Towards quantum chemistry on a quantum computer

B. P. Lanyon[1,2]*, J. D. Whitfield[4], G. G. Gillett[1,2], M. E. Goggin[1,5], M. P. Almeida[1,2], I. Kassal[4], J. D. Biamonte[4†], M. Mohseni[4†], B. J. Powell[1,3], M. Barbieri[1,2†], A. Aspuru-Guzik[4]* and A. G. White[1,2]

Exact first-principles calculations of molecular properties are currently intractable because their computational cost grows exponentially with both the number of atoms and basis set size. A solution is to move to a radically different model of computing by building a quantum computer, which is a device that uses quantum systems themselves to store and process data. Here we report the application of the latest photonic quantum computer technology to calculate properties of the smallest molecular system: the hydrogen molecule in a minimal basis. We calculate the complete energy spectrum to 20 bits of precision and discuss how the technique can be expanded to solve large-scale chemical problems that lie beyond the reach of modern supercomputers. These results represent an early practical step toward a powerful tool with a broad range of quantum-chemical applications.

One of the most important and difficult problems in science is the accurate calculation of material properties from first principles. Indeed, this can be viewed as the central challenge for theorists from an enormous range of fields, including chemistry, condensed matter physics, materials science, biophysics and biochemistry. The challenge of first-principles calculations arises because the computational resources required to obtain exact (that is, full configuration interaction) solutions of the Schrödinger equation on a conventional computer generally increase exponentially with the number of atoms involved[1,2]. This renders such calculations intractable for all but the smallest systems, even using the latest supercomputers.

In chemistry, the most important task is to calculate the energy of molecules to within chemical accuracy ($k_b T \sim 1$ kJ mol^{-1}), which is required to predict reaction rates. The difficulty of solving this problem exactly has led to the development of several approximate methods for first-principles quantum chemistry[3]. Tremendous methodological progress and the growth of the power of classical computers have allowed for spectacular computational successes, such as density functional theory (DFT) calculations of entire proteins or nanostructures[4,5]. Nevertheless, there are numerous important instances where approximate methods fail. For example, DFT is unreliable for systems involving excited-state charge transfer[6], conical intersections[7] (which are important in green fluorescent protein and rhodopsin) or strong electronic correlations (such as those in high-temperature superconductors[8]). Although some of these effects are captured by post-Hartree-Fock wavefunction methods, there are weaknesses associated with these methods as well. For example, coupled cluster methods fail for bond dissociation[9].

In 1982, Richard Feynman[1] suggested that an avenue for the tractable simulation of quantum systems would be to build quantum mechanical simulators. The most versatile quantum simulator is a universal quantum computer—a device that uses quantum systems themselves to store and process data. It has been proposed that a quantum computer can simulate many-body physical quantum systems (such as molecules) and calculate their energies to a fixed accuracy with a number of quantum computational

resources that increases only polynomially with the number of particles[2,10–12]. Such a device would provide an extremely powerful tool for new science and technology because essentially exact (within a specified basis) molecular properties would be available[13].

Here, we demonstrate a proposed quantum algorithm for quantum chemical calculations in a linear optic quantum computing architecture for the smallest problem: obtaining the energies of the hydrogen molecule (H_2) in a minimal basis. We perform a key algorithmic step—the iterative phase estimation algorithm[14–17]—in full, achieving a high level of precision and robustness to error. We implement other algorithmic steps with assistance from a classical computer and explain how this non-scalable approach could be avoided, given more quantum computational resources than are currently available. Finally, we provide new theoretical results, which lay the foundations for the next generation of simulation experiments using quantum computers. Pioneering experiments in the context of quantum simulation with quantum systems were first performed using NMR-based systems to simulate quantum oscillators[18], leading up to recent simulations of a pairing Hamiltonian[16,19]. The phase transitions of a two-spin quantum magnet[20] and the Dirac equation[21] have been simulated using ion-trap systems. Here we take this work further by both simulating a many-body quantum system and calculating its energy.

The fundamental concept of quantum computation is the quantum bit, or qubit—a two level quantum system that is analogous to the classical bit. Quantum computing requires the ability to prepare a large number of qubits into some initial state, manipulate the state in an arbitrary way (for example, using quantum logic gates) and read out the final state, much like the requirements of classical computing. A key difference between classical and quantum computing is that whereas the state of a bit can be either 0 or 1 at any given time, a qubit can exist in any arbitrary superposition state $\alpha|0\rangle + \beta|1\rangle$ (where the Dirac notation has been used and α and β are complex numbers such that $|\alpha|^2 + |\beta|^2 = 1$). The ability of qubits to exist in superposition states is a precursor to a uniquely quantum mechanical phenomenon called entanglement: multiple qubits can exist in a superposition of their joint states that cannot be written as a

product of states of the individual systems—the individuals have no well-defined state on their own but their joint state can be perfectly defined, hence they are 'entangled'. It has been shown that, for the conventional model of quantum computing, the ability for a quantum computer to exploit these correlations is a necessary condition for any computational advantage[22].

Key challenges in building a quantum computer are to isolate qubits in physical systems, prepare arbitrary states, implement arbitrary qubit evolution, read out qubit states, overcome noise and to do all this on a large scale; that is, with a large number of qubits. Today there is a large ongoing international effort to build a quantum computer in a wide range of different physical systems. Significant progress has been made with ionic[23,24], photonic[25], superconducting[26] and solid state[27] systems, for example. Alongside ionic methods, photonic approaches represent one of the most mature experimental quantum computing architectures. Not only does this approach offer a legitimate path to building a large-scale quantum computer[28,29], but it also provides a test bed in which to explore the next steps in experimental quantum computing. Today, the available photonic quantum computer technology is reaching the stage where the first small-scale quantum algorithms can be explored[30,31].

Results

Molecular energies are represented as the eigenvalues of an associated time-independent Hamiltonian $\hat{H}$ and can be efficiently obtained to fixed accuracy using a quantum algorithm with three distinct steps[15]: encoding a molecular wavefunction into qubits; simulating its time evolution using quantum logic gates; and extracting the approximate energy using the phase estimation algorithm[15,32]. The latter is a general-purpose quantum algorithm for evaluating the eigenvalues of arbitrary Hermitian or unitary operators. The algorithm estimates the phase, ϕ, accumulated by a molecular eigenstate, $|\psi\rangle$, under the action of the time-evolution operator, $\hat{U} = e^{-(i\hat{H}t/\hbar)}$; that is,

$$e^{-(i\hat{H}t/\hbar)}|\psi\rangle = e^{-(iEt/\hbar)}|\psi\rangle = e^{-i2\pi\phi}|\psi\rangle \qquad (1)$$

where E is the energy eigenvalue of $|\psi\rangle$. Therefore, estimating the phase for each eigenstate amounts to estimating the eigenvalues of the Hamiltonian.

We implement the iterative phase estimation algorithm[15,33] (IPEA), which minimizes the number of qubits and quantum logic gates required. Figure 1a shows the IPEA at iteration k. Key algorithmic features are the encoding of an eigenstate wave function into an n–qubit register and the implementation of powers of the time evolution operator $\hat{U}$ on this register conditional on the state of the upper single qubit, as in Fig. 1a. By conditional we mean that if the top qubit is in the logical 1 (0) state then this evolution does (does not) occur; that is, $|1, \psi\rangle \rightarrow \hat{U}^p|1, \psi\rangle$ and $|0, \psi\rangle \rightarrow |0, \psi\rangle$, where $p = 2^{k-1}$.

At the end of the kth iteration, a measurement of the logical state of the top qubit is performed, yielding the result 0 or 1. This represents an estimate of the kth bit of the binary expansion[34] of ϕ. As each iteration provides one bit, the total number of iterations (m) determines the precision with which the phase is estimated. An important feature is that least significant bits are estimated first (that is, k is iterated backwards from m to 1) and measured bits are used to improve the estimation of more significant bits. This information transfer between iterations is achieved via a single qubit rotation $\hat{R}_z(\omega_k)$ on the upper qubit, whose angle is determined by all previously measured bits, as described in the caption of Fig. 1a. A detailed explanation of the IPEA can be found in ref. 33, but it is useful to consider the following: the IPEA allows the phase ϕ to be determined by letting the wave function evolve in time and then (iteratively) performing a Fourier

transform to obtain the frequency of the phase oscillation. Consequently, to obtain higher precision in the frequency domain, the wave function must be propagated longer in the time domain; that is, higher powers of $\hat{U}$ enable more precision in the phase estimation.

In the case where the phase to be calculated has an exact expansion to m bits, a perfect implementation of m iterations of the algorithm (that is, without experimental error) will deterministically extract the exact phase. This is not generally the case and this leads to a non-zero probability of intrinsic algorithmic error: let m bits of a binary expansion of ϕ be $\tilde{\phi} = 0.\phi_1\phi_2\ldots\phi_m$, such that $\phi = \tilde{\phi} + \delta 2^{-m}$ where δ is a remainder $0 \leq \delta < 1$. A perfect implementation of the algorithm achieves a precision of $\pm 2^{-m}$ with error probability[33] $\varepsilon \leq 1 - (8/\pi^2) \approx 0.19$, which is independent of m (the bound is saturated for $\delta = 0.5$).

In the case of perfect algorithm implementation, this error can always be eliminated by simply repeating each IPEA iteration multiple (n) times, yielding n possible values for the corresponding bit. A majority vote is then performed by selecting the mode of the set of these samples—that is, choosing the result that is observed most often (see Methods). This technique only requires repeating the algorithm a fixed number of times, which does not affect its overall efficiency.

We take the standard approach to quantum-chemical calculations by solving an approximate Hamiltonian created by using the Born-Oppenheimer approximation (where the electronic Hamiltonian is parameterized by the nuclear configuration) and choosing a suitable truncated basis set in which to describe the non-relativistic electronic system. Typical sets consist of a finite number of single-electron atomic orbitals, which are combined to form antisymmetric multi-electron molecular states (configurations)[35]. Calculating the eigenvalues of the electronic Hamiltonian using all configurations gives the exact energy in the basis set and is referred to as full configuration interaction (FCI). For N orbitals and m electrons there are

$$\binom{N}{m} \approx \frac{N^m}{m!}$$

ways to allocate the electrons among the orbitals. This exponential growth is the handicap of FCI calculations on classical computers.

As described in the Methods section, the Hamiltonian for the H_2 molecule in a minimal basis is block diagonalized in the symmetry adapted basis, with 2×2 submatrices ($\hat{H}^{(1,6)}$ and $\hat{H}^{(3,4)}$). We map the configurations spanning each subspace to the qubit computational basis. As the subspaces are two-dimensional, one qubit suffices to represent the wavefunction. The corresponding time-evolution operators $\hat{U}^{(p,q)} = e^{-i(\hat{H}^{(p,q)}t/\hbar)}$, where $(p,q) = (1,6)$ or $(3,4)$, are therefore one-qubit operators. Finding the eigenvalues of each separately, using a phase estimation algorithm, amounts to performing FCI. For the purpose of our demonstration we encode exact eigenstates, obtained via a preliminary calculation on a classical computer. In the Supplementary Information we show that the algorithm is in fact robust to imperfect eigenstate encoding.

Overview of experiment

We implement the IPEA in a photonic architecture, encoding qubits into the polarization of single photons generated in pairs through the process of spontaneous parametric downconversion. An experimental schematic is shown in Fig. 1c. In our encoding, the logical state $|0\rangle$ ($|1\rangle$) corresponds to the horizontal, $|H\rangle$ (vertical, $|V\rangle$) polarization. Single qubit gates can be easily implemented using birefringant wave plates to perform arbitrary polarization rotations. Multi-qubit gates are possible by combining linear optic elements (phase shifters and beam splitters) with projective measurement, as originally proposed in 2001 and demonstrated in 2003[25,28]. Such
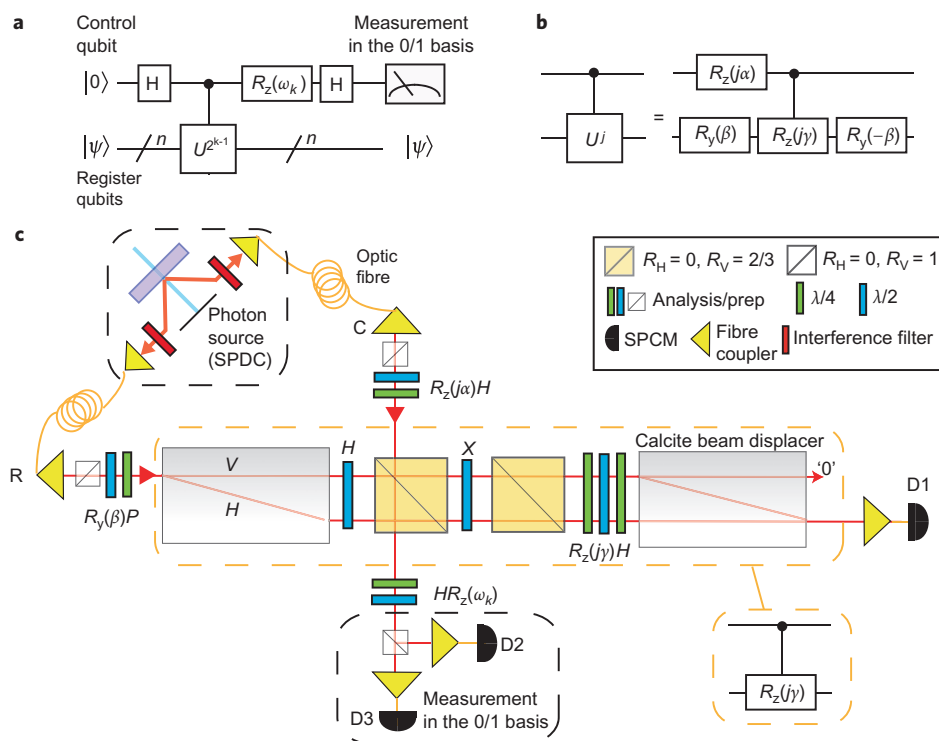
Figure 1 | The quantum algorithm for calculating energies of many-body quantum systems and our experimental implementation with linear optics. a, The iterative phase estimation algorithm[15,33] at iteration k represented using quantum circuit notation[34]. To produce an m-bit approximation to the phase of the eigenstate, ϕ (see Eqn 1), the algorithm is iterated m times. Each iteration obtains one bit of $\phi(\phi_k)$; starting from the least significant (ϕ_m), k is iterated backwards from m to 1. The angle ω_k depends on all previously measured bits as $\omega_k = -2\pi b$, where b, in the binary expansion of ϕ, is $b = 0.0\phi_{k+1}\phi_{k+2} \ldots \phi_m$ and $\omega_m = 0$. H is the standard Hadamard gate[34]. **b,** Quantum circuit model representation of our gate network for a two-qubit-controlled $\hat{U}^j$ gate, as discussed in the Methods section. **c,** Two-qubit optical implementation of the IPEA. Photon pairs are generated by spontaneous parametric down-conversion (SPDC), coupled into a single-mode optical fibre and launched into free space optical modes C (control) and R (register). The operation of the optical-controlled $\hat{R}_z$ gate is described by Lanyon[31] et al. Coincident detection events (3.1 ns window) between single photon counting modules (SPCMs) D1 and D3 (D2 and D3) herald a successful run of the circuit and result in 0 (1) for ϕ_k. Waveplates are labelled with their corresponding operations. PPBS: partially polarizing beam splitter. X is the standard Pauli X gate ($0 \leftrightarrow 1$) and $R_{\hat{n}}$ is the rotation of the qubit about the $\hat{n}$ axis of the Bloch sphere[34]. $\lambda/2$ ($\lambda/4$) is a half (quarter) wave plate.

gates are high quality, well characterized and have several in-principle paths to scalable optical quantum computing[29]. The key device in our experiment is a photonic two-qubit controlled-unitary quantum logic gate, recently developed in our laboratory[31]. This device exploits multiple photonic degrees of freedom to extend the range of possible operations. We note that our implementation of the IPEA is the first, in any physical system, to use quantum logic gates that can generate the non-classical correlations required for efficient quantum computing. We acknowledge a previous implementation in a liquid-state ensemble NMR architecture[36], which is arguably an in-principle non-scalable approach[37]. Another implementation in ions[38] was semi-classical in the sense that it combined single-qubit measurement and feedback in place of entangling gates. We now provide a simple conceptual explanation of our experimental method.

Each iteration of our IPEA implementation requires two qubits, and therefore two photons, to be injected into the optical circuit shown in Fig. 1c; one in each optical input mode. We call these photons the 'control' and 'register', after the qubits that they encode (Fig. 1a). Each optical mode then passes through a polarizing beam splitter set to prepare (transmit) the fiducial initial state $|H\rangle$. The required IPEA input states are then prepared using wave plates: the register polarization is rotated into a state that represents (encodes) an eigenstate of our molecular Hamiltonian ($\hat{H}^{(i,j)}$) and the control photon is rotated into the equal superposition state $(|H\rangle + |V\rangle)/\sqrt{2}$. Both modes then pass through an optical network that implements a controlled-$\hat{U}^{2^{k-1}}$ gate: if the control is horizontally polarized then nothing is changed (the mathematical

'identity' operation is implemented); if the control is vertically polarized then the operation $\hat{U}^{2^{k-1}}$ is implemented on the register polarization—that is, the evolution of the register polarization state is equivalent to (simulates) the evolution of a molecular eigenstate due to the molecular Hamiltonian, for a specific time interval. In this case, the joint polarization state will pick up a phase determined by k and the energy of the encoded eigenstate. Note that this nonlinear process (that is, conditional dynamics at the single photon level) occurs only if a single photon exits each port, hence the circuit is non-deterministic (we note that this can be overcome and efficient scalable linear optical quantum computing is possible[28,39]). Finally, the control polarization is measured in the H/V basis, yielding the result H (0) or V (1). This result is the estimate for one bit of the binary expansion of the phase. This same circuit (one iteration of the IPEA) is repeated with many photon pairs and a majority vote is used to determine the estimated bit value.

The circuit is changed in two ways for the next iteration: the molecule is simulated for a different time period (that is, k is changed) and previous results are fed back (the $\hat{R}_z(\omega_k)$ gate is set to an angle determined by all previously measured bits, see Fig. 1). Both changes are implemented by changing the settings of wave plates.

To resolve energy differences relevant to chemical processes, absolute molecular energies must be computed to an accuracy greater than $k_b T \approx 1$ kJ mol^{-1}. Furthermore, because we choose to encode the molecular Hamiltonian in atomic units, the phase ϕ is read out in hartrees, 1 E_h = 2,625.5 kJ mol^{-1}. Therefore, to achieve

a precision of 1 kJ mol^{-1}, the IPEA must be able to resolve the phase to 14 bits, as $(2\pi \times 2^{-14})E_h$ corresponds to 1.01 kJ mol^{-1}. To demonstrate that the necessary precision can be achieved, we carried out the IPEA to 20 bits, using $n = 31$ samples per bit. Figure 2 shows our results: H_2 energies calculated over a range of internuclear separations, reconstructing the potential energy surfaces. The ground-state energy obtained at the equilibrium bond length of 73.48 pm is -535.58 ± 0.03 kJ mol^{-1}, which agrees exactly with the result obtained on a classical computer to an uncertainty in the least significant bit.

Discussion

Achieving such high precision will become a far more significant challenge for large-scale implementations. Owing to the small scale of our demonstration, we are able to implement each power of $\hat{U}^{(p,q)}$ directly, by re-encoding the same number of gates (see Methods). Therefore, the probability of error introduced by gate imperfections remains a constant for each bit. This is the main algorithmic feature that allows for the high precision obtained. However, following current proposals, the circuit network for $\hat{U}$ will not generally have the same form as $\hat{U}^J$ for larger implementations (detailed in our Supplementary Information). For each additional digit of precision sought, the gate requirements of the algorithm are roughly doubled, thereby amplifying any gate error.

Important next experimental steps are to demonstrate the two parts of the quantum algorithm that we implemented with assistance from a classical computer. First, encoding even low fidelity eigenstate approximations into qubits is a highly non-trivial step for molecules much larger than H_2. In many cases this problem could be overcome using a heuristic adiabatic state preparation algorithm[15,20,40]. Here, ground state approximations, for example, can be efficiently obtained as long as the energy gap between the ground state and excited states is sufficiently large along the path of the adiabatic evolution[41]. Second, directly calculating and decomposing the molecular evolution operator into quantum logic gates does not scale efficiently with molecular size[2] and an alternative scheme must be used. The proposed solution exploits the fact that the general molecular Hamiltonian is a sum of fixed-sized one- and two-electron terms that can be efficiently simulated individually and combined to approximate the global evolution[2]. We give an overview of this 'operator-splitting' technique in our Supplementary Information and find that the total number of elementary quantum gates required to simulate the evolution of an arbitrary molecule, without error correction, scales as N^5, where N is the number of single-particle basis functions used to describe the molecular system. In this scheme, N is also the number of qubits necessary. We also present the quantum logic circuits required to simulate each term in the general molecular Hamiltonian—these are the building blocks of a universal quantum molecular simulator. Finally, we perform an accurate resource count to reproduce our H_2 simulation in this fully scalable way: four qubits and 522 perfect gates are required to simulate the full unitary propagator such that the error of the simulated evolution is within chemical accuracy. We note that experiments have been performed that involve logic gate operations on four qubits[30]. However, implementing 522 sequential one- and two-qubit quantum logic gates is currently beyond the capabilities of any quantum computing architecture, and the most immediate reasons are different in each case. In linear optics a large number of additional photons would be required, which we are not yet able to generate in a controlled fashion.

Other major challenges in the path to scalability include those associated with scaling up the 'hardware'—that is, achieving more qubits, gates and longer coherence times. Much progress is being made on developing the necessary technology for a large-scale
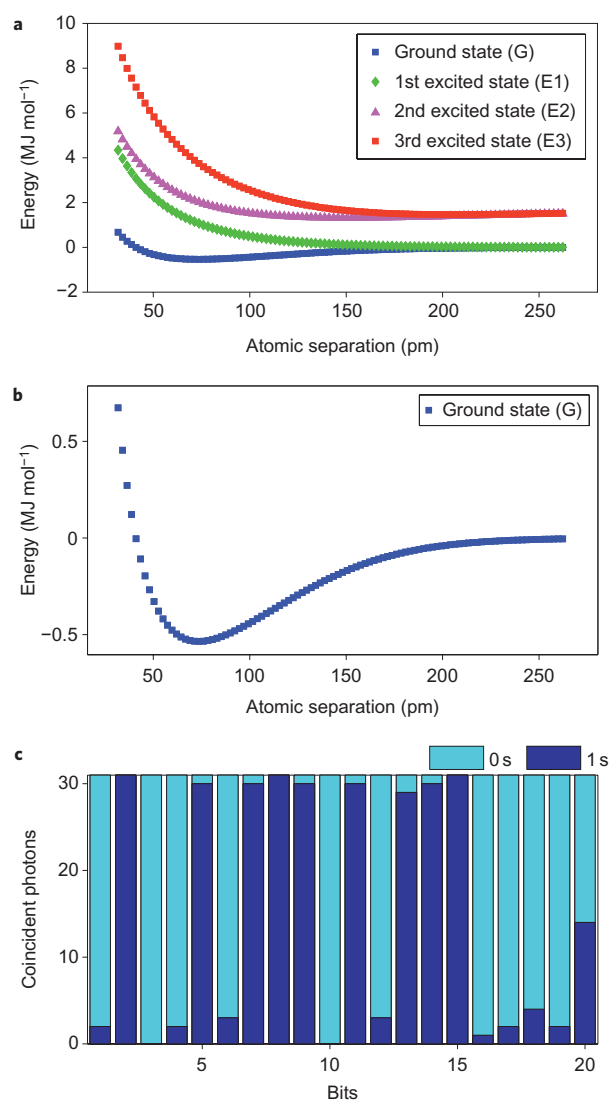
Figure 2 | Experimental quantum algorithm results. a, H_2 potential energy curves in a minimal basis. Each point was calculated with a 20-bit iterative phase estimation algorithm, using $n = 31$ samples per bit (repetitions of each iteration). Measured phases are converted to energies via $E = 2\pi\phi E_h + (e^2/4\pi\varepsilon R) - E_\infty$; the second term accounts for the proton–proton Coulomb energy at atomic separation R and E_∞ is the ground state energy of two hydrogen atoms at infinite separation. Each case achieved the target precision of 20 bits corresponding to $\pm(2^{-20} \times 2\pi)E_h \approx 16$ J mol^{-1}. Curve G (E3) is the low (high) eigenvalue of $\hat{H}^{(1,6)}$. Curve E1 is a triply degenerate spin-triplet state, corresponding to the lower eigenvalue of and $\hat{H}^{(3,4)}$ as well as the eigenvalues $\hat{H}^{(2)}$ and $\hat{H}^{(5)}$. Curve E2 is the higher (singlet) eigenvalue of $\hat{H}^{(3,4)}$. **b**, Curve G rescaled to highlight the bound state. **c**, Example of raw data for the ground state energy, obtained at the equilibrium bond length 73.48 pm. The measured binary phase is $\phi = 0.01001011101011100000$, which is equal to the exact value in our minimal basis, to a precision of $\pm 2^{-20} E_h$. Note that the exact value has a remainder of $\delta \approx 0.5 E_h$ after a 20-bit expansion, hence the low contrast in the measured 20th bit.

photonic quantum computer[42,43]. The influence of noise is perhaps the most serious consideration[44] and must be overcome using error-correction and fault-tolerant constructions[17,34]. We note that an alternative promising path to efficient quantum simulators is to exploit controllable quantum systems that can be used to directly implement model Hamiltonians, thereby potentially avoiding the aforementioned resource-intensive approximation techniques and error correction[2,45].

Methods

Our model molecular Hamiltonian. We use the minimal STO-3G basis[46] for H_2, consisting of one $|1s\rangle$–type atomic orbital per atom. These two functions are combined to form the bonding (antibonding) molecular orbitals[47]. The four corresponding single-electron molecular spin-orbitals are combined antisymmetrically to form the six two-electron configurations, $\{\Phi_i\}_{i=1}^6$ that form the basis for our simulation. Owing to symmetry, the Hamiltonian is block-diagonal in this basis, with blocks acting on each of the four subspaces spanned by $\{\Phi_1, \Phi_6\}$, $\{\Phi_2\}$, $\{\Phi_3, \Phi_4\}$ and $\{\Phi_5\}$ (see Supplementary Information for the particular forms of each configuration). Therefore, finding the eigenvalues of the two 2×2 submatrices in the Hamiltonian, $\hat{H}^{(1,6)}$ and $\hat{H}^{(3,4)}$, amounts to performing the FCI. Estimating the eigenvalues of 2×2 matrices is the simplest problem for the IPEA.

We use a propagator time step of $t = 1(\hbar/E_h)$ (the hartree, $E_h =$ 2,625.5 kJ mol^{-1}, is the atomic unit of energy), chosen so that $0 \leq Et/2\pi\hbar \leq 1$. For our proof-of-principle demonstration, all necessary molecular integrals are evaluated classically (see Supplementary Information) using the Hartree–Fock procedure[47]. We use these integrals to calculate the matrix elements of $\hat{H}$ and $\hat{U}$, then directly decompose each $\hat{U}^{(p,q)}$ operator into a logic gate network. We decompose the $\hat{U}^{(p,q)}$ operators into a global phase and a series of rotations of the one-qubit Hilbert space[34]:

$$\hat{U} = e^{i\alpha}\hat{R}_y(\beta)\hat{R}_z(\gamma)\hat{R}_y(-\beta) \qquad (2)$$

where α, β, and γ, are real angles. $\hat{U}^j$ is achieved by keeping β constant and replacing angles α and γ with $j\alpha$ and $j\gamma$, respectively. Our decomposition of the controlled-$\hat{U}^j$ is shown in Fig. 1b.

Details of classical computational methods. Restricted Hartree–Fock calculations were carried out on a classical computer using the STO-3G basis[46]. The software used was the PyQuante quantum chemistry package, version 1.6. The molecular integrals from the Hartree–Fock procedure are used to evaluate the matrix elements of the Hamiltonians $\hat{H}^{(1,6)}$ and $\hat{H}^{(3,4)}$, described in the main text.

Classical error correction technique. When running the IPEA, the probability of correctly identifying any individual bit with a single sample ($n = 1$) is reduced from unity by both theoretical (inexact phase expansion to k bits) and experimental factors (such as imperfect gates). However, as long as it remains above 0.5, repeated sampling and a majority vote will reduce the probability of error exponentially with n, in accordance with the Chernoff bound[34]. This technique allows for a significant increase in success probability, at the expense of repeating the experiment a fixed number of times. We note that this simple classical error correction technique can only play a small role when it comes to dealing with errors in large-scale implementations. The numerous errors in very large quantum logic circuits will make achieving a bit success probability over 0.5 a significant challenge. This will need to be addressed with quantum error correction techniques[17,34].

Count rates. We operate with a low-brightness optical source (spontaneous parametric downconversion pumping power 50 mW) to reduce the effects of unwanted multi-photon-pair emissions (which cannot be distinguished by our non-photon-number-resolving detectors and introduce error into the circuit operation). This yields about 15 coincident detection events per second at the output of our optical circuit. Therefore each iteration can be repeated 15 times a second. Reconfiguring the circuit for different iterations takes approximately seven seconds, largely due to the finite time required to rotate standard computer-controlled waveplate mounts. Therefore, obtaining a 20-bit estimation of a phase takes about three minutes, when using $n = 31$ samples to determine the logical state of each bit (as was used to achieve the results shown in Fig. 2). Note that approximately 95% of this time is spent rotating waveplates. In future implementations, this time could be reduced significantly using integrated photonics, for example, qubit manipulation using an electro-optically-controlled waveguide Mach–Zehnder interferometer[48].

References

1. Feynman, R. P. Simulating physics with computers. *Int. J. Theor. Phys.* **21,** 467–488 (1982).
2. Lloyd, S. Universal quantum simulators. *Science* **273,** 1073–1078 (1996).
3. Head-Gordon, M. & Artacho, E. Chemistry on the computer. *Physics Today* **61**(4), 58–63 (2008).
4. Hung, L. & Carter, E. A. Accurate simulations of metals at the mesoscale: Explicit treatment of 1 million atoms with quantum mechanics. *Chem. Phys. Lett.* **475,** 163–170 (2009).
5. Chelikowsky, J. R. *et al.* Pseudopotentials on grids: Application to the optical, and vibrational properties of silicon nanocrystals. *J. Comput. Theor. Nanosci.* **6,** 1247–1261 (2009).
6. Dreuw, A. & Head-Gordon, M. Failure of time-dependent density functional theory for long-range charge-transfer excited states: The zincbacteriochlorin-bacterlochlorin and bacteriochlorophyll-spheroidene complexes. *J. Am. Chem. Soc.* **126,** 4007–4016 (2004).
7. Levine, B. G. & Martinez, T. J. Isomerization through conical intersections. *Ann. Rev. Phys. Chem.* **58,** 613–634 (2007).
8. Lee, P. A., Nagaosa, N. & Wen, X.-G. Doping a Mott insulator: Physics of high-temperature superconductivity. *Rev. Mod. Phys.* **78,** 17–85 (2006).
9. Van Voorhis, T. & Head-Gordon, M. Benchmark variational coupled cluster doubles results. *J. Chem. Phys.* **113,** 8873–8879 (2000).
10. Abrams, D. & Lloyd, S. Simulation of many-body fermi systems on a universal quantum computer. *Phys. Rev. Lett.* **79,** 2586–2586 (1997).
11. Kassal, I. Jordan, S. P., Love, P. J., Mohseni, M. & Aspuru-Guzik, A. Polynomial-time quantum algorithm for the simulation of chemical dynamics *Proc. Natl Acad. Sci. USA* **105,** 18681–18686 (2008).
12. Zalka. C. Efficient simulation of quantum systems by quantum computers. *Proc. R. Soc. Lond. A* **454,** 313–322 (1998).
13. Kassal, I. & Aspuru-Guzik, A. Quantum algorithm for molecular properties and geometry optimization. *J. Chem. Phys.* (in the press); preprint at http://arxiv.org/abs/0908.1921 (2009).
14. Lidar D. A. & Wang, H. Calculating the thermal rate constant with exponential speedup on a quantum computer. *Phys. Rev. E* **59,** 2429–2438 (1999).
15. Aspuru-Guzik, A., Dutoi, A., Love, P. J. & Head-Gordon, M. Simulated quantum computation of molecular energies. *Science* **309,** 1704–1707 (2005).
16. Brown, K. R., Clark, R. J. & Chuang, I. L. Limitations of quantum simulation examined by simulating a pairing Hamiltonian using nuclear magnetic resonance. *Phys. Rev. Lett.* **97,** 050504 (2006).
17. Clark C. R., Metodi, T. S., Gasster, S. D. & Brown, K. R. Resource requirements for fault-tolerant quantum simulation: The ground state of the transverse Ising model. *Phys. Rev. A* **79,** 062314 (2009).
18. Somaroo, S., Tseng, C. H., Havel, T. F., Laflamme, R. & Cory, D. G. Quantum simulations on a quantum computer. *Phys. Rev. Lett.* **82,** 5381–5384 (1999).
19. Yang, X., Wang, A. M., Xu, F. & Du. J. Experimental simulation of a pairing Hamiltonian on an NMR quantum computer. *Chem. Phys. Lett.* **422,** 20–24 (2006).
20. Friedenauer, A., Schmitz, H., Glueckert, J. T., Porras, D. & Schaetz, T. Simulating a quantum magnet with trapped ions. *Nature Phys.* **4,** 757–761 (2008).
21. Gerritsma, R. *et al.* Quantum simulation of the Dirac equation. Preprint at: http://arxiv.org/abs/0909.0674 (2009).
22. Vidal, G. Efficient classical simulation of slightly entangled quantum computations. *Phys. Rev. Lett.* **91,** 147902 (2003).
23. Schmidt-Kaler, F. *et al.* Realization of the Cirac–Zoller controlled-NOT quantum gate *Nature* **422,** 408–411 (2003).
24. Leibfried, D. *et al.* Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate. *Nature* **422,** 412–415 (2003).
25. O'Brien, J. L., Pryde, G. J., White, A. G., Ralph, T. C. & Branning, D. Demonstration of an all-optical quantum controlled-NOT gate. *Nature* **426,** 264–267 (2003).
26. Plantenberg, H. J., de Groot, P. C., Harmans, C. J. P. M. & Mooij, J. E. Demonstration of controlled-NOT quantum gates on a pair of superconducting quantum bits. *Nature* **447,** 836–839 (2007).
27. Pashkin, Y. A. *et al.* Quantum oscillations in two coupled charge qubits. *Nature* **421,** 823–826 (2003).
28. Knill, E., Laflamme, R. & Milburn, G. J. A scheme for efficient quantum computation with linear optics. *Nature* **409,** 46–52 (2001).
29. Kok, P. *et al.* Linear optical quantum computing with photonic qubits. *Rev. Mod. Phys.* **79,** 135 (2007).
30. Lanyon, B. P. *et al.* Experimental demonstration of a compiled version of Shor's Algorithm with quantum entanglement. *Phys. Rev. Lett.* **99,** 250505 (2007).
31. Lanyon, B. P. *et al.* Simplifying quantum logic using higher-dimensional Hilbert spaces. *Nature Phys.* **5,** 134–140 (2009).
32. Kitaev, A. Quantum measurements and the Abelian Stabilizer Problem. Preprint at: http://arxiv.org/abs/quant-ph/9511026 (1995).
33. Dobsicek, M., Johansson, G., Shumeiko, V. S. & Wendin, G. Arbitrary accuracy iterative phase estimation algorithm as a two qubit benchmark. *Phys. Rev. A* **76,** 030306(R) (2007).
34. Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information* (Cambridge Univ. Press, 2001).
35. Helgaker, T., Jorgensen, P. & Olsen, J. *Modern Electronic Structure Theory* (Wiley, 2000).
36. Xiu-Mei, L., Jun, L. & Xian-Ping, S. Experimental realization of arbitrary accuracy iterative phase estimation algorithms on ensemble quantum computers. *Chinese Phys. Lett.* **24,** 3316–3319 (2007).
37. Braunstein, S. L. *et al.* Separability of very noisy mixed states and implications for NMR quantum computing. *Phys. Rev. Lett.* **83,** 1054–1057 (1999).
38. Chiaverini J. *et al.* Implementation of the semiclassical quantum Fourier transform in a scalable system. *Science* **308,** 997–1000 (2005).
39. Nielsen, M. A. Optical quantum computation using cluster states. *Phys. Rev. Lett.* **93,** 040503 (2004).
40. Wu, L.-A., Byrd, M. S. & Lidar, D. A. Polynomial-time simulation of pairing models on a quantum computer. *Phys. Rev. Lett.* **89,** 057904 (2002).
41. Farhi, E., Goldstone, J., Gutmann, S. & Sipser, M. Quantum adiabatic evolution algorithms with different paths. *Science* **292,** 472–475 (2000).
42. Grangier, P., Sanders, B. & Vuckovic, J. (eds) Special issue: Focus on single photons on demand *New J. Phys.* **6,** (2004).

43. Cheung, J., Migdal, A. & Rastello, M.-L. (eds) Special issue: Single-photon: detectors, applications, and measurement methods. *J. Mod. Opt.* **56**, 2–3 (2009).
44. Dür, W., Bremner, M. J. & Briegel, H. J. Quantum simulation of interacting high-dimensional systems: the influence of noise. *Phys. Rev. A* **78**, 052325 (2008).
45. Jané, E., Vidal, G., Dür, W., Zoller, P. & Cirac, J. I. Simulation of quantum dynamics with quantum optical systems. *Quant. Inf. Comp.* **3**, 15–37 (2003).
46. Hehre, W. J., Stewart, R. F. & Pople, J. A. Self-consistent molecular orbital methods I. Use of Gaussian expansions of slater type atomic orbitals. *J. Chem. Phys.* **51**, 2657–2664 (1969).
47. Szabo, A. & Ostlund, N. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory* (Dover Publications, 1996).
48. Liao, Y. *et al*. Electro-optic integration of embedded electrodes and waveguides in LiNbO$_3$using a femtosecond laser. *Opt. Lett.* **33**, 2281–2283 (2008).

Author contributions

B.P.L., J.D.W., I.K., M.M., A.A.G. and A.G.W. conceived and designed the experiments, B.P.L., G.G.G., M.E.G. and M.P.A. performed the experiments, B.P.L. and G.G.G. analysed the data, J.D.W. performed the classical preprocessing. All authors discussed the results and co-wrote the manuscript.

Additional information

The authors declare no competing financial interests. Supplementary information accompanies this paper at www.nature.com/naturechemistry. Reprints and permission information is available online at http://npg.nature.com/reprintsandpermissions/. Correspondence and requests for materials should be addressed to B.P.L. and A.A.-G.