

ARTICLE

Received 9 Dec 2013 | Accepted 27 May 2014 | Published 23 Jul 2014

DOI: 10.1038/ncomms5213

OPEN

A variational eigenvalue solver on a photonic quantum processor

Alberto Peruzzo^{1,*†}, Jarrod McClean^{2,*}, Peter Shadbolt¹, Man-Hong Yung^{2,3}, Xiao-Qi Zhou¹, Peter J. Love⁴, Alán Aspuru-Guzik² & Jeremy L. O'Brien¹

Quantum computers promise to efficiently solve important problems that are intractable on a conventional computer. For quantum systems, where the physical dimension grows exponentially, finding the eigenvalues of certain operators is one such intractable problem and remains a fundamental challenge. The quantum phase estimation algorithm efficiently finds the eigenvalue of a given eigenvector but requires fully coherent evolution. Here we present an alternative approach that greatly reduces the requirements for coherent evolution and combine this method with a new approach to state preparation based on ansätze and classical optimization. We implement the algorithm by combining a highly reconfigurable photonic quantum processor with a conventional computer. We experimentally demonstrate the feasibility of this approach with an example from quantum chemistry—calculating the ground-state molecular energy for He-H⁺. The proposed approach drastically reduces the coherence time requirements, enhancing the potential of quantum resources available today and in the near future.

¹Centre for Quantum Photonics, H.H. Wills Physics Laboratory & Department of Electrical and Electronic Engineering, University of Bristol, Bristol BS8 1UB, UK. ²Department of Chemistry and Chemical Biology, Harvard University, Cambridge, Massachusetts 02138, USA. ³Center for Quantum Information, Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing 100084, P. R. China. ⁴Department of Physics, Haverford College, Haverford, Pennsylvania 19041, USA. * These authors contributed equally to this work. † Present address: School of Physics, University of Sydney, Sydney, New South Wales 2006, Australia. Correspondence and requests for materials should be addressed to A.P. (email: Alberto.Peruzzo@sydney.edu.au) or to A.A.-G. (email: Aspuru@chemistry.harvard.edu) or to J.L.O'B. (email: Jeremy.obrien@bristol.ac.uk).

In chemistry, the properties of atoms and molecules can be determined by solving the Schrödinger equation. However, because the dimension of the problem grows exponentially with the size of the physical system under consideration, exact treatment of these problems remains classically infeasible for compounds with more than 2–3 atoms¹. Many approximate methods² have been developed to treat these systems, but efficient, exact methods for large chemical problems remain out of reach for classical computers. Beyond chemistry, the solution of large eigenvalue problems³ would have applications ranging from determining the results of internet search engines⁴ to designing new materials and drugs⁵.

Recent developments in the field of quantum computation offer a way forward for determining efficient solutions of many instances of large eigenvalue problems that are classically intractable^{6–12}. Quantum approaches to finding eigenvalues have previously relied on the quantum phase estimation (QPE) algorithm. The QPE algorithm offers an exponential speedup over classical methods and requires a number of quantum operations $O(p^{-1})$ to obtain an estimate with precision p (refs 13–18). In the standard formulation of QPE, one assumes the eigenvector $|\psi\rangle$ of a Hermitian operator \mathcal{H} is given as input and the problem is to determine the corresponding eigenvalue λ . The time the quantum computer must remain coherent is determined by the necessity of $O(p^{-1})$ successive applications of $e^{-i\mathcal{H}t}$, each of which can require on the order of millions or billions of quantum gates for practical applications^{17,19}, as compared to the tens to hundreds of gates achievable in the short term.

Here we introduce an alternative to QPE that significantly reduces the requirements for coherent evolution. We have developed a reconfigurable quantum processing unit (QPU), which efficiently calculates the expectation value of a Hamiltonian (\mathcal{H}), providing an exponential speedup over exact diagonalization, the only known exact solution to the problem on a traditional computer. The QPU has been experimentally implemented using integrated photonics technology with a spontaneous parametric downconversion single-photon source and combined with an optimization algorithm run on a classical processing unit (CPU), which variationally computes the eigenvalues and eigenvectors of \mathcal{H} . By using a variational algorithm, this approach reduces the requirement for coherent evolution of the quantum state, making more efficient use of quantum resources, and may offer an alternative route to practical quantum-enhanced computation.

Results

Quantum expectation estimation. The quantum expectation estimation (QEE) algorithm computes the expectation value of a given Hamiltonian \mathcal{H} for an input state $|\psi\rangle$. Any Hamiltonian may be written as

$$\mathcal{H} = \sum_{i\alpha} h_{\alpha}^i \sigma_{\alpha}^i + \sum_{ij\alpha\beta} h_{\alpha\beta}^{ij} \sigma_{\alpha}^i \sigma_{\beta}^j + \dots \quad (1)$$

for real h , where Roman indices identify the subsystem on which the operator acts, and Greek indices identify the Pauli operator, for example, $\alpha = x$. Note that no assumption about the dimension or structure of the hermitian Hamiltonian is needed for this expansion to be valid. By exploiting the linearity of quantum observables, it follows that

$$\langle \mathcal{H} \rangle = \sum_{i\alpha} h_{\alpha}^i \langle \sigma_{\alpha}^i \rangle + \sum_{ij\alpha\beta} h_{\alpha\beta}^{ij} \langle \sigma_{\alpha}^i \sigma_{\beta}^j \rangle + \dots \quad (2)$$

We consider Hamiltonians that can be written as a polynomial number of terms, with respect to the system size. This class of Hamiltonians encompasses a wide range of physical systems,

including the electronic structure Hamiltonian of quantum chemistry, the quantum Ising Model, the Heisenberg Model^{20,21}, matrices that are well approximated as a sum of n -fold tensor products^{22,23}, and more generally any k -sparse Hamiltonian without evident tensor product structure (see Supplementary Methods for details). Thus, the evaluation of $\langle \mathcal{H} \rangle$ reduces to the sum of a polynomial number of expectation values of simple Pauli operators for a quantum state $|\psi\rangle$, multiplied by some real constants. A quantum device can efficiently evaluate the expectation value of a tensor product of an arbitrary number of simple Pauli operators²³. Therefore, with an n -qubit state we can efficiently evaluate the expectation value of this $2^n \times 2^n$ Hamiltonian.

One might attempt this using a classical computer by separately optimizing all reduced states corresponding to the desired terms in the Hamiltonian, but this would suffer from the N -representability problem, which is known to be intractable for both classical and quantum computers (it is in the quantum complexity class QMA-Hard²⁴). The power of our approach derives from the fact that quantum hardware can store a global quantum state with exponentially fewer resources than required by classical hardware, and as a result the N -representability problem does not arise.

The expectation value of a tensor product of an arbitrary number of Pauli operators can be estimated by local measurement of each qubit⁶. Such independent measurements can be performed in parallel, incurring a constant cost in time. Furthermore, since these operators are normalized and finite-dimensional, their spectra are bounded. As a result, each $\langle H_i^m \rangle = h_{\alpha\beta\dots}^{ij\dots} \langle \sigma_{\alpha}^i \otimes \sigma_{\beta}^j \dots \rangle$ can be estimated to a precision p of an individual element with coefficient h , which is an arbitrary element from the set of constants $\{h_{\alpha\beta\dots}^{ij\dots}\}$, at a cost of $O(|h_{\max}|^2 M p^{-2})$ repetitions. Here M is the number of terms in the decomposition of the Hamiltonian and h_{\max} is the coefficient with maximum norm in the decomposition of the Hamiltonian. The advantage of this approach is that the coherence time to make a single measurement after preparing the state is $O(1)$. Conversely, the disadvantage of this approach with respect to QPE is the scaling in the total number of operations, as a function of the desired precision is quadratically worse ($O(p^{-2})$ versus $O(p^{-1})$). Moreover, this scaling will also reflect the number of state preparation repetitions required, whereas in QPE the number of state preparation steps is constant. In essence, we dramatically reduce the coherence time requirement while maintaining an exponential advantage over the classical case, by adding a polynomial number of repetitions with respect to QPE.

Quantum variational eigensolver. The procedure outlined above replaces the long coherent evolution required by QPE by many short coherent evolutions. In both QPE and QEE we require a good approximation to the ground-state wavefunction to compute the ground-state eigenvalue, and we now consider this problem. Previous approaches have proposed to prepare ground states by adiabatic evolution¹⁵, or by the quantum Metropolis algorithm^{25,26}. Unfortunately both of these require long coherent evolution. The quantum variational eigensolver (QVE) algorithm is a variational method to prepare the eigenstate and, by exploiting QEE, requires short coherent evolution. QEE and QVE and their relationship are shown in Fig. 1 and detailed in the Supplementary Methods.

It is well known that the eigenvalue problem for an observable represented by an operator \mathcal{H} can be restated as a variational problem on the Rayleigh–Ritz quotient^{27,28}, such that the eigenvector $|\psi\rangle$ corresponding to the lowest eigenvalue is the




Figure 1 | Architecture of the quantum-variational eigensolver. In QEE, quantum states that have been previously prepared are fed into the quantum modules, which compute $\langle H_i \rangle$, where H_i is any given term in the sum defining \mathcal{H} . The results are passed to the CPU, which computes $\langle \mathcal{H} \rangle$. In the quantum variational eigensolver, the classical minimization algorithm, run on the CPU, takes $\langle \mathcal{H} \rangle$ and determines the new state parameters, which are then fed back to the QPU.

$|\psi\rangle$ that minimizes

$$\frac{\langle \psi | \mathcal{H} | \psi \rangle}{\langle \psi | \psi \rangle}. \tag{3}$$

By varying the experimental parameters in the preparation of $|\psi\rangle$ and computing the Rayleigh–Ritz quotient using QEE as a subroutine in a classical minimization, one may prepare unknown eigenvectors. At the termination of the algorithm, a simple prescription for the reconstruction of the eigenvector is stored in the final set of experimental parameters that define $|\psi\rangle$.

If a quantum state is characterized by an exponentially large number of parameters, it cannot be prepared with a polynomial number of operations. The set of efficiently preparable states are therefore characterized by polynomially many parameters, and we choose a particular set of ansatz states of this type. Under these conditions, a classical search algorithm on the experimental parameters that define $|\psi\rangle$ needs only explore a polynomial number of dimensions—a requirement for the search to be efficient. One example of a quantum state parameterized by a polynomial number of parameters for which there is no known efficient classical implementation is the unitary coupled cluster ansatz²⁹

$$|\Psi\rangle = e^{T - T^\dagger} |\Phi\rangle_{\text{ref}}. \tag{4}$$

where $|\Phi\rangle_{\text{ref}}$ is some reference state, usually the Hartree Fock ground state, and T is the cluster operator for an N electron system, defined by

$$T = T_1 + T_2 + T_3 + \dots + T_N, \tag{5}$$

where

$$T_1 = \sum_{pr} t_p^r \hat{a}_p^\dagger \hat{a}_r \tag{6}$$

$$T_2 = \sum_{pqrs} t_{pq}^{rs} \hat{a}_p^\dagger \hat{a}_q^\dagger \hat{a}_r \hat{a}_s \tag{7}$$

and higher-order terms follow logically. It is clear that by construction the operator $(T - T^\dagger)$ is anti-hermitian, and exponentiation maps it to a unitary operator $U = e^{(T - T^\dagger)}$. For any fixed excitation level k , the reduced cluster operator is

written as

$$T^{(k)} = \sum_{i=1}^k T_i. \tag{8}$$

In general no efficient implementation of this ansatz has yet been developed for a classical computer, even for low-order cluster operators, due to the non-truncation of the BCH series²⁹. However, this state may be prepared efficiently on a quantum device. The reduced anti-hermitian cluster operator $(T^{(k)} - T^{(k)\dagger})$ is the sum of a polynomial number of terms—namely, it contains a number of terms $O(N^k(M - N)^k)$, where M is the number of single-particle orbitals. By defining an effective Hermitian Hamiltonian $\mathcal{H} = i(T^{(k)} - T^{(k)\dagger})$ and performing the Jordan–Wigner transformation to reach a Hamiltonian that acts on the space of qubits, $\tilde{\mathcal{H}}$, we are left with a Hamiltonian that is a sum of polynomially many products of Pauli operators. The problem then reduces to the quantum simulation of this effective Hamiltonian, $\tilde{\mathcal{H}}$, which can be done in polynomial time using the procedure outlined by Ortiz *et al.*²³ We note that while this state preparation procedure utilizes tools from quantum simulation, the total effective time of evolution is fixed by the expansion coefficients t_{pq}^{rs} . This is in contrast to the normal difficulties encountered in QPE, where simulations must be carried out for times that are exponential in the desired bits of precision.

While there is currently no known efficient classical algorithm based on these ansatz states, non-unitary coupled cluster ansatz is sometimes referred to as the ‘gold standard of quantum chemistry’ as it is the standard of accuracy to which other methods in quantum chemistry are often compared. The unitary version of this ansatz is thought to yield superior results to even this ‘gold standard’²⁹.

Prototype demonstration. We have implemented the QPU using integrated quantum photonics technology³⁰. Our device, shown schematically in Fig. 2, is a reconfigurable waveguide chip that can prepare and measure arbitrary two-bit pure states using several single-qubit rotations and one two-qubit entangling gate. The state is path-encoded using photon pairs generated via a spontaneous parametric downconversion process. State

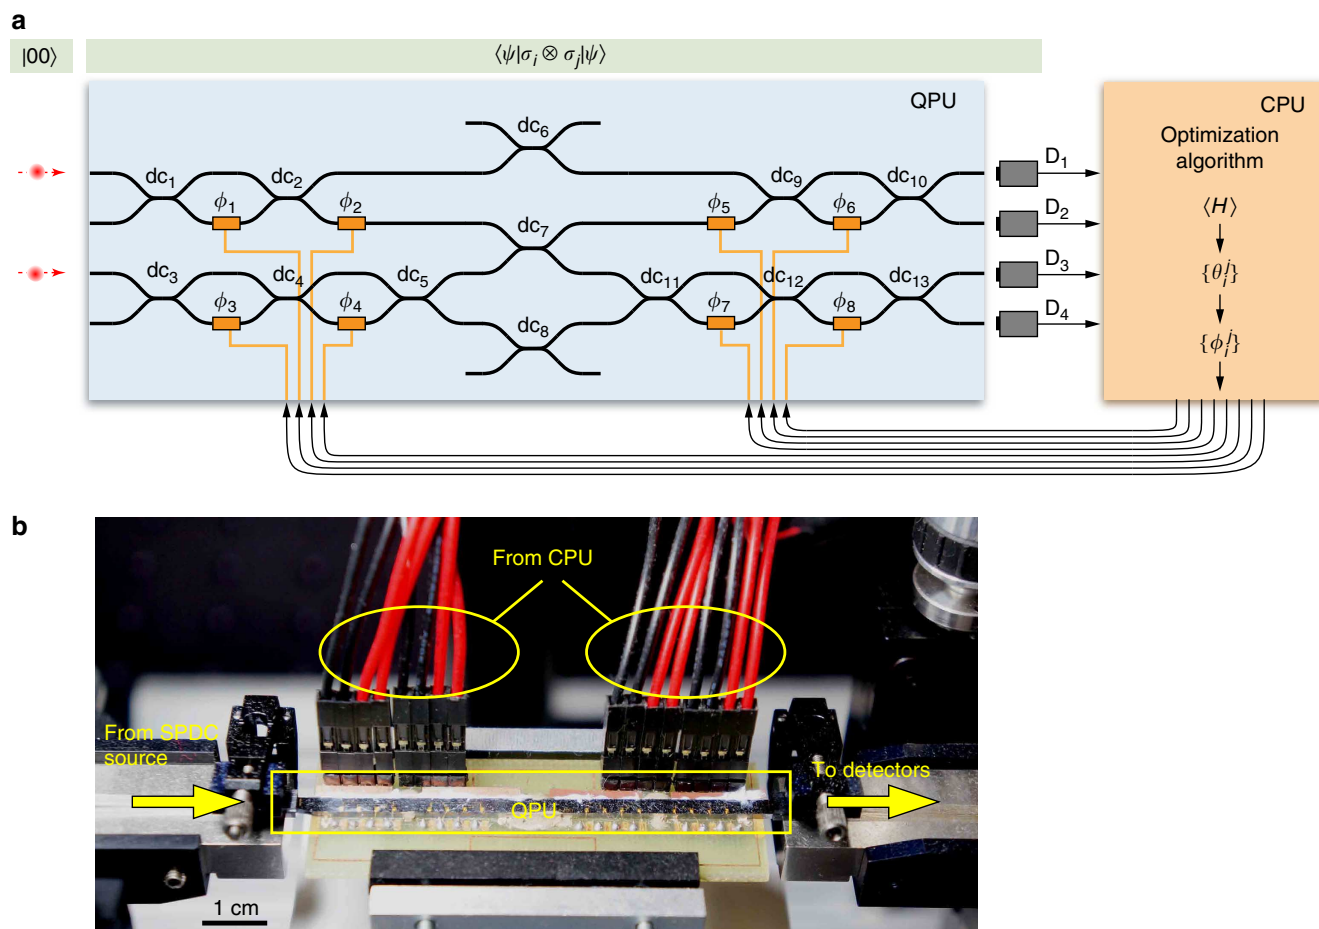


Figure 2 | Experimental implementation of our scheme. (a) Quantum-state preparation and measurement of the expectation values $\langle \psi | \sigma_i \otimes \sigma_j | \psi \rangle$ are performed using a quantum photonic chip. Photon pairs, generated using spontaneous parametric downconversion, are injected into the waveguides encoding the $|00\rangle$ state. The state $|\psi\rangle$ is prepared using thermal phase shifters ϕ_{1-8} (orange rectangles) and one CNOT gate and measured using photon detectors. $dc_{\{1-4,9-13\}}$ (dc_{5-7}) are 50% (30%) reflectivity directional couplers. Coincidence count rates from the detectors D_{1-4} are passed to the CPU running the optimization algorithm. This computes the set of parameters for the next state and writes them to the quantum device. (b) A photograph of the QPU.

preparation and measurement in the Pauli basis is achieved by setting 8 voltage-driven phase shifters and counting photon detection events with silicon single-photon detectors³¹.

The ability to prepare an arbitrary two-qubit separable or entangled state enables us to investigate 4×4 Hamiltonians. For the experimental demonstration of our algorithm we choose a problem from quantum chemistry—namely, determining the bond dissociation curve of the molecule He-H^+ in a minimal basis. The full configuration interaction Hamiltonian for this system has dimension 4, and can be written compactly as

$$\mathcal{H}(R) = \sum_{i\alpha} h_{\alpha}^i(R) \sigma_{\alpha}^i + \sum_{ij\alpha\beta} h_{\alpha\beta}^{ij}(R) \sigma_{\alpha}^i \sigma_{\beta}^j. \quad (9)$$

The coefficients $h_{\alpha}^i(R)$ and $h_{\alpha\beta}^{ij}(R)$ were determined using the PSI3 computational package³² and are tabulated in Supplementary Table 2.

In order to compute the bond dissociation of the molecule, we use QVE to compute its ground state for a range of values of the nuclear separation R . In Fig. 3 we report a representative optimization run for a particular nuclear separation, demonstrating the convergence of our algorithm to the ground state of $\mathcal{H}(R)$ in the presence of experimental noise. Figure 3a demonstrates the convergence of the average energy, while Fig. 3b demonstrates the convergence of the overlap $|\langle \psi | \psi^G \rangle|$ of the current state $|\psi^j\rangle$

with the target state $|\psi^G\rangle$. The colour of each entry in Fig. 3a represents the tangle (absolute concurrence squared) of the state at that step of the algorithm. It is known that the volume of separable states is doubly exponentially small with respect to the rest of state space³³. Thus, the ability to traverse non-separable state space increases the number of paths by which the algorithm can converge and will be a requirement for future large-scale implementations. Moreover, it is clear that the ability to produce entangled states is a necessity for the accurate description of general quantum systems where eigenstates may be non-separable—for example, the ground state of the He-H^+ Hamiltonian has small but not negligible tangle.

Repeating this procedure for several values of R , we obtain the bond dissociation curve, which is reported in Fig. 4. After the computed energies have been corrected for experimental errors, the determination of the equilibrium bond length of the molecule was found to be $R = 92.3 \pm 0.1$ pm, with a corresponding ground-state electronic energy of $E = -2.865 \pm 0.008$ MJ mol⁻¹. Full details of the correction for systematic errors and estimation of the uncertainty on E are reported in the Supplementary Methods. The corresponding theoretical curve shows the numerically exact energy derived from a full configuration interaction calculation of the molecular system in the same basis. More than 96% of the experimental data are within chemical accuracy with respect to

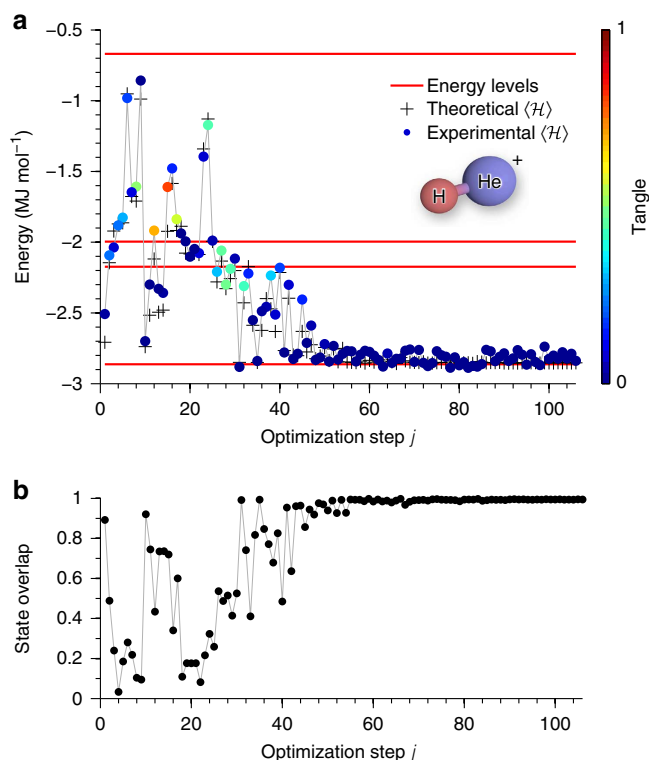


Figure 3 | Finding the ground state of He-H⁺ for a specific molecular separation $R = 90$ pm. (a) Experimentally computed energy $\langle \mathcal{H} \rangle$ (coloured dots) as a function of the optimization step j . The colour represents the tangle (degree of entanglement) of the physical state, estimated directly from the state parameters $\{\phi_i^j\}$. The red lines indicate the energy levels of $\mathcal{H}(R)$. The optimization algorithm clearly converges to the ground state of the molecule, which has small but non-zero tangle. The crosses show the energy calculated at each experimental step, assuming an ideal quantum device. (b) Overlap $|\langle \psi^j | \psi^G \rangle|$ between the experimentally computed state $|\psi^j\rangle$ at each optimization step j and the theoretical ground state of \mathcal{H} , $|\psi^G\rangle$. Error bars are smaller than the data points. Further details are provided in the Methods section, Supplementary Table 1 and Supplementary Methods.

the theoretical values. At the conclusion of the optimization, we retain full knowledge of the experimental parameters, which can be used for efficient reconstruction of the state $|\psi\rangle$ in the event that additional physical or chemical properties are required.

Discussion

QEE uses relatively few quantum resources compared to QPE. Broadly speaking, QPE requires a large number of n -qubit quantum controlled operations to be performed in series—placing considerable demands on the number of components and coherence time—while the inherent parallelism of our scheme enables a small number of n -qubit gates to be exploited many times, drastically reducing these demands. Moreover, adding control to arbitrary unitary operations in practice is difficult, if not impossible, for current quantum architectures (although a proposed scheme to add control to arbitrary unitary operations has recently been demonstrated³⁴). To give a numerical example, the QPE circuit for a 4×4 Hamiltonian such as that demonstrated here would require at least 12 CNOT gates, while our method only requires one. We note that the resource saving provided by QEE incurs a cost of polynomial repetitions of the state preparation, as compared to the single copy required by




Figure 4 | Bond dissociation curve of the He-H⁺ molecule. This curve is obtained by repeated computation of the ground-state energy (as shown in Fig. 3) for several $\mathcal{H}(R)$ values. The magnified plot shows that after correction for the measured systematic error the data overlap with the theoretical energy curve, and, importantly, we can resolve the molecular separation of minimal energy. Error bars show the standard deviation of the computed energy, as described in the Methods section.

QPE. In many cases (for example, our photonic implementation), repeated preparation of a state is not significantly harder than preparation of a single copy, requiring only a polynomial overhead in time without any modification of the device.

In implementing QVE, the device prepares ansatz states that are defined by a polynomial set of parameters. This ansatz might be chosen based on knowledge of the physical system of interest (as for the unitary coupled cluster and typical quantum chemistry ansätze), thus determining the device design. However, our architecture allows for an alternative and potentially more promising approach, where the device is first constructed based on the available resources and we define the set of states that the device can prepare as the ‘device ansatz’. Due to the quantum nature of the device, this ansatz can be very distinct from those used in traditional quantum chemistry. With this alternative approach the physical implementation is then given by a known sequence of quantum operations with adjustable parameters—determined at the construction of the device—with a maximum depth fixed by the coherence time of the physical qubits. This approach, while approximate, provides a variationally optimal solution for the given quantum resources and may still be able to provide qualitatively correct solutions, just as approximate methods do in traditional quantum chemistry (for example, Hartree Fock). The unitary coupled cluster ansatz (equation (4)) provides a concrete example where our approach provides an exponential advantage over known classical techniques. For this ansatz, with as few as 40–50 qubits, one expects to manipulate a state that is not efficient to simulate classically, and can provide a solution superior to the classical gold standard, non-unitary coupled cluster.

We have developed and experimentally implemented a new approach to solving the eigenvalue problem with quantum hardware. QEE shares with QPE the need to prepare a good

approximation to the ground state, but replaces a single long coherent evolution by a number of shorter coherent calculations proportional to the number of terms in the Hamiltonian. While the effect of errors on each of these calculations is the same as in QPE, the reliance on a number of separate calculations makes the algorithm sensitive to variations in state preparation between the separate quantum calculations. This effect requires further investigation. The most general local Hamiltonian problem is QMA-complete³⁵. However, under the reasonable assumption that a good approximation to the state can be prepared, our method and QPE can both efficiently estimate the energy of the state, and it is in this setting that we compare them. In QVE, we experimentally implemented a ground-state preparation procedure through a direct variational algorithm on the control parameters of the quantum hardware. The prepared state could be utilized in either QEE 1 or QPE if desired. Larger calculations will require a choice of ansatz, for which there are two possibilities. One could experimentally implement chemically motivated ansatz such as the unitary coupled cluster method. Alternatively, one could pursue those ansätze that are most easy to implement experimentally—creating a new set of device ansatz states that would require classification in terms of their overlap with chemical ground states. Such a classification would be a good way to determine the value of a given experimental advance—for ground-state problems it is best to focus limited experimental resources on those efforts that will most enhance the overlap of preparable states with chemical ground states. In addition to the above issues, which we leave to future work, an interesting avenue of research is to ask whether the conceptual approach described here could be used to address other intractable problems with quantum-enhanced computation. Examples that can be mapped to the ground-state problem, and where the N -representability problem does not occur, include search engine optimization and image recognition. It should be noted that the approach presented here requires no control or auxiliary qubits, relying only on measurement techniques that are already well established. For example, in the two-qubit case, these measurements are identical to those performed in Bell inequality experiments.

Quantum simulators with only a few tens of qubits are expected to outperform the capabilities of conventional computers, not including open questions regarding fault tolerance and errors/precision. Our scheme would allow such devices to be implemented using dramatically less resources than the current best known approach.

Methods

Classical optimization algorithm. For the classical optimization step of our integrated processor we implemented the Nelder–Mead (NM) algorithm³⁶, a simplex-based direct search (DS) method for unconstrained minimization of objective functions. Although in general NM can fail because of the deterioration of the simplex geometry or lack of sufficient decrease, the convergence of this method can be greatly improved by adopting a restarting strategy. Although other DS methods, such as the gradient descent, can perform better for smooth functions, these are not robust to the noise, which makes the objective function non-smooth under experimental conditions. NM has the ability to explore neighbouring valleys with better local optima, and likewise this exploring feature usually allows NM to overcome non-smoothnesses. We verified that the gradient descent minimization algorithm is not able to converge to the ground state of our Hamiltonian under the experimental conditions, mainly due to the Poissonian noise associated with our photon source and the accidental counts of the detection system, while NM converged to the global minimum in most optimization runs.

Mapping from the state parameters to the chip phases. The set of phases $\{\theta_i\}$, which uniquely identifies the state $|\psi\rangle$, is not equivalent to the phases that are written to the photonic circuit $\{\phi_i\}$, since the chip phases are also used to implement the desired measurement operators $\sigma_x \otimes \sigma_\beta$. Therefore, knowing the desired state parameters and measurement operator we compute the appropriate values of the chip phases on the CPU at each iteration of the optimization algorithm. The algorithm for finding the state parameters $\{\theta_i\}$ for an arbitrary two-qubit state is

described in the Methods. These phases are then applied to the CNOT-MZ chip using $\phi_{1,2,3,4,7,8}$. Here $\phi_{7,8}$ are modified to account for the choice of measurement setting at the target qubit. (Any single-qubit projective measurement can be performed using an MZI together with two phase shifters.) The measurement setting for the control qubit is implemented using $\phi_{5,6}$.

Estimation of the error on $\langle \mathcal{H} \rangle$. We performed measurements of the statistical and systematic errors that affect our computation of $\langle \mathcal{H} \rangle$.

Statistical errors. Statistical errors due to the Poissonian noise associated with single-photon statistics are intrinsic to the estimation of expectation values in quantum mechanics.

These errors can be arbitrarily reduced at a sublinear cost of measurement time (that is, efficiently) since the magnitude of error is proportional to the square root of the count rate. We experimentally measured the standard deviation of an expectation value $\langle \mathcal{H}_i \rangle$ for a particular state using 50 trials. The total average coincidence rate was $\sim 1,500 \text{ s}^{-1}$. The standard deviation was found to be 37 kJ mol^{-1} , which is comparable to the error observed in the measurement of the ground-state energy shown in Fig. 4.

The minima of the potential energy curve was determined by a generalized least squares procedure to fit a quadratic curve to the experimental data points in the region $R = (80, 100) \text{ pm}$, as is common in the use of trust region searches for minima³⁷, using the inverse experimentally measured variances as weights. Covariances determined by the generalized least squares procedure were used as input to a Monte Carlo sampling procedure to determine the minimum energy and equilibrium bond distance as well as their uncertainties assuming Gaussian random error. The uncertainties reported represent standard deviations. Sampling error in the Monte Carlo procedure was $3 \times 10^{-4} \text{ pm}$ for the equilibrium bond distance and $3 \times 10^{-8} \text{ MJ mol}^{-1}$ for the energy.

In Fig. 4, the large deviations from the theoretical line result from the coincidental impact of noise resulting in premature optimization termination. These points could have been rerun or eliminated using the prior knowledge of smoothness of the dissociation curve. However, to accurately portray the performance of the algorithm exactly as described, with no expert interference, these points are retained.

Systematic errors. In all the measurements described above we observed a constant and reproducible small shift, $\epsilon = 50 \text{ kJ mol}^{-1}$, of the expectation value with respect to the theoretical value of the energy. There are at least three effects that contribute to this systematic error.

Firstly, the downconversion source that we use in our experiment does not produce the pure two-photon state that is required for high-fidelity quantum interference. In particular, higher-order photon number terms and, more significantly, photon distinguishability both degrade the performance of our entangling gate and thus the preparation of the state $|\psi\rangle$. This results in a shift of the measured energy $\langle \psi | \mathcal{H} | \psi \rangle$. Higher-order terms could be effectively eliminated by use of true single-photon sources (such as quantum dots or nitrogen vacancy centers in diamond), and there is no fundamental limit to the degree of indistinguishability that can be achieved through improved state engineering.

Secondly, imperfections in the implementation of the photonic circuit also reduce the fidelity with which $|\psi\rangle$ is prepared and measured. Small deviations from designed beamsplitter reflectivities and interferometer path lengths, as well as imperfections in the calibration of voltage-controlled phase shifters used to manipulate the state, all contribute to this effect. However, these are technological limitations that can be greatly improved in future realizations.

Finally, unbalanced input and output coupling efficiency also results in skewed two-photon statistics, again shifting the measured expectation value of $\langle \mathcal{H} \rangle$.

Another systematic effect that can be noted in Fig. 4 is that the magnitude of the error on the experimental estimation of the ground-state energy increases with R . This is due to the fact that as R increases the first and second excited eigenstates of this Hamiltonian become degenerate, resulting in increased difficulty for the classical minimization, generating mixtures of states that increases the overall variance of the estimation.

Quantum-state fidelity. In a previous work³¹, we measured the average state fidelity of states generated by the CNOT gate, estimated by quantum process tomography, to be 0.873 ± 0.001 . The average quantum-state fidelity over four Bell states was 0.93. The average fidelity across 995 configurations (equivalent to many truth tables in many bases) was 0.990 ± 0.009 , with 96% of configurations producing photon statistics with $f > 0.97$.

Count rate. In our experiment the mean count rate, which directly determines the statistical error, was $\sim 2,000$ – $4,000$ twofold events per measurement. The expectation value of a given Hamiltonian was reconstructed at each point from four two-qubit Pauli measurements. For the bond dissociation curve we measured about 100 points per optimization run. In the full dissociation curve we found the ground states of 79 Hamiltonians. The full experiment was performed in about 158 h.

State preparation is relatively fast, requiring a few milliseconds to set the phases on the chip. However, 17 s is required for cooling the chip, resulting in a duty cycle of ~5%. The purpose of this is to overcome the instability of the fibre-to-chip coupling due to thermal expansion of the chip during operation. This will not be an issue in future implementations, where fibres will be permanently fixed to the chip's facets. Moreover, the thermal phase shifters used here will also likely be replaced by alternative technologies based on the electro-optic effect.

Brighter single-photon sources will considerably reduce the measurement time.

References

- Thøgersen, L. & Olsen, J. A coupled cluster and full configuration interaction study of cn and cn^- . *Chem. Phys. Lett.* **393**, 36–43 (2004).
- Helgaker, T., Jørgensen, P. & Olsen, J. *Mol. Electronic Struct. Theory* (Wiley, Sussex, 2002).
- Saad, Y. *Numerical Methods for Large Eigenvalue Problems* Vol. 158 (SIAM, 1992).
- Page, L., Brin, S., Motwani, R. & Winograd, T. *The Pagerank Citation Ranking: Bringing Order to the Web*. Technical Report 1999-66 (Stanford InfoLab, 1999).
- Golub, G. H. & van der Vorst, H. A. Eigenvalue computation in the 20th century. *J. Comput. Appl. Math.* **123**, 35–65 (2000).
- Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information* (Cambridge University Press, 2000).
- Kitaev, A. Quantum measurements and the Abelian stabilizer problem. *Electronic Colloquium on Computational Complexity (ECCC)* **3** (1996).
- Griffiths, R. B. & Niu, C.-S. Semiclassical fourier transform for quantum computation. *Phys. Rev. Lett.* **76**, 3228–3231 (1996).
- Neven, H., Rose, G. & Macready, W. G. Image recognition with an adiabatic quantum computer I. Mapping to quadratic unconstrained binary optimization. Preprint at <http://arxiv.org/abs/0804.4457> (2008).
- Harrow, A., Hassidim, A. & Lloyd, S. Quantum algorithm for linear systems of equations. *Phys. Rev. Lett.* **103**, 150502 (2009).
- Berry, D. W. High-order quantum algorithm for solving linear differential equations. *J. Phys. A* **47**, 105301 (2014).
- Garnerone, S., Zanardi, P. & Lidar, D. A. Adiabatic quantum algorithm for search engine ranking. *Phys. Rev. Lett.* **108**, 230506 (2012).
- Abrams, D. S. & Lloyd, S. Simulation of many-body fermi systems on a universal quantum computer. *Phys. Rev. Lett.* **79**, 2586–2589 (1997).
- Abrams, D. S. & Lloyd, S. Quantum algorithm providing exponential speed increase for finding eigenvalues and eigenvectors. *Phys. Rev. Lett.* **83**, 5162–5165 (1999).
- Aspuru-Guzik, A., Dutoi, A. D., Love, P. J. & Head-Gordon, M. Simulated quantum computation of molecular energies. *Science* **309**, 1704–1707 (2005).
- Lanyon, B. P. *et al.* Towards quantum chemistry on a quantum computer. *Nat. Chem.* **2**, 106–111 (2010).
- Whitfield, J. D., Biamonte, J. & Aspuru-Guzik, A. Simulation of electronic structure hamiltonians using quantum computers. *Mol. Phys.* **109**, 735–750 (2011).
- Aspuru-Guzik, A. & Walther, P. Photonic quantum simulators. *Nat. Phys.* **8**, 285–291 (2012).
- Jones, N. C. *et al.* Faster quantum chemistry simulation on fault-tolerant quantum computers. *New J. Phys.* **14**, 115023 (2012).
- Lloyd, S. Computational capacity of the universe. *Phys. Rev. Lett.* **88**, 237901 (2002).
- Ma, X.-s., Dakic, B., Naylor, W., Zeilinger, A. & Walther, P. Quantum simulation of the wavefunction to probe frustrated Heisenberg spin systems. *Nat. Phys.* **7**, 399–405 (2011).
- Oseledets, I. Approximation of $2^d \times 2^d$ matrices using tensor decomposition. *SIAM J. Matrix Anal. A* **31**, 2130–2145 (2010).
- Ortiz, G., Gubernatis, J. E., Knill, E. & Laflamme, R. Quantum algorithms for fermionic simulations. *Phys. Rev. A* **64**, 022319 (2001).
- Liu, Y.-K., Christandl, M. & Verstraete, F. Quantum computational complexity of the n -representability problem: Qma complete. *Phys. Rev. Lett.* **98**, 110503 (2007).
- Temme, K., Osborne, T. J., Vollbrecht, K. G., Poulin, D. & Verstraete, F. Quantum Metropolis sampling. *Nature* **471**, 87–90 (2011).
- Yung, M.-H. & Aspuru-Guzik, A. A quantum-quantum Metropolis algorithm. *Proc. Natl Acad. Sci. USA* **109**, 754–759 (2012).
- Rayleigh, J. W. In finding the correction for the open end of an organ-pipe. *Phil. Trans.* **161**, 77 (1870).
- Ritz, W. Über eine neue Methode zur Lösung gewisser variationsprobleme der mathematischen physik. *J. Reine Angew. Math.* **135**, 1–61 (1908).
- Taube, A. G. & Bartlett, R. J. New perspectives on unitary coupled-cluster theory. *Int. J. Quant. Chem.* **106**, 3393–3401 (2006).
- O'Brien, J. L., Furusawa, A. & Vuckovic, J. Photonic quantum technologies. *Nat. Photon.* **3**, 687–695 (2009).
- Shadbolt, P. *et al.* Generating, manipulating and measuring entanglement and mixture with a reconfigurable photonic circuit. *Nat. Photon.* **6**, 45–49 (2011).
- Crawford, T. D. *et al.* Psi3: an open-source ab initio electronic structure package. *J. Comp. Chem.* **28**, 1610–1616 (2007).
- Szarek, S. J. Volume of separable states is super-doubly-exponentially small in the number of qubits. *Phys. Rev. A* **72**, 032304 (2005).
- Zhou, X.-Q. *et al.* Adding control to arbitrary unknown quantum operations. *Nat. Commun.* **2**, 413 (2011).
- Kempe, J., Kitaev, A. & Regev, O. The complexity of the local hamiltonian problem. *SIAM J. Comput.* **35**, 1070–1097 (2006).
- Nelder, J. A. & Mead, R. A simplex method for function minimization. *Comput. J.* **7**, 308–313 (1965).
- Conn, A. R., Gould, N. I. & Toint, P. L. *Trust Region Methods* Vol. 1 (Society for Industrial Mathematics, 1987).

Acknowledgements

We thank Scott Aaronson, Robert Chapman, Seth Lloyd, Tim Ralph, Terry Rudolph, Joe Fitzsimons and James Whitfield for discussions. We acknowledge financial support from the UK EPSRC, ERC, QUANTIP, PHORBITECH, QESSENCE, Nokia, NSQI, the Templeton Foundation and the EU DIQIP. A.P. acknowledges a Royal Academy of Engineering Research Fellowship and a ARC Discovery Early Career Researcher Award under project number DE140101700. J.M. is supported by the DOE Computational Science Graduate Fellowship under grant number DE-FG02-97ER25308. M.-H.Y. acknowledges the support by the National Basic Research Program of China Grants 2011CBA00300 and 2011CBA00301, the National Natural Science Foundation of China Grants 61033001 and 61361136003, and the Youth 1000-talent program. P.J.L. is supported by NSF award PHY-0955518 and by AFOSR award no. FA9550-12-1-0046. A.A.-G. acknowledges support from the NSF CCI award no. CHE-1037992, the Air Force Office of Scientific Research award no. FA9550-12-1-0046, the Camille and Henry Dreyfus foundation and the Alfred P. Sloan Foundation. J.L.O'B. acknowledges a Royal Society Wolfson Merit Award and a Royal Academy of Engineering Chair in Emerging Technologies.

Author contributions

All authors contributed extensively to the work presented in this paper.

Additional information

Supplementary information accompanies this paper at <http://www.nature.com/naturecommunications>

Competing financial interests: The authors declare no competing financial interests.

Reprints and permission information is available online at <http://npng.nature.com/reprintsandpermissions/>

How to cite this article: Peruzzo, A. *et al.* A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.* **5**:4213 doi: 10.1038/ncomms5213 (2014).



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 4.0 International License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>