

Experimental Bayesian Quantum Phase Estimation on a Silicon Photonic Chip

S. Paesani,¹ A. A. Gentile,¹ R. Santagati,¹ J. Wang,¹ N. Wiebe,^{2,*} D. P. Tew,³ J. L. O'Brien,¹ and M. G. Thompson^{1,†}
¹Quantum Engineering Technology Labs, H. H. Wills Physics Laboratory and Department of Electrical and Electronic Engineering,
 University of Bristol, BS8 1FD, United Kingdom

²Quantum Architectures and Computation Group, Microsoft Research, Redmond, Washington 98052, USA

³School of Chemistry, University of Bristol, Bristol BS8 1TS, United Kingdom

(Received 17 October 2016; revised manuscript received 13 January 2017; published 7 March 2017)

Quantum phase estimation is a fundamental subroutine in many quantum algorithms, including Shor's factorization algorithm and quantum simulation. However, so far results have cast doubt on its practicability for near-term, nonfault tolerant, quantum devices. Here we report experimental results demonstrating that this intuition need not be true. We implement a recently proposed adaptive Bayesian approach to quantum phase estimation and use it to simulate molecular energies on a silicon quantum photonic device. The approach is verified to be well suited for prethreshold quantum processors by investigating its superior robustness to noise and decoherence compared to the iterative phase estimation algorithm. This shows a promising route to unlock the power of quantum phase estimation much sooner than previously believed.

DOI: 10.1103/PhysRevLett.118.100503

Introduction.—Quantum algorithms promise exponential advantages over their classical counterparts, allowing the possibility to accomplish tasks otherwise unachievable on a classical computer [1,2]. A fundamental tool in quantum computing is the quantum phase estimation algorithm (PEA), necessary for harnessing many of its main applications, e.g factorization of large numbers [1,3–6] and simulation of molecular properties [7–11]. An efficient implementation of PEA will thus be a crucial subroutine for quantum computers. Kitaev's iterative phase estimation algorithm (IPEA) [12] and its adaptive version [13,14] have been employed in proof-of-principle implementations of PEA, as they solely rely on a relatively small number of qubits and logic gates [4–10]. However, they require exponentially long coherence of the quantum hardware and are very susceptible to experimental noise [9,14–16]. This means that conventional quantum phase estimation algorithms rapidly become impractical if the quantum computer is not fully error corrected, limiting the feasibility in near-term, pre-fault-tolerant quantum machines.

A new efficient Bayesian phase estimation algorithm, called *rejection filtering phase estimation* (RFPE), has been recently proposed to overcome this issue [16]. The algorithm applies an approximate form of Bayesian inference to efficiently estimate the correct eigenphase. Theoretical results suggest that RFPE has a number of desirable features: high robustness to noise, a well-motivated confidence interval for the estimated eigenphase, minuscule

memory requirement for the classical control, and speed-up over the standard IPEA.

The potential of RFPE to exhibit these properties has so far been suggested by numerical simulations of the algorithm [16]. Experimental evidence with realistic noise sources is now required to demonstrate the high performance of the approach, potentially vindicating RFPE as a scalable, practical, and quadratically faster alternative to other statistical methods [17,18].

Here we exploit a fully reconfigurable silicon quantum photonic device to investigate the experimental viability of RFPE. We compare and contrast the performance of both RFPE and IPEA under the action of different controllable experimental noises. Specifically, we find experimental evidence that RFPE is robust against realistic sources of errors, making it very appealing for near-term useful applications. Our results are made possible by our ability to generate and manipulate single photons states on-chip via an arbitrary controlled-unitary operation as well as the high level of precision and reconfigurability offered by the photonic chip.

Phase estimation.—The goal of phase estimation is, given a unitary \hat{U} and a quantum state $|\psi\rangle$, to learn an eigenvalue $e^{i\phi}$ of \hat{U} within the support of $|\psi\rangle$. Standard algorithms work by interfering paths in which either $\hat{1}$ or \hat{U}^j is applied to $|\psi\rangle$, for integer values of j , and then recombining the paths together to allow them to interfere [3]. Iterative phase estimation works by pooling the results of many such experiments and using a classical inference algorithm to estimate an eigenphase of \hat{U} from the resultant interference pattern.

IPEA allows eigenvalues of \hat{U} to be learned quadratically faster than by statistical sampling and requires exponentially fewer measurements. However, it typically requires long evolution times which can reduce its utility in pre-fault

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.




FIG. 1. (a) Quantum circuit for standard iterative and Bayesian phase estimation algorithms. (b) Experimental setup and the integrated silicon quantum photonic device. The quantum chip can perform any controlled- $U(2)$ operation and any single-qubit state preparation and analysis. Photons are produced and guided in the silicon waveguides (black wires) and reconfigurably controlled by thermo-optical phase shifters. Coherent light was used to generate photons and superconducting nanowire detectors were used for the detection, both coupled to the chip through lensed-single mode fibers. The implementation of the algorithms was achieved by interfacing the quantum device with a classical CPU.

tolerant hardware [12–14]. The most commonly used IPEA algorithm works by inferring each of the bits in a binary expansion of the eigenphase ϕ in reverse order [12–14]. The method uses the circuit in Fig. 1(a), where the measurement on the control qubit gives output 0 or 1 with probabilities $\cos^2(\pi M[\phi - \theta])$ and $\sin^2(\pi M[\phi - \theta])$, respectively. As each bit is learned iteratively, the algorithm applies a fixed policy for updating M and θ (see Supplemental Material [19]).

RFPE is in many ways simpler. This Bayesian approach uses a Gaussian probability distribution $P(\phi)$ (the prior) representing the confidence that the current hypotheses is the correct eigenphase. The result of each new measurement is used to update the mean μ and standard deviation σ according to Bayes' theorem, approximated with rejection sampling (see Supplemental Material [19]). Specifically, a host of particles are drawn from the prior distribution and then probabilistically discarded based on the likelihood function. The remaining samples model the posterior probability distribution, which becomes the new prior. Since the number of surviving particles decreases exponentially, the posterior distribution is refit to a new Gaussian at each step and fresh particles added, drawn from this new distribution. Rather than learning each of the bits of ϕ individually, the RFPE algorithm gains information about every bit simultaneously.

In more detail, if an outcome $E \in \{0, 1\}$ is obtained from an experiment with parameters M and θ , the likelihood function for the host of particles $\{x_i\}$ is calculated:

$$P(E|x_i; M, \theta)/\kappa = \begin{cases} \cos^2[M\pi(x_i - \theta)]/\kappa & E = 0 \\ \sin^2[M\pi(x_i - \theta)]/\kappa & E = 1, \end{cases}$$

where κ is a rescaling constant [16]. Particles from the prior are discarded probabilistically based on this function. In order to maximize the information gain at each step, new values of M and θ can be extracted from the prior distribution using various optimization methods. A near-optimal choice is provided by particle guess heuristics [22,23], giving $\theta \sim P(\phi)$ and $M = \lceil 1.25/\sigma \rceil$. If the likelihood function fails to describe the data (due to experimental noise) then the algorithm estimates the best model for the experimental results within the assumptions, and gives a firm estimate of the uncertainty in the eigenphase. While in standard IPEA any error that occurs in inferring a bit cannot be corrected in subsequent algorithm steps, RFPE does not suffer this issue because it does not try to infer the bits sequentially. While these expectations have been born out in simulation [16], they have not been verified in practice. We provide such verification below using an integrated quantum photonic device. Further details for RFPE and IPEA can be found in Supplemental Material [19].

Integrated quantum photonic device.—Silicon quantum photonics have emerged as a promising approach for the realization of quantum hardware, since, in principle, all the necessary photonic components (sources, circuits, filters, and detectors) for quantum information processing can be integrated on a single platform [24]. We developed a quantum photonic device in silicon waveguides for the experimental implementation of RFPE, shown in Fig. 1(b). The chip was manufactured on a standard silicon-on-insulator platform and is capable of performing arbitrary two-qubit controlled unitary operations. Using spontaneous four-wave mixing, photon pairs were created in two spiral sources pumped with ≈ 20 mW bright light near 1550 nm

wavelength [25]. Separating the photons by multimode interferometer beam splitters (BS) and swapping the two modes by a waveguide crosser yield a maximally entangled two-photon postselected state $(|0\rangle_C |0\rangle_T + |1\rangle_C |1\rangle_T)/\sqrt{2}$, where $|0\rangle$ ($|1\rangle$) indicates the photon state either in its upper or lower spatial mode [26,27], whereas $C(T)$ subscript refers to the control (target) register. Two additional spatial modes are then added to the target register, obtaining a path encoded qubit for each of the two modes of the target wave function. After an initial state preparation $|\phi\rangle_T$, each qubit is manipulated with a separate transformation, depending on which path the photon is traveling on the identity \hat{I} for the upper modes (the ones corresponding to a $|0\rangle_C$) and an arbitrary unitary \hat{V} for the bottom modes (the ones corresponding to a $|1\rangle_C$). The path information is erased by two BSs and the state is finally postselected obtaining the equivalent photon count statistics of an arbitrary control-unitary operation $(|0\rangle_C \otimes |\phi\rangle_T + |1\rangle_C \otimes \hat{V} |\phi\rangle_T)/\sqrt{2}$ [10,28]. The quantum logic of Fig. 1(a) is completed by performing a single qubit operation on the control photon. All the processes required for state preparation, manipulation, and measurement are achieved through thermo-optical phase shifters and Mach-Zehnder interferometers, as shown in Fig. 1(b) [27,29]. Controlling the electric power supplied to the phase shifters, each phase φ can be driven with an average precision of ≈ 0.01 rad (see Supplemental Material [19]). Finally, photons were detected by superconducting nanowire single photon detectors (SNSPD) and coincidence counts obtained by a time interval analyzer. The photon statistics were used to measure the projectors $\Pi(0, 1)$ on the computational basis of the photonic qubits. The Bayesian update and changes to the controls of the quantum device required by RFPE can be calculated and fed to the quantum system using an interfaced classical computer.

Experimental results.—The rapid reconfigurability and the high precision of the silicon photonic device were crucial to the practical implementation and testing of RFPE. As is usual in photonic experiments, where measurements provide probability distributions rather than single-shot data, the value of E was determined by the majority voting method on the projective measurement statistics for both RFPE and IPEA (see Supplemental Material [19]). For the stochastic representation of the prior, 1000 particles was found to be reliable. The initial prior distribution was set to a Gaussian $\mathcal{N}(\pi, \pi^2)$, which approximates a uniform prior.

In Fig. 2(a) we report the results from a single RFPE run, which demonstrates that the estimation converges exponentially to the true eigenphase. For this case an error as low as 2.4×10^{-4} rad is achieved within 50 experiments on the quantum device. This error is in good agreement with the standard deviation of the final posterior Gaussian distribution [$\approx 4.2 \times 10^{-4}$ rad after 50 experimental steps, see the inset of Fig. 2(a)], thus confirming that the algorithm provides a reliable uncertainty estimate. We

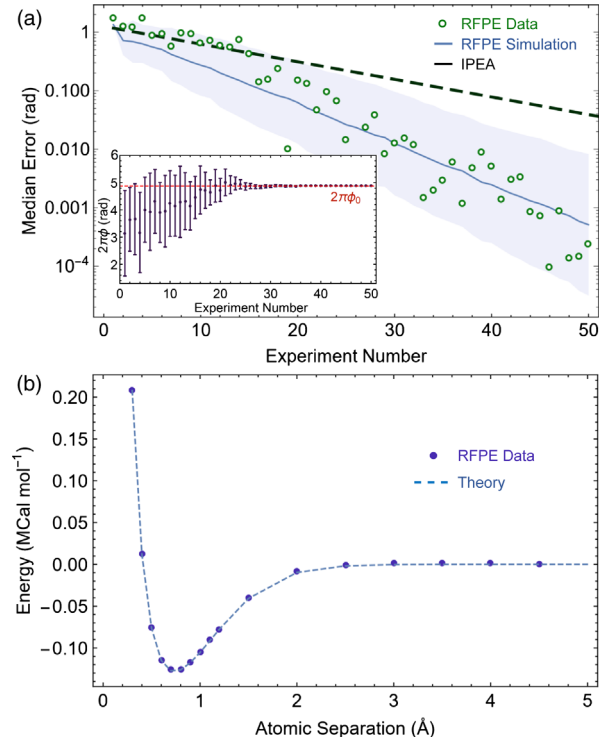


FIG. 2. (a) Convergence of RFPE to the true phase value $2\pi\phi_0 = 4.8741$ rad related to the energy of the dissociated H_2 molecule. The initial prior distribution is $\mathcal{N}(\pi, \pi^2)$. Data points show an exponential shrink of the error, compatible with simulations (blue line) of the device performance averaged over 1000 runs of the RFPE algorithm (shaded area: 67.5% credible interval). The dashed black line denotes the convergence of the standard IPEA using the parameters discussed in the main text. Inset: convergence of the phase estimation to ϕ_0 (red line), where errors are evaluated using the s.d. of the prior distribution. Error bars are obtained from the standard deviation of the median. (b) Bonding energies of the H_2 molecule for various atomic distances using RFPE with 50 steps. Energy estimations are achieved within chemical accuracy. Errors are smaller than the markers and neglected in the plot for more clarity. The dashed line represents theoretical energies.

remark that this feature is not present in other phase estimation techniques, which do not provide a rigorous estimation of the phase uncertainty that arises in non-fault-tolerant devices [16]. A possible strategy for obtaining such uncertainty estimates from IPEA would be to determine the mean and the standard deviation of the measured eigenphase from repeated experiments. Figure 2(a) shows that when adopting this strategy with a reasonable cardinality for the experiments, 10 for the curve reported, RFPE outperforms IPEA.

To extensively test how RFPE performs in key applications, we scanned the energy of molecular H_2 for different atomic separations, using 50 iterations of RFPE for each eigenphase evaluation. The eigenstates of the molecular Hamiltonian are mapped into a the qubit basis using the Jordan-Wigner transformation and the eigenphases are

directly related to the binding energy [7,9–11]. As shown in Fig. 2(b), the estimated energies match the theoretical values with high precision. The average error for the data set is 0.72 kCal/mol, providing a precision higher than chemical accuracy (≈ 1 kCal/mol).

The main advantage of the Bayesian approach over traditional methods comes from its expected reliability on non-fault-tolerant devices. Here we investigate experimentally the robustness of the protocol against two main controllable sources of noises: *gate errors* and *decoherence*.

The infidelity of unitary operations is a well-known problem existing in quantum hardware, and is typically given by a noisy control and imperfect manufacture and calibration of the components. On integrated quantum photonic devices it is mainly due to control noise and residual crosstalk of the phase gates, which are expected to occur on the device in Fig. 1 [10]. Our electronic phase-shifter’s driver allows us to precisely control the thermo-optical phase gates as residual thermal cross talk can be compensated by further calibration (see Supplemental Material [19]). The high controllability allows us to add a tunable level of noise on the phase gates to study the robustness of RFPE. This is achieved by replacing the correct phases $\bar{\varphi}$ required to implement the unitary transformation with synthetic values φ sampled from a Gaussian distribution $\varphi \sim \mathcal{N}(\bar{\varphi}, \sigma_{\text{phase}})$. The parameter σ_{phase} mimics a Gaussian noise in the phases, which in turn introduces a controllable noise in the fidelity of both the implemented state preparation and the unitary evolution. Figure 3(a) shows the convergence of both RFPE and IPEA scanning σ_{phase} up to 0.55 rad, corresponding to average state fidelity 94% and gate fidelity 91% expected in the chip (see

Supplemental Material [19]). We report the performance of RFPE with 100 steps, compared to a 16-bit IPEA, averaged over 10 runs to estimate the error bars, i.e., 160 total experiments per data point. We remark that since a reasonable error bar estimation requires a higher number of experiments for IPEA than RFPE, the relative rates of convergence are not immediately obvious from these figures. Rather, these plots illustrate how the performance of each algorithm is affected by increasing errors, to compare their robustness to noise. For $\sigma_{\text{phase}} \geq 0.05$ rad IPEA dramatically decreases in accuracy and becomes quickly unreliable. This occurs because while the majority voting scheme provides error resilience for small error rates, it can diverge rapidly once the error rate crosses a threshold (see Supplemental Material [19]).

On the other hand, in this regime the performance of RFPE is initially only slightly affected, maintaining a very high level of precision even when IPEA fails. This is expected because RFPE does not make hard decisions about bits as each experiment yields information about all bits at once. Thus errors are unlikely to be critical. In order for RFPE to be substantially affected we require σ_{phase} higher than 0.3 rad, a value much higher than the actual experimental noise in our device.

Decoherence is an important limitation in many quantum computing experiments but it plays a minor role in integrated quantum photonics. It must then be introduced artificially in our experiment. In order to simulate it, coincidence counts provided by the SNSPDs for the $\Pi(0)$ and $\Pi(1)$ projectors were progressively flattened out by classical postprocessing and combined with Poissonian noise in the measurements. In agreement with the depolarizing noise model we mimicked the presence of

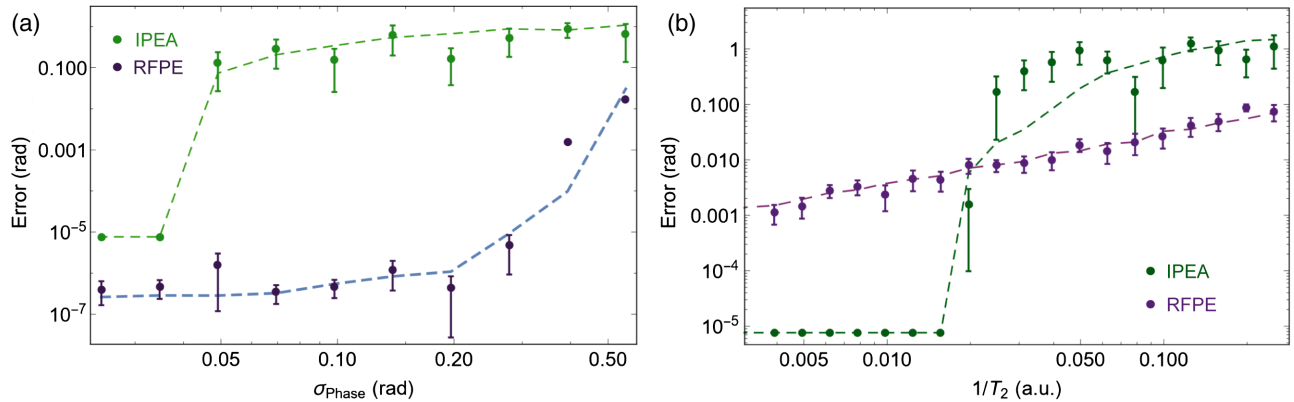


FIG. 3. Effects of different experimental noises on phase estimation strategies. (a) Infidelity of quantum operation. Each of the correct phases $\bar{\varphi}_i$ for the phase gates is synthetically replaced with a Gaussian distributed $\varphi_i \sim \mathcal{N}(\bar{\varphi}_i, \sigma_{\text{phase}})$, where σ_{phase} represents a noise in the phases. (b) Decoherence. For IPEA data, experiments were repeated 10 times with 16-bit accuracy to evaluate median error and error bars, while the RFPE data were collected from a single run, after 100 measurements, and directly used to evaluate the error and uncertainty within the algorithm. Error bars for the estimated phase represent in both plots either a 67.5% credible region for RFPE, either a 67.5% confidence interval for IPEA. In the cases where error bars are smaller than the markers they have been omitted for clarity. Points are experimental data and dashed lines are simulations averaged over 1000 runs. The simulations take into account the characterized residual phase noise in the device.




FIG. 4. Experimental data (points) and simulation (lines) of the behavior of RFPE under different decoherence times T_2 . Shaded areas represent a 67.5% credible interval from data simulated averaging over 500 runs of RFPE.

a normalized decoherence time T_2 (see Supplemental Material [19]) such that

$$P'(D|\phi) = e^{-M/T_2}P(D|\phi) + \frac{1 - e^{-M/T_2}}{2},$$

where $P(D|\phi)$ is the data obtained from the photon coincidence counts for the outcome $D \in \{0, 1\}$. We introduce this noise model while processing online the output data during the iterative process, thus affecting the choice of the experiments. In this way it is possible to simulate the behavior of RFPE and IPEA in systems that are prone to this model of decoherence.

We studied the action of the depolarizing noise up to $T_2 = 4$ for both IPEA and RFPE, as shown in Fig. 3(b). The performance of IPEA has a substantial and sharp deterioration at T_2 close to 32, whereas the median error of 100-step RFPE decreases only polynomially with $1/T_2$, maintaining an error $\mathcal{O}(10^{-2})$ even in the regime where conventional IPEA fails to provide any reliable estimate of the phase. In the presence of characterized depolarizing noise an optimized value for M is given by $\min(\lceil 1.25/\sigma \rceil, T_2)$ [16], which, however, implies that, when decoherence is significant, the performance of RFPE degrades significantly. This behavior is exhibited by the experimental data in Fig. 4, where the convergence of RFPE is reported under the action of various T_2 . We observe that RFPE ceases to learn exponentially quickly when $1/\sigma \sim T_2$, after which the algorithm continues to learn at a polynomial rate, unlike IPEA [16].

Conclusion.—Our work shows how the precision and controllability developed in quantum technologies, here in particular integrated photonics, allows us to go beyond the basic proof-of-principle demonstrations of quantum algorithms and to enter a regime where they can be extensively tested and compared. We experimentally verified the Bayesian phase estimation algorithm on a fully programmable silicon quantum photonic device and demonstrated its superior performance in presence of noise. Although in

this work the experiment is performed using a small-scale unitary and a photonic device, more complex future implementation can be efficiently performed on any scalable quantum architecture. The Bayesian approach remarkably lowers the requirements for the implementation of quantum phase estimation on pre-fault-tolerant devices, showing a new route for practical and useful quantum information processing in the near future. We remark that the implications of the Bayesian approach to phase estimation are not limited to photonic implementations: all other quantum information processing platforms, e.g., superconducting qubits and ions traps, can benefit as well from its enhanced performance and noise resilience.

We thank A. Laing, J. W. Silverstone, D. Bonneau, C. E. Granade, and X. Zhou for insightful discussion. We thank K. Ohira, N. Suzuki, H. Yoshida, N. Iizuka, and M. Ezaki for the device fabrication. We acknowledge the support from the Engineering and Physical Sciences Research Council (EPSRC) (Grants No. UF130574, EP/J017175/1, EP/K021931/1, EP/K033085/1), European Research Council (ERC) (Grants No. #608062, No. #641039), Photonic Integrated Compound Quantum Encoding (PICQUE), FP7 Action: Beyond the Barriers of Optical Integration (BBOI) (Grant No. FP7-ICT-2013-C), Quantum Simulation on a Photonic Chip (QuChip), U. S. Army Research Office (ARO), and Centre for Nanoscience and Quantum Information (NSQI). D. P. T thanks the Royal Society for a University Research Fellowship. J. L. O'B. acknowledges a Royal Society Wolfson Merit Award and a Royal Academy of Engineering Chair in Emerging Technologies. M. G. T. acknowledges the support from an EPSRC Early Career Fellowship and the Toshiba Research Fellowship scheme.

S. P. and A. A. G. contributed equally to this work.

*nawiebe@microsoft.com

†mark.thompson@bristol.ac.uk

- [1] P. Shor, Algorithms for quantum computation: Discrete logarithms and factoring, in *Proceedings of the 35th Annual Symposium on Foundations of Computer Science* (IEEE Computer Society Press, New York, 1994).
- [2] S. Lloyd, Universal quantum simulators, *Science* **273**, 1073 (1996).
- [3] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, England, 2001).
- [4] B. P. Lanyon, T. J. Weinhold, N. K. Langford, M. Barbieri, D. F. V. James, A. Gilchrist, and A. G. White, Experimental Demonstration of a Compiled Version of Shor's Algorithm with Quantum Entanglement, *Phys. Rev. Lett.* **99**, 250505 (2007).
- [5] E. Martin-Lopez, A. Laing, T. Lawson, R. Alvarez, X.-Q. Zhou, and J. L. O'Brien, Experimental realization of Shor's quantum factoring algorithm using qubit recycling, *Nat. Photonics* **6**, 773 (2012).

- [6] T. Monz, D. Nigg, E. A. Martinez, M. F. Brandl, P. Schindler, R. Rines, S. X. Wang, I. L. Chuang, and R. Blatt, Realization of a scalable Shor algorithm, *Science* **351**, 1068 (2016).
- [7] B. P. Lanyon *et al.*, Towards quantum chemistry on a quantum computer, *Nat. Chem.* **2**, 106 (2010).
- [8] J. Du, N. Xu, X. Peng, P. Wang, S. Wu, and D. Lu, NMR Implementation of a Molecular Hydrogen Quantum Simulation with Adiabatic State Preparation, *Phys. Rev. Lett.* **104**, 030502 (2010).
- [9] P. J. J. O'Malley, R. Babbush *et al.*, Scalable Quantum Simulation of Molecular Energies, *Phys. Rev. X* **6**, 031007 (2016).
- [10] R. Santagati, J. Wang *et al.*, Quantum simulation of Hamiltonian spectra on a silicon chip, [arXiv:1611.03511v3](https://arxiv.org/abs/1611.03511v3).
- [11] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, Simulated quantum computation of molecular energies, *Science* **309**, 1704 (2005).
- [12] A. Kitaev, Quantum measurements and the Abelian Stabilizer Problem, *Electron. Colloq. Comput. Complex.* **3** (1996).
- [13] R. B. Griffiths and C.-S. Niu, Semiclassical Fourier Transform for Quantum Computation, *Phys. Rev. Lett.* **76**, 3228 (1996).
- [14] M. Dobšíček, G. Johansson, V. Shumeiko, and G. Wendin, Arbitrary accuracy iterative quantum phase estimation algorithm using a single ancillary qubit: A two-qubit benchmark, *Phys. Rev. A* **76**, 030306 (2007).
- [15] M. A. Rubin and S. Kaushik, Loss-induced limits to phase measurement precision with maximally entangled states, *Phys. Rev. A* **75**, 053805 (2007).
- [16] N. Wiebe and C. Granade, Efficient Bayesian Phase Estimation, *Phys. Rev. Lett.* **117**, 010503 (2016).
- [17] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5**, 4213 (2014).
- [18] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, *New J. Phys.* **18**, 023023 (2016).
- [19] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.118.100503> for additional experimental details and further theoretical discussion on the implemented algorithms, which includes Refs. [20,21].
- [20] A. Politi, M. J. Cryan, J. G. Rarity, S. Yu, and J. L. O'Brien, Silica-on-silicon waveguide quantum circuits, *Science* **320**, 646 (2008).
- [21] M. Veldhorst *et al.*, A two-qubit logic gate in silicon, *Nature (London)* **526**, 410 (2015).
- [22] C. Ferrie, C. E. Granade, and D. G. Cory, How to best sample a periodic probability distribution, or on the accuracy of Hamiltonian finding strategies?, *Quantum Inf. Process.* **12**, 611 (2013).
- [23] N. Wiebe, C. Granade, C. Ferrie, and D. G. Cory, Hamiltonian Learning and Certification Using Quantum Resources, *Phys. Rev. Lett.* **112**, 190501 (2014).
- [24] J. Silverstone, D. Bonneau, J. L. O'Brien, and M. G. Thompson, Silicon Quantum Photonics, *IEEE J. Sel. Top. Quantum Electron.* **22**, 390 (2016).
- [25] M. D. Reid and D. F. Walls, Quantum theory of nondegenerate four-wave mixing, *Phys. Rev. A* **34**, 4929 (1986).
- [26] J. Wang *et al.*, Chip-to-chip quantum photonic interconnect by path-polarization interconversion, *Optica* **3**, 407 (2016).
- [27] J. W. Silverstone, R. Santagati, D. Bonneau, M. J. Strain, M. Sorel, J. L. O'Brien, and M. G. Thompson, Qubit entanglement between ring-resonator photon-pair sources on a silicon chip, *Nat. Commun.* **6**, 7948 (2015).
- [28] X. Zhou, P. Kalasuwan, T. C. Ralph, and J. L. O'Brien, Calculating unknown eigenvalues with a quantum algorithm, *Nat. Photonics* **7**, 223 (2013).
- [29] G. T. Reed, G. Mashanovich, F. Y. Gardes, and D. J. Thomson, Silicon optical modulators, *Nat. Photonics* **4**, 518 (2010).