be tapered out.[133] Finally, we use the frozen-core approximation which brings the final qubit count for our embedded system down to 16 qubits (see Section 4 for further detail on these methods).

For the QM/MM simulation, we place an additional 96 water molecules randomly around the Zundel system and perform a molecular dynamics (MD) propagation of these atoms, whilst the coordinates of the QM atoms with a specified (d_{OO}, r) parameter set remain fixed. Embedded qubit Hamiltonian data is recorded for 6 steps at a time interval of 2.0 fs per step. The system size and number of steps have been chosen to produce a minimally-sized simulation which can be used as a test-bed for the multiscale embedding and subspace methods which are the main focus of this perspective. After all steps have been performed, an averaged Hamiltonian for that parameter set is produced by calculating the mean Pauli coefficient for each Pauli string in the set of Hamiltonians. The set of mean Hamiltonians are then passed to the IQM superconducting device for the FCI-in-DFT-in-MM energy evaluations.

The energy of the quantum region is calculated on the 20-qubit IQM quantum chip QExa20 using quantum-selected configuration interaction (QSCI).[36–43] A set of electron configurations (determinants) $\mathcal{D} = \{|\Phi_k\rangle\}_{k=0}^{K-1}$ is sampled from the quantum device. We then form the configuration subspace projection operator $\mathbb{P} := \sum_{k=0}^{K-1} |\Phi_k\rangle\langle\Phi_k|$ and project the electronic structure Hamiltonian into this space

$$H \mapsto \mathbb{P}H\mathbb{P} = \sum_{k,\ell=0}^{K-1} \langle\Phi_k|H|\Phi_\ell\rangle|\Phi_k\rangle\langle\Phi_\ell|. \tag{15}$$

It is then possible to diagonalise the $K \times K$ matrix $\mathbf{H}$ with entries $\mathbf{H}_{k\ell} = \langle\Phi_k|H|\Phi_\ell\rangle$ using classical compute resources; by solving $\mathbf{H}v_j = \varepsilon_j v_j$ we obtain eigenstates $|\Psi_j\rangle = \sum_{k=0}^{K-1} v_{jk}|\Phi_k\rangle$ of the projected Hamiltonian that satisfy $\mathbb{P}H|\Psi_j\rangle = \varepsilon_j|\Psi_j\rangle$. For the hardware results in Fig. 9 we allowed a shot-budget of 10^5, which produced $K < 500$ unique configurations for each simulation, requiring only modest compute resources to diagonalise. Since the electronic energy is obtained *via* classical diagonalisation,
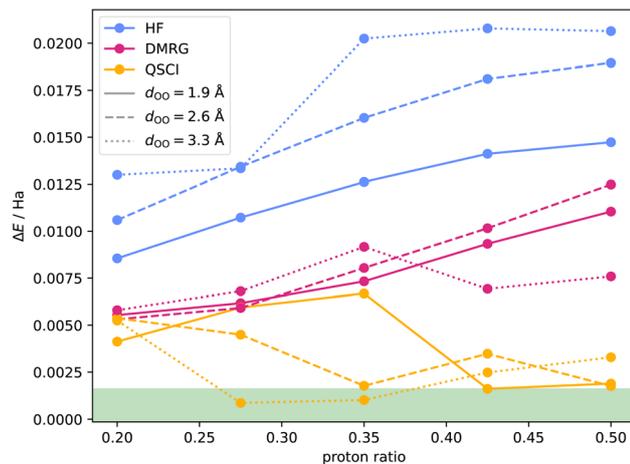
Fig. 9 Energy errors of the Zundel system embedded region at different O−O separations (indicated by the line style) and proton ratio values, evaluated with three different methods: HF and DMRG use purely classical compute resources, whereas the QSCI energies are found *via* sampling from the IQM device. All energies are displayed relative to the exact FCI-in-DFT-in-MM energy. The green region indicates the target accuracy of within 1.6 mHa to the exact energy. The proton ratio is relative to the O−O separation, for example a ratio of 0.5 indicates $d_{OH} = d_{OO}/2$, *i.e.* the proton is equidistant between two oxygens.

QSCI is more robust to hardware noise than algorithms such as VQE, where the energy estimation itself is susceptible to corruption by noise. By contrast, in QSCI it is only the quality of the configuration subspace that suffers.

The QSCI method requires preparing a quantum state on the quantum chip that approximates the ground state solution or, more generally, shares sufficient support with the ground state to provide an accurate approximation to its energy. The best way to prepare such a state remains an open question. Many previous works fix the parameters of a local unitary cluster Jastrow (LUCJ) ansatz using excitation amplitudes obtained from coupled-cluster calculations.[36–40] For this proof-of-concept demonstration, we instead use a direct matrix product state (MPS) to circuit mapping to warm-start the quantum chip with a low-bond-dimension DMRG solution. We note there is no evidence that either of these approaches will be sufficient for quantum advantage without additional work being done by the quantum device. This is beyond the scope of this article and will be addressed in later work; instead, it serves here as a validation of our QM/MM approach and a proof-of-concept that quantum computational resources may be deployed within large-scale chemical workflows.

The results of the quantum energy evaluations are presented in Fig. 9 alongside the DMRG energy and the HF energy (obtained by direct computation, $\langle \psi_{HF}|H|\psi_{HF}\rangle$). The DMRG energy corresponds to the DMRG solution with bond-dimension truncated to 2 that is prepared on the quantum chip. We do not claim that our results outperform DMRG in general; indeed, with a moderately high bond-dimension DMRG can be used to find the exact ground state solution for this system. However, the QSCI results obtained from this reference state improve upon its accuracy which motivates the use of this method in the

regime where exact DMRG becomes intractable. This accuracy improvement is attributed to the fact that the QSCI energy calculation is optimal within the subspace spanned by the sampled configurations whereas the DMRG energy is only optimal within the manifold of accessible MPS states with a capped maximum bond dimension. In this demonstration, it may also be the case that device noise enables us to sample from a larger set of relevant configurations than are present in the DMRG solution. While this effect of noise has been highlighted elsewhere as a potential benefit to QSCI,[44] we do not expect it to provide a reliable mechanism for larger systems and encourage the development of methods that expand the available configuration space in a more systematic way.[41–43]

The energies were evaluated at three different oxygen separation values, indicated by the line dashing, and at five different proton ratios. For this proof-of-concept demonstration, the Hamiltonian has been chosen such that it can be evaluated with exact diagonalisation, so we report the energies relative to this reference. We also mark the region known as chemical accuracy (within 1.6mHa to this reference), which is motivated by the limits of real-world experimental precision.[154] However, since we are using the minimal STO-3G basis set, it may be more suitable to describe this region as representing algorithmic accuracy, that is, accuracy within the limits of the basis set as opposed to true quantitative chemical accuracy.[155,156] STO-3G is almost always employed as the default basis in quantum chemistry studies on quantum computers, but as the quality of quantum hardware increases we expect that more effort will be made to move to larger basis sets where results can be considered more chemically meaningful.

As can be seen, the QSCI energies are comfortably within this region at two points in the $d_{OO} = 3.3$ Å evaluation, and near the boundary at several more evaluations. Whilst this precision to the exact solution is not seen across all of our energy evaluations, it is true that our implementation consistently yields more accurate energies than the HF and DMRG direct computation results, the latter of which was used to prepare the quantum circuit which is sampled from.

The proof-of-concept demonstration presented here combines classical molecular dynamics, quantum embedding and subspace methods into a single workflow, which allows quantum energy evaluations of chemical systems beyond the gas-phase regime.

In order to extend this workflow beyond the current fixed quantum geometry scheme, it is necessary to implement evaluations of the energy gradient with respect to nuclear coordinates of the QM nuclei. Within the PBE formalism, analytical nuclear gradients can be computed by introducing a total energy Lagrangian and finding its derivative when minimised with respect to the molecular orbitals.[157,158] Various quantum algorithms have been proposed for evaluating energy gradients on quantum hardware, both analytical[159–161] and numerical[162,163] with requirements ranging from NISQ-friendly to fault-tolerant. Energy gradients may also be approximated directly with a finite-difference approach, although this raises the quantum overhead substantially due to repeated use of the device. We anticipate that challenges surrounding the evaluation of energy

gradients can be tackled with a mixture of finite-difference approaches and algorithmic advances. With the ability to compute gradients, the quantum device can then be used to inform the updates of the quantum region nuclei, which inherently include the effect of the point-charge molecular mechanics atoms under the electrostatic QM/MM embedding. In such a setup, analysis of chemical properties beyond ground state energies is possible. For example, relaxation of the O–H–O atoms in the proof-of-concept example would allow for direct study of the proton transfer over a series of MD steps, allowing for accurate estimations of the energy barrier and consequent estimation of the hopping rate.

6 Outlook

The combination of methods described in this perspective have so far been placed in the context of near-term quantum algorithms and processors, applied to small molecular systems embedded within larger classical environments. The advantage of employing QM/MM as the overarching embedding procedure is that, as previously discussed, it is broadly applicable across many fields, including large-scale biomolecular simulations,[61,62] photochemistry,[63] surface chemistry,[64] and condensed matter physics.[65] As we look to the future development of quantum-classical hybrid multiscale models which can tackle systems like these, there are some important considerations for the scaling of such methodologies with quantum hardware improvement.

As quantum processors and error mitigation techniques develop, it will be possible to include larger subsystems within quantum algorithms, and to reduce the number of circuit measurements required to obtain accurate results. However, the hybrid methods we describe are still dependent on reasonable scaling of their components. Although these methods are flexible in how computation is partitioned between classical and quantum processors, the overall method may be constrained by either component. For example, in the case of PBE, a full-system DFT calculation is required to initialise the embedding. Similarly, determination of an initial reference state by DMRG constrains the QSCI procedure. As such, admissible systems are those for which approximate classical method can practically be performed.

On the other hand, circuit measurement and readout of quantum processors pose a challenge in scaling. It is important to scrutinise the computational overhead of quantum algorithms to understand their scalability. For example, a popular prediction for Quantum Phase Estimation (QPE) is it could take 25 hours to compute a single electronic point energy for the P450 cytochrome molecule, based on the availability of a fully error-corrected device comprised of 500 000 physical qubits each with a very low error rate of 0.001%.[164] Already, this is somewhat prohibitive for practical use-cases, and furthermore this is only for the calculation of electronic energies – to obtain other ground state properties of interest, such as dipole moments or forces, requires considerable additional overhead in the QPE framework. This is because QPE does not permit direct access to the wavefunction, so further processing must be performed on the QPU itself to extract these quantities. By contrast, in VQE, as long as we can write down an observable for the quantity of interest, it can be estimated with no to little overhead beyond the base electronic energy calculation (assuming there is considerable overlap between the Pauli measurements required for these observables and the Hamiltonian). In QSCI, we gain direct access to the wavefunction owing to the classical diagonalisation step, although this imposes a strict classical limit on the size of configuration subspace we can hope to solve, which is the bottleneck of the algorithm. However, this is a limitation of all CI-based methods; the objective of QSCI is instead to identify higher quality subspaces that lead to more compact representations of the wavefunction and thus allow us to extend the limits of CI applicability in terms of system scale.

Beyond intrinsic algorithmic considerations, the scaling of quantum algorithms is consistently challenged by hardware noise. Device fidelities and coherence times are constantly improving, enabling deeper and more complex quantum circuits to be reliably executed. However, in the pre-fault-tolerant regime, error mitigation and error suppression will always be essential in obtaining useful measurement results from the quantum device. For example, for QSCI error mitigation/suppression techniques are essential to maximise the yield of configurations within the correct symmetry sector. Some techniques such as dynamical decoupling[165,166] are widely useful and do not significantly increase the computational overhead as the system scales. On the other hand, for readout error mitigation, while being best practice for any quantum circuit execution, care must be taken in its implementation as naive methods exert an additional exponentially-scaling classical overhead to quantum computations.[167,168] Finally, algorithm-specific error mitigation schemes such as zero-noise extrapolation[169,170] and probabilistic error cancellation[170,171] each come with their own trade-off between accuracy improvements and computational overhead. As we look to scale implementations of hybrid methods, a well-chosen combination of error mitigation/suppression techniques should improve the reliability of quantum computations without accumulating a prohibitive classical overhead.

The integration of these components introduces additional challenges for hybrid workflows. If QPUs are not tightly integrated into classical HPC resources, communication between devices may well undo any practical benefit of hybridisation.[172]

Overcoming these scaling challenges will be crucial to the development of hybrid quantum-classical algorithms. This is so that as quantum hardware continues to improve, we can begin to study systems of sufficient complexity to challenge both the available quantum and classical compute resources and obtain scientifically important results beyond the reach of classical computers alone.

7 Conclusions

In this perspective, we have outlined a route towards integrating quantum computing into mainstream scientific computing, particularly within the chemical sciences, by demonstrating

how quantum devices can serve as supplementary resources that complement conventional approaches to molecular electronic structure calculations. The earliest demonstrations of quantum advantage for problems of real-world interest are likely to emerge through hybrid quantum–classical architectures, where quantum processing units (QPUs) are tightly coupled with high-performance computing (HPC) systems and play roles analogous to GPU accelerators. This integration is already well underway and holds promise for addressing computationally intensive tasks within multiscale and multi-physics simulation frameworks.

To this end, we highlighted a range of classical techniques that can significantly extend the applicability of current noisy intermediate-scale quantum (NISQ) devices (Sections 2, 3, 4). By identifying subdomains within larger simulations that are suitable for quantum treatment, facilitated by the QM/MM method demonstrated here for proton hopping (Section 5), we can offload the quantum-relevant components to QPUs. Furthermore, we presented a flexible and scalable simulation workflow that leverages many of these techniques in tandem, adapting to available quantum and classical resources. Such algorithmic and architectural advances are critical to realising quantum utility well before the advent of fully fault-tolerant quantum computers.

Author contributions

TMB: conceptualisation, investigation, software, writing – review & editing. AM: conceptualisation, investigation, software, writing – review & editing. TW: conceptualisation, investigation, software, writing – review & editing. MWdlB: conceptualisation, investigation, software, writing – review & editing. SW: writing – review & editing. MN: software. PS: software. AR: writing – review & editing. PJL: supervision. MC: supervision. MHV: supervision. LS: supervision. PVC: conceptualisation, supervision, writing – review & editing.

Conflicts of interest

The authors declare no conflicts of interest.

Data availability

The data and scripts supporting the proof-of-concept demonstration are openly available on GitHub (https://github.com/UCL-CCS/extending_digitaldiscovery_data) and have been archived on Zenodo (https://doi.org/10.5281/zenodo.17304716). The remaining results are generated with the open-source packages Nbed, Symmer, and PySCF, as cited throughout.

Acknowledgements

References

1 W. Humphrey, A. Dalke and K. Schulten, *J. Mol. Graphics*, 1996, **14**, 33–38.

2 A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik and J. L. O'Brien, *Nat. Commun.*, 2014, **5**, 1–7.

3 Y. Shen, X. Zhang, S. Zhang, J.-N. Zhang, M.-H. Yung and K. Kim, *Phys. Rev. A*, 2017, **95**, 020501.

4 P. J. J. O'Malley, *et al.*, *Phys. Rev. X*, 2016, **6**, 031007.

5 R. Santagati, J. Wang, A. A. Gentile, S. Paesani, N. Wiebe, J. R. McClean, S. Morley-Short, P. J. Shadbolt, D. Bonneau, J. W. Silverstone, D. P. Tew, X. Zhou, J. L. O'Brien and M. G. Thompson, *Sci. Adv.*, 2018, **4**, 1–12.

6 A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow and J. M. Gambetta, *Nature*, 2017, **549**, 242–246.

7 J. I. Colless, V. V. Ramasesh, D. Dahlen, M. S. Blok, M. E. Kimchi-Schwartz, J. R. McClean, J. Carter, W. A. de Jong and I. Siddiqi, *Phys. Rev. X*, 2018, **8**, 011021.

8 C. Hempel, C. Maier, J. Romero, J. McClean, T. Monz, H. Shen, P. Jurcevic, B. P. Lanyon, P. Love, R. Babbush, *et al.*, *Phys. Rev. X*, 2018, **8**, 031022.

9 A. Kandala, K. Temme, A. D. Córcoles, A. Mezzacapo, J. M. Chow and J. M. Gambetta, *Nature*, 2019, **567**, 491–495.

10 Y. Nam, J.-S. Chen, N. C. Pisenti, K. Wright, C. Delaney, D. Maslov, K. R. Brown, S. Allen, J. M. Amini, J. Apisdorf, K. M. Beck, A. Blinov, V. Chaplin, M. Chmielewski, C. Collins, S. Debnath, K. M. Hudek, A. M. Ducore, M. Keesan, S. M. Kreikemeier, J. Mizrahi, P. Solomon, M. Williams, J. D. Wong-Campos, D. Moehring, C. Monroe and J. Kim, *npj Quantum Inf.*, 2020, **6**, 33.

11 S. E. Smart and D. A. Mazziotti, *Phys. Rev. A*, 2019, **100**, 022517.

12 A. J. McCaskey, Z. P. Parks, J. Jakowski, S. V. Moore, T. D. Morris, T. S. Humble and R. C. Pooser, *npj Quantum Inf.*, 2019, **5**, 99.

13 J. E. Rice, T. P. Gujarati, M. Motta, T. Y. Takeshita, E. Lee, J. A. Latone and J. M. Garcia, *J. Chem. Phys.*, 2021, **154**, 134115.

14 F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, S. Boixo, M. Broughton, B. B. Buckley, D. A. Buell, *et al.*, *Science*, 2020, **369**, 1084–1089.

15 Q. Gao, G. O. Jones, M. Motta, M. Sugawara, H. C. Watanabe, T. Kobayashi, E. Watanabe, Y.-y. Ohnishi, H. Nakamura and N. Yamamoto, *npj Comput. Mater.*, 2021, 7, 70.

16 Y. Kawashima, E. Lloyd, M. P. Coons, Y. Nam, S. Matsuura, A. J. Garza, S. Johri, L. Huntington, V. Senicourt, A. O. Maksymov, J. H. V. Nguyen, J. Kim, N. Alidoust, A. Zaribafiyan and T. Yamazaki, *Commun. Phys.*, 2021, **4**, 245.

17 A. Eddins, M. Motta, T. P. Gujarati, S. Bravyi, A. Mezzacapo, C. Hadfield and S. Sheldon, *PRX Quantum*, 2022, **3**, 010309.

18 K. Yamamoto, D. Z. Manrique, I. T. Khan, H. Sawada and D. M. Ramo, *Phys. Rev. Res.*, 2022, **4**, 033110.

19 J. J. M. Kirsopp, C. Di Paola, D. Z. Manrique, M. Krompiec, G. Greene-Diniz, W. Guba, A. Meyder, D. Wolf, M. Strahm and D. Muñoz Ramo, *Int. J. Quantum Chem.*, 2022, **122**, 1–16.

20 K. Huang, X. Cai, H. Li, Z.-Y. Ge, R. Hou, H. Li, T. Liu, Y. Shi, C. Chen, D. Zheng, *et al.*, *J. Phys. Chem. Lett.*, 2022, **13**, 9114–9121.

21 P. Lolur, M. Skogh, W. Dobrautz, C. Warren, J. Biznárová, A. Osman, G. Tancredi, G. Wendin, J. Bylander and M. Rahm, *J. Chem. Theory Comput.*, 2023, **19**, 783–789.

22 V. Leyton-Ortega, S. Majumder and R. C. Pooser, *Quantum Sci. Technol.*, 2022, **8**, 014008.

23 Z. Liang, J. Cheng, H. Ren, H. Wang, F. Hua, Z. Song, Y. Ding, F. T. Chong, S. Han, X. Qian and Y. Shi, *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, 2024, **43**, 1834–1847.

24 M. Motta, G. O. Jones, J. E. Rice, T. P. Gujarati, R. Sakuma, I. Liepuoniute, J. M. Garcia and Y.-y. Ohnishi, *Chem. Sci.*, 2023, **14**, 2915–2927.

25 T. E. O'Brien, G. Anselmetti, F. Gkritsis, V. E. Elfving, S. Polla, W. J. Huggins, O. Oumarou, K. Kechedzhi, D. Abanin, R. Acharya, *et al.*, *Nat. Phys.*, 2023, **19**, 1787–1792.

26 I. T. Khan, M. Tudorovskaya, J. J. M. Kirsopp, D. Muñoz Ramo, P. Warrier, D. K. Papanastasiou and R. Singh, *J. Chem. Phys.*, 2023, **158**, 214114.

27 L. Zhao, J. Goings, K. Shin, W. Kyoung, J. I. Fuks, J.-K. Kevin Rhee, Y. M. Rhee, K. Wright, J. Nguyen, J. Kim, *et al.*, *npj Quantum Inf.*, 2023, **9**, 60.

28 S. Guo, J. Sun, H. Qian, M. Gong, Y. Zhang, F. Chen, Y. Ye, Y. Wu, S. Cao, K. Liu, *et al.*, *Nat. Phys.*, 2024, **20**, 1240–1246.

29 T. Weaving, A. Ralli, W. M. Kirby, P. J. Love, S. Succi and P. V. Coveney, *Phys. Rev. Res.*, 2023, **5**, 043054.

30 P. Liu, R. Wang, J.-N. Zhang, Y. Zhang, X. Cai, H. Xu, Z. Li, J. Han, X. Li, G. Xue, *et al.*, *Phys. Rev. X*, 2023, **13**, 021028.

31 E. Dimitrov, G. Sanchez-Sanz, J. Nelson, L. O'Riordan, M. Doyle, S. Courtney, V. Kannan, H. Naseri, A. G. Garcia, J. Tricker, *et al.*, *arXiv*, 2023, preprint, arXiv:2311.01242v1, DOI: 10.48550/arXiv.2311.01242.

32 M. A. Jones, H. J. Vallury and L. C. Hollenberg, *Phys. Rev. Appl.*, 2024, **21**, 064017.

33 Z. Liang, Z. Song, J. Cheng, H. Ren, T. Hao, R. Yang, Y. Shi and T. Li, *Proceedings of the 61st ACM/IEEE Design Automation Conference*, New York, NY, USA, 2024.

34 T. Weaving, A. Ralli, P. J. Love, S. Succi and P. V. Coveney, *npj Quantum Inf.*, 2025, **11**, 25.

35 J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth, *et al.*, *Phys. Rep.*, 2022, **986**, 1–128.

36 J. Robledo-Moreno, M. Motta, H. Haas, A. Javadi-Abhari, P. Jurcevic, W. Kirby, S. Martiel, K. Sharma, S. Sharma, T. Shirakawa, I. Sitdikov, R.-Y. Sun, K. J. Sung, M. Takita, M. C. Tran, S. Yunoki and A. Mezzacapo, *Sci. Adv.*, 2025, **11**, eadu9991.

37 D. Kaliakin, A. Shajan, F. Liang, J. Robledo Moreno, Z. Li, A. Mitra, M. Motta, C. Johnson, A. A. Saki, S. Das, I. Sitdikov, A. Mezzacapo and K. M. Merz Jr, *Commun. Phys.*, 2025, **8**, 396.

38 I. Liepuoniute, K. D. Doney, J. Robledo Moreno, J. A. Job, W. S. Friend and G. O. Jones, *J. Chem. Theory Comput.*, 2025, **21**, 5062–5070.

39 S. Barison, J. Robledo Moreno and M. Motta, *Quantum Sci. Technol.*, 2025, **10**, 025034.

40 A. Shajan, D. Kaliakin, A. Mitra, J. Robledo Moreno, Z. Li, M. Motta, C. Johnson, A. A. Saki, S. Das, I. Sitdikov, A. Mezzacapo and K. M. Merz, *J. Chem. Theory Comput.*, 2025, **21**, 6801–6810.

41 M. Mikkelsen and Y. O. Nakagawa, *Phys. Rev. Res.*, 2025, **7**, 043043.

42 K. Sugisaki, S. Kanno, T. Itoko, R. Sakuma and N. Yamamoto, *Phys. Chem. Chem. Phys.*, 2025, **27**, 20869–20884.

43 J. Yu, J. R. Moreno, J. Iosue, L. Bertels, D. Claudino, B. Fuller, P. Groszkowski, T. S. Humble, P. Jurcevic, W. Kirby, T. A. Maier, M. Motta, B. Pokharel, A. Seif, A. Shehata, K. J. Sung, M. C. Tran, V. Tripathi, A. Mezzacapo and K. Sharma, *arXiv*, 2023, preprint, arXiv:2501.09702v2, DOI: 10.48550/arXiv.2501.09702.

44 S. Piccinelli, A. Baiardi, M. Rossmannek, A. C. Vazquez, F. Tacchino, S. Mensa, E. Altamura, A. Alavi, M. Motta, J. Robledo-Moreno, W. Kirby, K. Sharma, A. Mezzacapo and I. Tavernelli, *arXiv*, 2025, preprint, arXiv:2508.02578v1, DOI: 10.48550/arXiv.2508.02578.

45 N. S. Blunt, J. Camps, O. Crawford, R. Izsák, S. Leontica, A. Mirani, A. E. Moylett, S. A. Scivier, C. Sünderhauf, P. Schopf, J. M. Taylor and N. Holzmann, *J. Chem. Theory Comput.*, 2022, **18**, 7001–7023.

46 H. Ma, J. Liu, H. Shang, Z. Fan, Z. Li and J. Yang, *Chem. Sci.*, 2023, **14**, 3190–3205.

47 M. Capone, M. Romanelli, D. Castaldo, G. Parolin, A. Bello, G. Gil and M. Vanzan, *ACS Phys. Chem. Au*, 2024, **4**, 202–225.

48 R. Santagati, A. Aspuru-Guzik, R. Babbush, M. Degroote, L. González, E. Kyoseva, N. Moll, M. Oppel, R. M. Parrish, N. C. Rubin, M. Streif, C. S. Tautermann, H. Weiss, N. Wiebe and C. Utschig-Utschig, *Nat. Phys.*, 2024, **20**, 549–557.

49 L. P. Weisburn, M. Cho, M. Bensberg, O. R. Meitei, M. Reiher and T. Van Voorhis, *J. Chem. Theory Comput.*, 2025, **21**, 4591–4603.

50 J. Günther, T. Weymuth, M. Bensberg, F. Witteveen, M. S. Teynor, F. E. Thomasen, V. Sora, W. Bro-Jørgensen, R. T. Husistein, M. Erakovic, M. Miller, L. Weisburn, M. Cho, M. Eckhoff, A. W. Harrow, A. Krogh,

T. V. Voorhis, K. Lindorff-Larsen, G. Solomon, M. Reiher and M. Christandl, *arXiv*, 2025, preprint, arXiv:2506.20587v1, DOI: **10.48550/arXiv.2506.20587**.

51 R. Izsák, C. Riplinger, N. S. Blunt, B. de Souza, N. Holzmann, O. Crawford, J. Camps, F. Neese and P. Schopf, *J. Comput. Chem.*, 2023, **44**, 406–421.

52 E. G. Hohenstein, O. Oumarou, R. Al-Saadon, G.-L. R. Anselmetti, M. Scheurer, C. Gogolin and R. M. Parrish, *J. Chem. Phys.*, 2023, **158**, 114119.

53 W. Li, Z. Yin, X. Li, D. Ma, S. Yi, Z. Zhang, C. Zou, K. Bu, M. Dai, J. Yue, *et al.*, *Sci. Rep.*, 2024, **14**, 16942.

54 P. Ettenhuber, M. B. Hansen, P. P. Poier, I. Shaik, S. E. Rasmussen, N. K. Madsen, M. Majland, F. Jensen, L. Olsen and N. T. Zinner, *J. Chem. Theory Comput.*, 2025, **21**, 3493–3503.

55 K. Wintersperger, H. Safi and W. Mauerer, *International Conference on Architecture of Computing Systems*, 2022, pp. 100–114.

56 T. Beck, A. Baroni, R. Bennink, G. Buchs, E. A. C. Pérez, M. Eisenbach, R. F. da Silva, M. G. Meena, K. Gottiparthi, P. Groszkowski, *et al.*, *Future Gener. Comput. Syst.*, 2024, **161**, 11–25.

57 A. Elsharkawy, X. Guo and M. Schulz, *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*, 2024, pp. 774–783.

58 Y. Alexeev, M. Amsler, M. A. Barroca, S. Bassini, T. Battelle, D. Camps, D. Casanova, Y. J. Choi, F. T. Chong, C. Chung, *et al.*, *Future Gener. Comput. Syst.*, 2024, **160**, 666–710.

59 S. Lee, J. Lee, H. Zhai, Y. Tong, A. M. Dalzell, A. Kumar, P. Helms, J. Gray, Z.-H. Cui, W. Liu, M. Kastoryano, R. Babbush, J. Preskill, D. R. Reichman, E. T. Campbell, E. F. Valeev, L. Lin and G. K.-L. Chan, *Nat. Commun.*, 2023, **14**, 1952.

60 P. B. Calio, C. Li and G. A. Voth, *J. Am. Chem. Soc.*, 2021, **143**, 18672–18683.

61 H. M. Senn and W. Thiel, *Angew. Chem., Int. Ed.*, 2009, **48**, 1198–1229.

62 C. E. Tzeliou, M. A. Mermigki and D. Tzeli, *Molecules*, 2022, **27**, 2660.

63 E. Boulanger and J. N. Harvey, *Curr. Opin. Struct. Biol.*, 2018, **49**, 72–76.

64 T. S. Hofer and A. O. Tirler, *J. Chem. Theory Comput.*, 2015, **11**, 5873–5887.

65 D. Hunt, V. M. Sanchez and D. A. Scherlis, *J. Phys.: Condens. Matter*, 2016, **28**, 335201.

66 A. Warshel and M. Levitt, *J. Mol. Biol.*, 1976, **103**, 227–249.

67 F. Maseras and K. Morokuma, *J. Comput. Chem.*, 1995, **16**, 1170–1179.

68 M. Svensson, S. Humbel, R. D. J. Froese, T. Matsubara, S. Sieber and K. Morokuma, *J. Phys. Chem.*, 1996, **100**, 19357–19363.

69 G. Groenhof, in *Introduction to QM/MM Simulations*, ed. L. Monticelli and E. Salonen, Humana Press, Totowa, NJ, 2013, pp. 43–66.

70 L. Cao and U. Ryde, *Front. Chem.*, 2018, **6**, 1–15.

71 E. R. Kjellgren, P. Reinholdt, A. Fitzpatrick, W. N. Talarico, P. W. K. Jensen, S. P. A. Sauer, S. Coriani, S. Knecht and J. Kongsted, *J. Chem. Phys.*, 2024, **160**, 124114.

72 The CUDA-Q development team, *CUDA-Q*, DOI: **10.5281/zenodo.15407754**.

73 J. Tomasi, B. Mennucci and R. Cammi, *Chem. Rev.*, 2005, **105**, 2999–3094.

74 M. Cossi, G. Scalmani, N. Rega and V. Barone, *J. Chem. Phys.*, 2002, **117**, 43–54.

75 A. Klamt and G. Schüürmann, *J. Chem. Soc., Perkin Trans. 2 (1972-1999)*, 1993, 799–805.

76 B. Mennucci, *Wiley Interdisc. Rev.: Comput. Mol. Sci.*, 2012, **2**, 386–404.

77 A. V. Marenich, C. J. Cramer and D. G. Truhlar, *J. Phys. Chem. B*, 2009, **113**, 6378–6396.

78 A. Klamt, *COSMO-RS: From Quantum Chemistry to Fluid Phase Thermodynamics and Drug Design*, Elsevier Science, 2005.

79 G. Scalmani and M. J. Frisch, *J. Chem. Phys.*, 2010, **132**, 114110.

80 L. O. Jones, M. A. Mosquera, G. C. Schatz and M. A. Ratner, *J. Am. Chem. Soc.*, 2020, **142**, 3281–3295.

81 J. M. Herbert, *J. Chem. Phys.*, 2019, **151**, 170901.

82 S. Laricchia, E. Fabiano and F. D. Sala, *J. Chem. Phys.*, 2012, **137**, 014102.

83 Z. Amanollahi, L. Lampe, M. Bensberg, J. Neugebauer and M. Feldt, *Phys. Chem. Chem. Phys.*, 2023, **25**, 4635–4648.

84 A. Ralli, M. Williams de la Bastida and P. V. Coveney, *Phys. Rev. A*, 2024, **109**, 022418.

85 A. Katabarwa, K. Gratsea, A. Caesura and P. D. Johnson, *PRX Quantum*, 2024, **5**, 020101.

86 F. R. Manby, M. Stella, J. D. Goodpaster and T. F. Miller, *J. Chem. Theory Comput.*, 2012, **8**, 2564–2568.

87 S. J. R. Lee, M. Welborn, F. R. Manby and T. F. I. Miller, *Acc. Chem. Res.*, 2019, **52**, 1359–1368.

88 J. M. Waldrop, A. Panyala, D. Mejia-Rodriguez, T. L. Windus and N. Govind, *J. Comput. Chem.*, 2025, **46**, e70043.

89 M. Rossmannek, F. Pavošević, A. Rubio and I. Tavernelli, *J. Phys. Chem. Lett.*, 2023, **14**, 3491–3497.

90 M. Bensberg and J. Neugebauer, *J. Chem. Theory Comput.*, 2020, **16**, 3607–3619.

91 M. Bensberg and J. Neugebauer, *J. Chem. Phys.*, 2019, **150**, 214106.

92 M. Bensberg and M. Reiher, *J. Phys. Chem. Lett.*, 2023, **14**, 2112–2118.

93 G. Knizia, *J. Chem. Theory Comput.*, 2013, **9**, 4834–4843.

94 J. Pipek and P. G. Mezey, *J. Chem. Phys.*, 1989, **90**, 4916–4926.

95 D. Claudino and N. J. Mayhall, *J. Chem. Theory Comput.*, 2019, **15**, 1053–1064.

96 D. C. Claudino, R. L. Smith and N. J. Mayhall, *Reference Module in Chemistry, Molecular Sciences and Chemical Engineering*, Elsevier, 2023, p. B978012821978200132X.

97 M. W. Schmidt, E. A. Hull and T. L. Windus, *J. Phys. Chem. A*, 2015, **119**, 10408–10427.

98 D. Claudino and N. J. Mayhall, *J. Chem. Theory Comput.*, 2019, **15**, 6085–6096.

99 E. Kolodzeiski and C. J. Stein, *J. Chem. Theory Comput.*, 2023, **19**, 6643–6655.

100 B. Hégely, P. R. Nagy, G. G. Ferenczy and M. Kállay, *J. Chem. Phys.*, 2016, **145**, 064107.

101 D. V. Chulhai and J. D. Goodpaster, *J. Chem. Theory Comput.*, 2018, **14**, 2023.

102 J. M. Waldrop, T. L. Windus and N. Govind, *J. Phys. Chem. A*, 2021, **125**, 6384–6393.

103 Q. Sun and G. K.-L. Chan, *Acc. Chem. Res.*, 2016, **49**, 2705–2712.

104 J. Inglesfield, *J. Phys. C: Solid State Phys.*, 1981, **14**, 3795.

105 G. Kotliar, S. Y. Savrasov, K. Haule, V. S. Oudovenko, O. Parcollet and C. Marianetti, *Rev. Mod. Phys.*, 2006, **78**, 865–951.

106 K. Haule and G. Kotliar, *Phys. Rev. B:Condens. Matter Mater. Phys.*, 2007, **76**, 104509.

107 K. Byczuk and D. Vollhardt, *Phys. Rev. B:Condens. Matter Mater. Phys.*, 2008, **77**, 235106.

108 K. Held, *Adv. Phys.*, 2007, **56**, 829–926.

109 S. Wouters, C. A. Jiménez-Hoyos, Q. Sun and G. K.-L. Chan, *J. Chem. Theory Comput.*, 2016, **12**, 2706–2719.

110 G. Knizia and G. K.-L. Chan, *Phys. Rev. Lett.*, 2012, **109**, 186404.

111 G. H. Booth and G. K.-L. Chan, *Phys. Rev. B*, 2015, **91**, 155107.

112 T. Tsuchimochi, M. Welborn and T. Van Voorhis, *J. Chem. Phys.*, 2015, **143**, 024107.

113 M. Welborn, T. Tsuchimochi and T. Van Voorhis, *J. Chem. Phys.*, 2016, **145**, 074102.

114 H.-Z. Ye, N. D. Ricke, H. K. Tran and T. Van Voorhis, *J. Chem. Theory Comput.*, 2019, **15**, 4497–4506.

115 H.-Z. Ye, H. K. Tran and T. Van Voorhis, *J. Chem. Theory Comput.*, 2020, **16**, 5035–5046.

116 O. R. Meitei and T. Van Voorhis, *J. Chem. Theory Comput.*, 2023, **19**, 3123–3130.

117 H. K. Tran, L. P. Weisburn, M. Cho, S. Weatherly, H.-Z. Ye and T. Van Voorhis, *J. Chem. Theory Comput.*, 2024, **20**, 10912–10921.

118 N. C. Rubin, *arXiv*, 2016, preprint, arXiv:1610.06910v2, DOI: **10.48550/arXiv.1610.06910**.

119 L. Mineh and A. Montanaro, *Phys. Rev. B*, 2022, **105**, 125117.

120 W. Li, Z. Huang, C. Cao, Y. Huang, Z. Shuai, X. Sun, J. Sun, X. Yuan and D. Lv, *Chem. Sci.*, 2022, **13**, 8953–8962.

121 C. Cao, J. Sun, X. Yuan, H.-S. Hu, H. Q. Pham and D. Lv, *npj Comput. Mater.*, 2023, **9**, 78.

122 J. Tilly, P. Sriluckshmy, A. Patel, E. Fontana, I. Rungger, E. Grant, R. Anderson, J. Tennyson and G. H. Booth, *Phys. Rev. Res.*, 2021, **3**, 033230.

123 Y. Liu, O. R. Meitei, Z. E. Chin, A. Dutt, M. Tao, I. L. Chuang and T. Van Voorhis, *J. Chem. Theory Comput.*, 2023, **19**, 2230–2247.

124 V. W.-z. Yu, J. Moussa and V. Blum, *J. Chem. Phys.*, 2021, **154**, 224107.

125 C. Sosa, J. Geertsen, G. W. Trucks, R. J. Bartlett and J. A. Franz, *Chem. Phys. Lett.*, 1989, **159**, 148–154.

126 P. Jordan and E. P. Wigner, *The Collected Works of Eugene Paul Wigner*, Springer, 1993, pp. 109–129.

127 S. B. Bravyi and A. Y. Kitaev, *Ann. Phys.*, 2002, **298**, 210–226.

128 E. Noether, *Nachr. Ges. Wiss. Gottingen*, 1918, **1918**, 235–257.

129 B. T. Gard, L. Zhu, G. S. Barron, N. J. Mayhall, S. E. Economou and E. Barnes, *npj Quantum Inf.*, 2020, **6**, 1–9.

130 X. Bonet-Monroig, R. Sagastizabal, M. Singh and T. O'Brien, *Phys. Rev. A*, 2018, **98**, 062339.

131 R. Sagastizabal, X. Bonet-Monroig, M. Singh, M. A. Rol, C. Bultink, X. Fu, C. Price, V. Ostroukh, N. Muthusubramanian, A. Bruno, *et al.*, *Phys. Rev. A*, 2019, **100**, 010302.

132 D. Gottesman, *Stabilizer Codes and Quantum Error Correction*, California Institute of Technology, 1997.

133 S. Bravyi, J. M. Gambetta, A. Mezzacapo and K. Temme, *arXiv*, 2017, preprint, arXiv:1701.08213v1, DOI: **10.48550/arXiv.1701.08213**.

134 K. Setia, R. Chen, J. E. Rice, A. Mezzacapo, M. Pistoia and J. D. Whitfield, *J. Chem. Theory Comput.*, 2020, **16**, 6091–6097.

135 D. Picozzi and J. Tennyson, *Quantum Sci. Technol.*, 2023, **8**, 035026.

136 W. M. Kirby, A. Tranter and P. J. Love, *Quantum*, 2021, **5**, 456.

137 T. Weaving, A. Ralli, W. M. Kirby, A. Tranter, P. J. Love and P. V. Coveney, *J. Chem. Theory Comput.*, 2023, **19**, 808–821.

138 N. D. Mermin, *Phys. Rev. Lett.*, 1990, **65**, 3373.

139 N. D. Mermin, *Rev. Mod. Phys.*, 1993, **65**, 803.

140 R. W. Spekkens, *Phys. Rev. A*, 2007, **75**, 032110.

141 R. W. Spekkens, *Phys. Rev. Lett.*, 2008, **101**, 020401.

142 W. M. Kirby and P. J. Love, *Phys. Rev. Lett.*, 2019, **123**, 200501.

143 A. Peres, *Phys. Lett. A*, 1990, **151**, 107–108.

144 R. W. Spekkens, in *Quasi-Quantization: Classical Statistical Theories with an Epistemic Restriction*, Springer Netherlands, Dordrecht, 2016, pp. 83–135.

145 W. M. Kirby and P. J. Love, *Phys. Rev. A*, 2020, **102**, 032418.

146 R. Raussendorf, J. Bermejo-Vega, E. Tyhurst, C. Okay and M. Zurel, *Phys. Rev. A*, 2020, **101**, 012350.

147 A. Ralli, T. Weaving, A. Tranter, W. M. Kirby, P. J. Love and P. V. Coveney, *Phys. Rev. Res.*, 2023, **5**, 013095.

148 M. Kiser, M. Beuerle and F. Šimkovic IV, *J. Chem. Theory Comput.*, 2025, **21**, 2256–2271.

149 Q. Sun, T. C. Berkelbach, N. S. Blunt, G. H. Booth, S. Guo, Z. Li, J. Liu, J. D. McClain, E. R. Sayfutyarova, S. Sharma, *et al.*, *Wiley Interdiscip. Rev.: Comput. Mol. Sci.*, 2018, **8**, e1340.

150 A. Ralli and T. Weaving, symmer, 2022, **https://github.com/UCL-CCS/symmer**.

151 A. P. Thompson, H. M. Aktulga, R. Berger, D. S. Bolintineanu, W. M. Brown, P. S. Crozier, P. J. in 't Veld, A. Kohlmeyer, S. G. Moore, T. D. Nguyen, R. Shan, M. J. Stevens, J. Tranchida, C. Trott and S. J. Plimpton, *Comput. Phys. Commun.*, 2022, **271**, 108171.

152 T. Barnes, A. Kohlmeyer, J. A. Nash and S. Mostafanejad, MDI Library, **https://github.com/MolSSI-MDI/MDI_Library**.

153 M. Williams de la Bastida and A. Ralli, Nbed, https://github.com/UCL-CCS/Nbed.

154 J. A. Pople, *Rev. Mod. Phys.*, 1999, **71**, 1267–1274.

155 M. Motta and J. E. Rice, *WIREs Comput. Mol. Sci.*, 2022, **12**, e1580.

156 A. Szabo and N. S. Ostlund, *Modern Quantum Chemistry*, McGraw-Hill Publishing Company, New York, First Edition, Revised edn, 1989.

157 S. J. R. Lee, F. Ding, F. R. Manby, I. Miller and F. Thomas, *J. Chem. Phys.*, 2019, **151**, 064112.

158 J. Csóka, B. Hégely, P. R. Nagy and M. Kállay, *J. Chem. Phys.*, 2024, **160**, 124113.

159 T. E. O'Brien, B. Senjean, R. Sagastizabal, X. Bonet-Monroig, A. Dutkiewicz, F. Buda, L. DiCarlo and L. Visscher, *npj Quantum Inf.*, 2019, **5**, 113.

160 T. E. O'Brien, M. Streif, N. C. Rubin, R. Santagati, Y. Su, W. J. Huggins, J. J. Goings, N. Moll, E. Kyoseva, M. Degroote, C. S. Tautermann, J. Lee, D. W. Berry, N. Wiebe and R. Babbush, *Phys. Rev. Res.*, 2022, **4**, 043210.

161 J. Lai, Y. Fan, Q. Fu, Z. Li and J. Yang, *J. Chem. Phys.*, 2023, **159**, 114113.

162 A. Delgado, J. M. Arrazola, S. Jahangiri, Z. Niu, J. Izaac, C. Roberts and N. Killoran, *Phys. Rev. A*, 2021, **104**, 052402.

163 K. Sugisaki, H. Wakimoto, K. Toyota, K. Sato, D. Shiomi and T. Takui, *J. Phys. Chem. Lett.*, 2022, **13**, 11105–11111.

164 J. J. Goings, A. White, J. Lee, C. S. Tautermann, M. Degroote, C. Gidney, T. Shiozaki, R. Babbush and N. C. Rubin, *Proc. Natl. Acad. Sci. U. S. A.*, 2022, **119**, e2203533119.

165 L. Viola and S. Lloyd, *Phys. Rev. A*, 1998, **58**, 2733.

166 L. Viola, E. Knill and S. Lloyd, *Phys. Rev. Lett.*, 1999, **82**, 2417.

167 F. B. Maciejewski, Z. Zimborás and M. Oszmaniec, *Quantum*, 2020, **4**, 257.

168 S. Bravyi, S. Sheldon, A. Kandala, D. C. Mckay and J. M. Gambetta, *Phys. Rev. A*, 2021, **103**, 042605.

169 Y. Li and S. C. Benjamin, *Phys. Rev. X*, 2017, 7, 021050.

170 K. Temme, S. Bravyi and J. M. Gambetta, *Phys. Rev. Lett.*, 2017, **119**, 180509.

171 S. Endo, S. C. Benjamin and Y. Li, *Phys. Rev. X*, 2018, **8**, 031027.

172 M. Mohseni, A. Scherer, K. G. Johnson, O. Wertheim, M. Otten, N. A. Aadit, Y. Alexeev, K. M. Bresniker, K. Y. Camsari, B. Chapman, S. Chatterjee, G. A. Dagnew, A. Esposito, F. Fahim, M. Fiorentino, A. Gajjar, A. Khalid, X. Kong, B. Kulchytskyy, E. Kyoseva, R. Li, P. A. Lott, I. L. Markov, R. F. McDermott, G. Pedretti, P. Rao, E. Rieffel, A. Silva, J. Sorebo, P. Spentzouris, Z. Steiner, B. Torosov, D. Venturelli, R. J. Visser, Z. Webb, X. Zhan, Y. Cohen, P. Ronagh, A. Ho, R. G. Beausoleil and J. M. Martinis, *arXiv*, 2025, preprint, arXiv:2411.10406v2, DOI: 10.48550/arXiv.2411.10406.