

Quantum Science and Technology



PAPER

A quantum-classical cloud platform optimized for variational hybrid algorithms

RECEIVED
4 December 2019

REVISED
9 February 2020

ACCEPTED FOR PUBLICATION
12 February 2020

PUBLISHED
26 March 2020

Peter J Karalekas¹, Nikolas A Tezak¹, Eric C Peterson¹, Colm A Ryan¹, Marcus P da Silva² and Robert S Smith¹

Rigetti Computing, 2919 Seventh Street, Berkeley, CA 94710 United States of America

¹ Current address: OpenAI, 3180 18th St, San Francisco, CA 94110 United States of America.

² Current address: Microsoft Quantum, One Microsoft Way, Redmond, WA 98052 United States of America.

E-mail: peter@karalekas.com

Keywords: cloud-based quantum computing, near-term quantum algorithms, quantum software engineering

Supplementary material for this article is available [online](#)

Abstract

In order to support near-term applications of quantum computing, a new compute paradigm has emerged—the quantum-classical cloud—in which quantum computers (QPUs) work in tandem with classical computers (CPUs) via a shared cloud infrastructure. In this work, we enumerate the architectural requirements of a quantum-classical cloud platform, and present a framework for benchmarking its runtime performance. In addition, we walk through two platform-level enhancements, parametric compilation and active qubit reset, that specifically optimize a quantum-classical architecture to support variational hybrid algorithms, the most promising applications of near-term quantum hardware. Finally, we show that integrating these two features into the Rigetti Quantum Cloud Services platform results in considerable improvements to the latencies that govern algorithm runtime.

1. Introduction

The first experimental realizations of quantum algorithms date back to over a decade ago [1–4], but in the last three years quantum computing has rapidly transitioned from a field of scientific research to a full-fledged technology industry. The recent demonstration of *quantum supremacy* over classical computing [5] is a considerable milestone, but there is still much progress to be made on the road to solving real-world problems with quantum computers and achieving *quantum advantage*. Improving the error rates of quantum devices [6, 7] and ultimately reaching the regime of *fault tolerance* [8] is necessary for unlocking the most powerful known applications of quantum computers. At the same time, the industry has increased its focus on finding ways to solve valuable problems using the noisy intermediate-scale quantum (NISQ) processors that are currently available [9].

The desire to provide the research community with access to scarce quantum hardware in order to shorten the path to quantum advantage resulted in the development of a new compute architecture—the *quantum cloud*. As part of this architecture, the concept of an Internet-accessible data center has been extended to include quantum devices. Infrastructure for the quantum cloud requires a slew of new specialized hardware, for example, dilution refrigerators to house superconducting qubits and racks of microwave instruments to control them. To build the quantum cloud, some developers of quantum computers have pivoted to being *full-stack*, using in-house infrastructure to offer cloud-based access to their quantum devices (e.g. IBM Quantum Experience). In addition, some traditional cloud providers have begun to add quantum backends through strategic hardware-software partnerships.

The first iteration of quantum cloud offerings employed a hybrid cloud model [10], in which users of the service submitted quantum programs using a web API to a queue hosted by the public cloud (e.g. Amazon Web Services). Then, a server colocated with a quantum processor would periodically pull jobs off of the queue, execute them, and return results back to the user. This approach was effective in offering worldwide, public

access to quantum resources, but suffers in terms of runtime efficiency due to the overhead of network connections and the use of a shared queue. In addition, the traditional web API model fails to capitalize on or adapt to any properties specific to using a quantum device for computation.

In particular, the most promising approach to effectively using near-term quantum devices is through *variational hybrid algorithms* (VHAs) [11] which employ a quantum-classical architecture, essentially leveraging the quantum computer as a co-processor alongside a powerful classical computer. These algorithms have been applied to areas such as combinatorial optimization [12, 13], quantum chemistry [14, 15], and machine learning [16], and numerous proposals for applications of the variational method continue to arise with increasing frequency [17, 18]. However, VHAs require a tight coupling between quantum and classical resources, and using a cloud-hosted queue is slow with respect to the scale of quantum operations, especially on a superconducting device [19]. In addition, a quantum cloud architecture must be specifically optimized in order to efficiently support the variational model of execution. In this work, we investigate architectural bottlenecks of this new *quantum-classical cloud*, and provide a benchmarking framework to analyze its runtime performance. We then use the benchmark to quantify the dramatic reduction in latency achieved by the Rigetti Quantum Cloud Services (QCS) platform via the implementation of specialized techniques for quantum program compilation and qubit register reset.

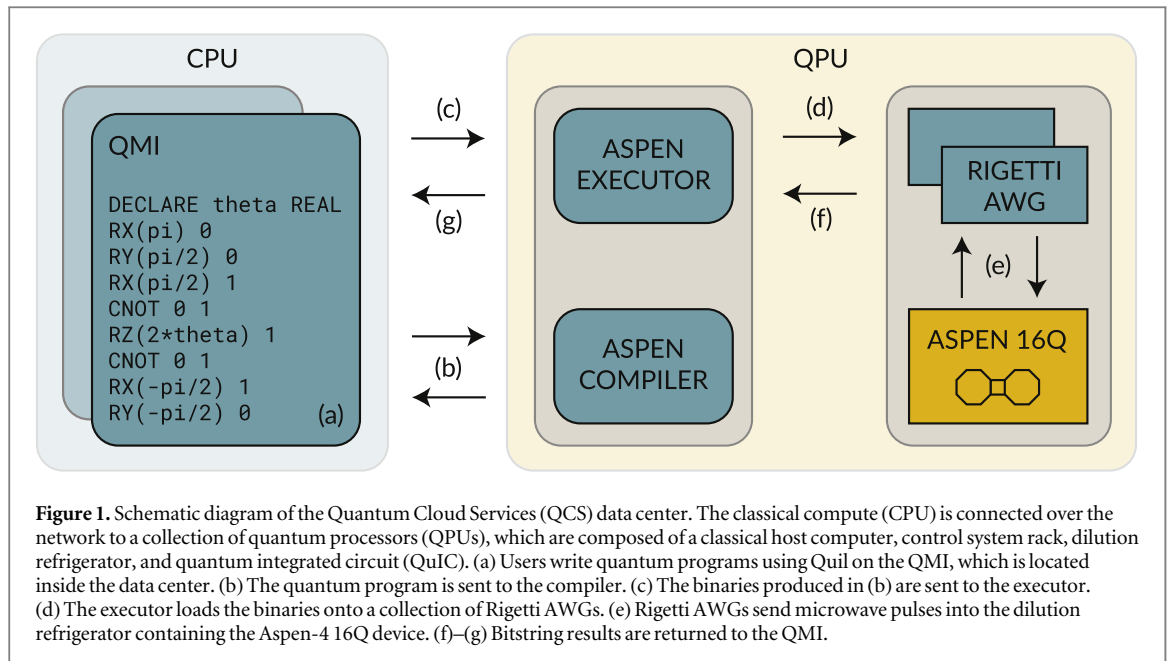
2. Runtime bottlenecks in the quantum-classical cloud

The job of a quantum cloud platform is to ingest programs written in a backend-independent high-level quantum programming language [20, 21], compile them into a platform-specific representation, run them on an available quantum device, and return the results to the user. Specifically, a quantum cloud platform has four essential components:

- (i) An apparatus that houses the physical objects that act as qubits (e.g. an optical table and trapping system for ions or neutral atoms).
- (ii) A *control system* containing instruments for manipulating that apparatus in order to drive the desired evolution and read out qubit measurement results.
- (iii) An *executor* that orchestrates the control system to run quantum programs and return measurement results to the user.
- (iv) A *compiler* that takes in quantum programs and produces instrument binaries for the executor.

To be categorized as quantum-classical cloud, a platform must also include access to classical compute resources. Depending on the particular qubit implementation used by the platform, the CPU-QPU interaction could become the largest bottleneck in the variational model of execution. For example, when using superconducting qubits (with gate times in the tens of nanoseconds) the CPU and QPU should be physically colocated in order to enable a low-latency link between the user and the quantum device. Although colocating user and compute is not a new concept in cloud computing in general, it has yet to take hold broadly in quantum computing, and drastically reduces overhead in VHAs. For Rigetti QCS, which uses superconducting qubits as the backend, users interact with the QPU via a preconfigured development environment called the quantum machine image (QMI) (figure 1(a)). The QMI is a virtual machine running on a classical compute cluster located inside the Rigetti *quantum data center* in Berkeley, CA, and contains the Forest SDK for building applications using the quantum instruction language Quil [21]. Once written, quantum programs are sent for compilation into pulse-level instructions (figure 1(b)). The information that encodes this gate-to-pulse mapping is contained within a calibration database, and is updated whenever the system drifts out of specification. The binaries that are returned by the compiler are then sent to the executor (figure 1(c)), which loads them onto a collection of Rigetti-custom microwave arbitrary waveform generators (AWGs) and receivers, and triggers the instruments to begin execution (figure 1(d)). The Rigetti AWGs then send microwave pulses into a dilution refrigerator to manipulate and read out the state of the Aspen-4 16-qubit quantum processor (figure 1(e)). Finally, the bitstrings resulting from readout are returned to the user's compute environment (figures 1(f)–(g)) for processing and analysis.

In the variational model of execution, one seeks to minimize an objective function that is expensive to compute on classical hardware, by embedding it as a subroutine on a quantum computer. To begin, a classical computer takes an initial guess and instructs a quantum computer to perform a predetermined computation. Then, from the output statistics of sampling the quantum program many times, a classical optimizer running on the CPU updates the QPU instructions for the next round of iteration. Depending on the difficulty of the



problem, and the quality of the QPU, this will repeat many times before finding a potential solution [22]. Thus, the structure of a variational hybrid algorithm is broken into two nested loops: an outer *variational iteration loop*, and an inner *quantum execution loop*. We denote each roundtrip completion of the variational iteration loop as an *iteration step*, and each roundtrip completion of the quantum execution loop as a *shot*. Each iteration step includes communicating with the QPU, waiting for it to return results, and using an optimizer to decide what to run on the QPU next. Each shot is defined by running a single program of quantum instructions and getting back a single bitstring of measurement outcomes. This structure gives us a way to think about the different components that contribute to hybrid algorithm runtime—they either occur once per step in the variational iteration loop, or once per shot in the quantum execution loop. In order to improve the performance of QCS and optimize it for the variational model of execution, we began by constructing a latency budget for each of the nested loops (table 1).

The components of the per-step latency budget are instrument initialization, network communication, and compilation, which is the largest contributor by at least one order of magnitude. The contribution from instrument initialization has already been substantially reduced by using custom control hardware. In addition, the runtime of a single shot is split between gates (for manipulating and entangling qubit states), readout (for measuring and extracting bitstring outcomes), and qubit reset. Qubit reset can be performed passively, by waiting for all qubits to relax to their ground states. However, this relaxation process is the same decoherence process that determines T_1 , one of the metrics for qubit lifetime. Therefore, this *passive qubit reset* time will only increase as qubit performance improves [24, 25], and it already dominates the per-shot latency budget. Thus, from examining the two latency budgets, compilation and qubit reset are the areas to which improvements would have the largest impact on algorithm runtime.

3. Optimizing for the variational execution model

Having identified compilation and qubit reset as potential bottlenecks for a quantum-classical cloud architecture, we walk through specific enhancements made to mitigate their contributions to the latency budgets of Rigetti's quantum cloud platform.

3.1. Parametric compilation

The underlying quantum program for a variational algorithm is parametric, meaning that from iteration step to iteration step, the sequence of instructions is static and only the instruction arguments change. Thus, a specialized compiler that preserves this *parametric ansatz* structure can leverage it to improve runtime. To build such a compiler, we need to understand the physical implementation of a high-level quantum program in order to determine which instructions are easy to modify parametrically at runtime. For superconducting quantum processors (and many other QPU implementations), qubits are controlled by shaped radio frequency (RF) or microwave pulses typically generated by an AWG. The pulse amplitude, duration, and phase control the rotation angle and axis effected on the qubits. Typically the phase of the pulse is determined by an abstract rotating

Table 1. Per-step and per-shot latency budgets for the Aspen-4 QPU via QCS, rounded to one significant figure. The per-step budget was collected by running Max-Cut QAOA [12] on two qubits, and measuring the completion time for each of the components. The per-shot budget contains the average gate and readout times for Aspen-4 from the calibration database. The reset time is five times the longest T_1 time measured on Aspen-4 (see figure 3 for details). The compilation task is further broken down into two components: the first which converts arbitrary quantum programs into ones that use the gateset and topology of the target quantum device (about 150 ms) [23]; and the second which converts this *nativized* program into pulse-level instrument binaries (about 50 ms).

Per-step latency budget	
Task	Time
Compilation	200 ms
AWG load and arm	8 ms
AWG trigger	10 ms
Network comms	5 ms
Per-shot latency budget	
Task	Time
Single-qubit gates	60 ns
Two-qubit gates	300 ns
Readout and capture	2 μ s
Passive reset	100 μ s

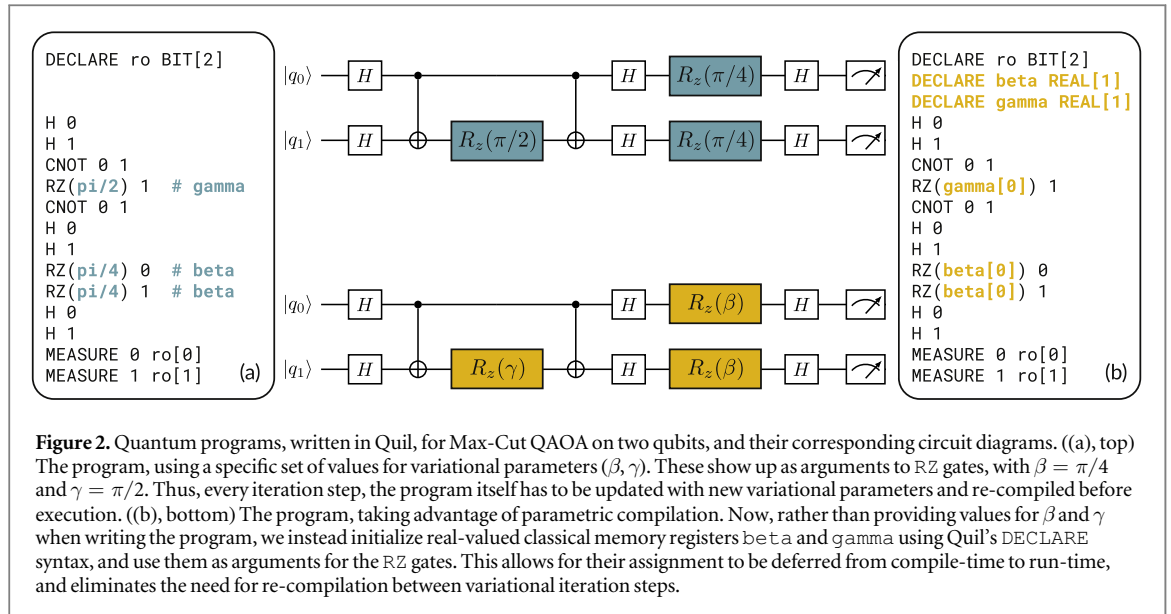
reference frame set to a frequency determined by the physics of the qubits. Z -rotations of a qubit then correspond to instantaneous reference frame update events that change the phase of subsequent microwave pulses. Without loss of generality, we can directly relate each parameter of a variational algorithm to Z -rotations or phase updates for one or more qubit reference frames. Previously, the arguments to these `SHIFT-PHASE` [26] operations were provided at compile-time, and hardcoded in the instrument binary. Thus, for every step in a VHA, new instrument binaries would need to be compiled, such that these arguments could be updated.

To circumvent this need for re-compiling, we implemented a feature called *parametric compilation*. Like a standard executable file, each instrument binary includes a header and a collection of sections [27]. The *instruction memory* section contains executable code, and the *waveform memory* section contains pulse shapes that are referenced by the instruction memory. Rather than having `SHIFT-PHASE` instructions use a static data field as an argument, each instrument binary now additionally has a *data memory* section which instructions can access by reference. Like classical shared memory, the data memory section can also be updated at runtime by an external process in order to change the effect of the binary [28]. The data memory section layout of a particular binary is prescribed by Quil's `DECLARE` syntax (figure 2), which can be used to initialize named memory registers of various data types (`BIT`, `OCTET`, `INTEGER`, and `REAL`). Once initialized, the memory registers can be provided as arguments to parameterized gates such as `RZ`. When these gates are compiled (figure 1(b)) they become `SHIFT-PHASE` operations that reference entries in an empty data memory section. Upon execution, the user provides these *parametric binaries* along with a map of memory register assignments, for example, variational parameters for the current iteration step of a VHA (figure 1(c)). Then, the executor fills in the data memory section using the memory map to produce *patched binaries* (figure 1(d)). For each step in a VHA, these binaries can be re-patched with a new set of variational parameters, and thus the variational iteration loop can repeat until termination without the need for compilation. For details on how to implement a variational algorithm using parametric compilation, see appendix A.

3.2. Active qubit reset

Rather than waiting for qubits to passively reset, we can implement a protocol called *active qubit reset* that sets all qubits to their ground states at the beginning of a computation. This has been previously demonstrated by controllably transferring qubit excited state population to a system with a much faster decay rate (such as the readout resonators) [5, 29, 30], or by performing a measurement and then quickly feeding back to conditionally apply an X gate [31, 32]. But, for this protocol to be useful, the conditional control flow needs to happen on a timescale comparable to the quantum gates, and it cannot introduce significant error into proceeding computations.

We chose to implement this classical feedback loop on our platform because it additionally unlocks more complex feedback/feedforward circuits that take advantage of mid-circuit measurements. To support this, the



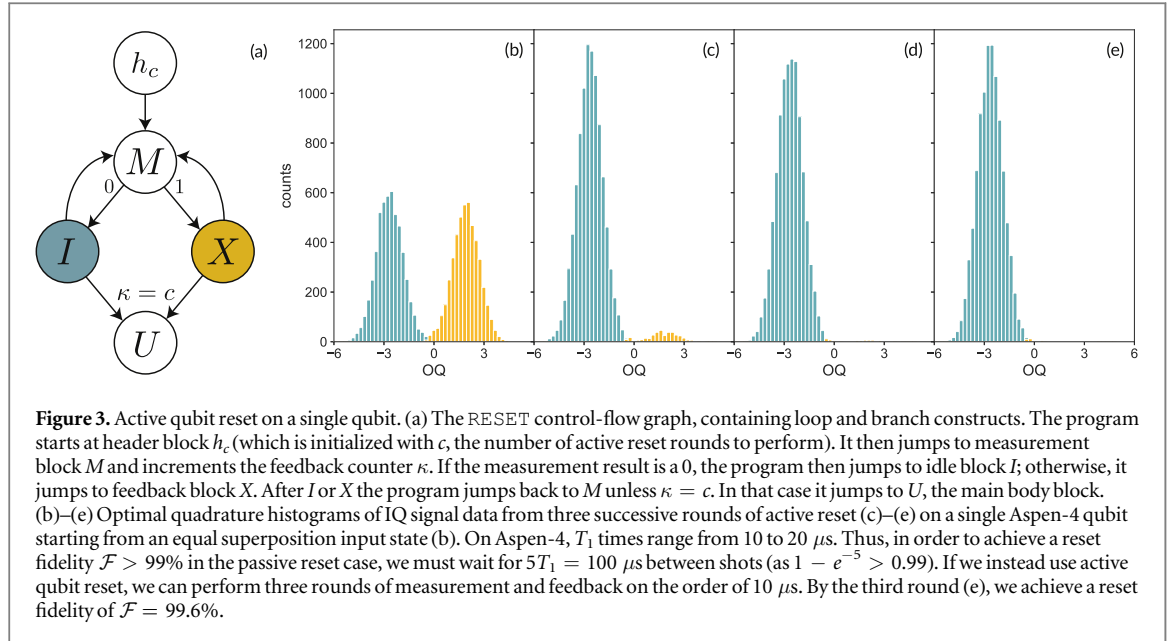
compilation toolchain can propagate control flow structures down to the hardware pulse sequencers. Coupled with the ability to rapidly broadcast qubit measurement results across the control system, active reset is not a special-case operation but instead a simple `if-then-else` control flow branching off of qubit measurement results. By providing the `RESET` instruction at the top of a Quil program, users can signal to the compilation toolchain that they would like to enable active qubit reset. When ingested by the compiler, the `RESET` directive is translated to control-flow graphs (CFGs) [33] for each qubit that encapsulate the branching and looping structure of the active reset protocol.

Control-flow graphs are composed of nodes called *basic blocks*, each of which includes a sequence of instructions without branching and up to two directed edges, or *jump targets*, to follow once the instructions are complete. If a block has two jump targets, it also contains a conditional expression for choosing between them. Each single-qubit `RESET` CFG contains four basic blocks—a header block, a measurement block, an idle block, and a feedback block (figure 3(a)). The program starts at the header block, which initializes a counter to track the number of active reset rounds performed, and then jumps to the measurement block and increments the counter. The measurement block contains readout and capture instructions for the qubit, and conditionally jumps to either the idle or feedback blocks dependent on the bit result produced by the measurement. If the result is 1, the qubit is in the excited state, and therefore the feedback block containing the `X` gate program is executed. Otherwise, if the result is 0, the program jumps to the idle block. After either the idle or feedback block is completed, the program then jumps to the measurement block again if the counter is less than the number of reset rounds requested. After each single-qubit program traverses its active reset CFG, the program jumps to the first basic block of the main body quantum circuit and proceeds until completion.

The time required to perform the active reset protocol is dependent on the underlying architecture of a QPU. For Aspen-4, we require three rounds of feedback to achieve a reset performance comparable to that of passive qubit reset (figures 3(b)–(e)). All active reset sequences must be completed before the main body program begins, and the readout operations on Aspen-4 are in the 2–3 μs range. Combined with a broadcast and feedback latency of around 1 μs , this results in an active reset time of 9–12 μs , approximately a tenfold improvement over passive reset times.

4. A volumetric framework for benchmarking runtime

A useful metric for determining the runtime performance of a quantum-classical cloud platform is QPU latency, which tells us how long the CPU has to wait between requesting execution on a quantum device and receiving back the results. Generically, this latency depends on the number of qubits involved in the computation, the number of shots requested, and the program being run. We can define a function $\mathcal{T}(m, n, \mathcal{P})$ which returns QPU latency and takes a number of qubits m , a number of shots n , and a function \mathcal{P} that produces a family of quantum programs for a given number of qubits. Some examples of \mathcal{P} are `GHZ_LINE`, which produces the straight-line program to create a GHZ state on m qubits [34], and `MAXCUTQAOA_COMPLETE`, which produces a QAOA program to solve Max-Cut on a fully-connected graph with m nodes [12]. In addition, by sweeping n and measuring \mathcal{T} with fixed m and \mathcal{P} , we can determine two asymptotic latencies of a platform for a particular



family of quantum programs and number of qubits: the single-shot limit ($n = 1$) tells us the *variational step latency* (\mathcal{T}_V) and the scaling of the latency in the limit of large n tells us the *quantum shot latency* (\mathcal{T}_Q). More formally, we assume that the total latency \mathcal{T} is straightforwardly related to the two components in the following manner

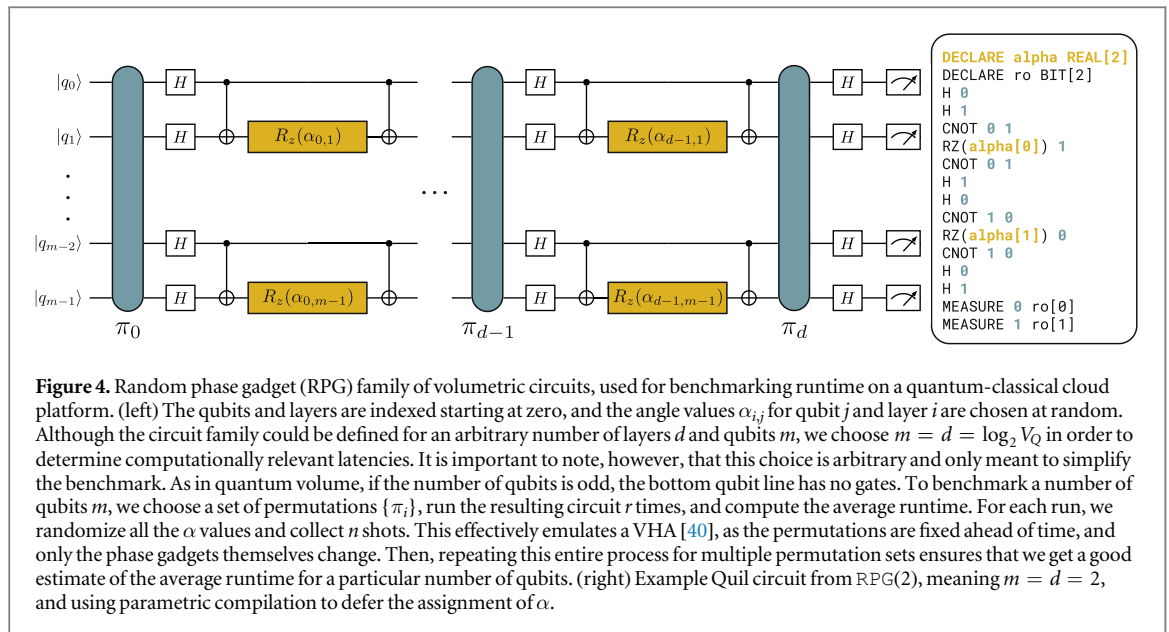
$$\mathcal{T}(m, n, \mathcal{P}) = \mathcal{T}_V(m, \mathcal{P}) + n\mathcal{T}_Q(m, \mathcal{P}). \tag{1}$$

However, \mathcal{T}_V and \mathcal{T}_Q are still functions that depend on the number of qubits and the requested quantum program, and therefore need further specification to be an effective cross-platform benchmark. Naively choosing $m = 1$ and \mathcal{P} to be circuits containing only MEASURE instructions makes it trivial to experimentally benchmark the system. By sweeping the number of shots, measuring latency, and fitting the data to a linear model, the resulting latency-axis intercept gives $\mathcal{T}_V(1, \text{MEASURE})$ and the resulting slope gives $\mathcal{T}_Q(1, \text{MEASURE})$. These numbers well-encapsulate some parts of the performance of a particular architecture, but fail to capture how the architecture fares as the number of qubits increases or as the program complexity changes. For example, potential pitfalls such as long gate times or a poorly scaling control system initialization routine are entirely omitted by the benchmark. But, simply setting m to be the maximum device size available on a particular platform is equally misleading, as current NISQ devices often have error rates such that they cannot produce useful entanglement across the entire QPU.

Although there is no unanimously supported benchmark for QPU performance within the community (and there may never be), log quantum volume [35, 36] ($\log_2 V_Q$) has been proposed as a reasonable near-term metric for the number of qubits that can be meaningfully used in a computation. We are interested in something that is similar to quantum volume (so that the choice of number of qubits remains relevant), but more appropriate for the near-term VHAs. Variational algorithms contain structures known as *phase gadgets* [37], which are RZ gates sandwiched between CNOTs. These structures are often the cornerstone of parametric-ansatz-style programs, and therefore we propose using a volumetric family of circuits [38] that we call *random phase gadgets* (RPG) to benchmark algorithm runtime. The RPG circuit family incorporates the permutation aspect of quantum volume for exercising connectivity, parallelism, and gateset [39], but replaces the random 2Q unitaries with phase gadgets that have RZ gates with randomly chosen arguments (figure 4). In addition, each permutation is followed by a layer of Hadamard gates on all qubits, to make it more difficult to compile away the phase gadgets. Setting $m = \log_2 V_Q$ and $\mathcal{P} = \text{RPG}$ gives us the *computationally relevant* step (\mathcal{T}_V) and shot (\mathcal{T}_Q) latency of a QPU

$$\mathcal{T}_V = \mathcal{T}_V(\log_2 V_Q, \text{RPG}), \quad \mathcal{T}_Q = \mathcal{T}_Q(\log_2 V_Q, \text{RPG}). \tag{2}$$

Fitting the resulting runtime data to the linear model $\mathcal{T}(n) = \mathcal{T}_V + n\mathcal{T}_Q$ then allows us to easily estimate variational algorithm runtimes using the computationally relevant QPU latency for a particular device available on a quantum cloud platform.

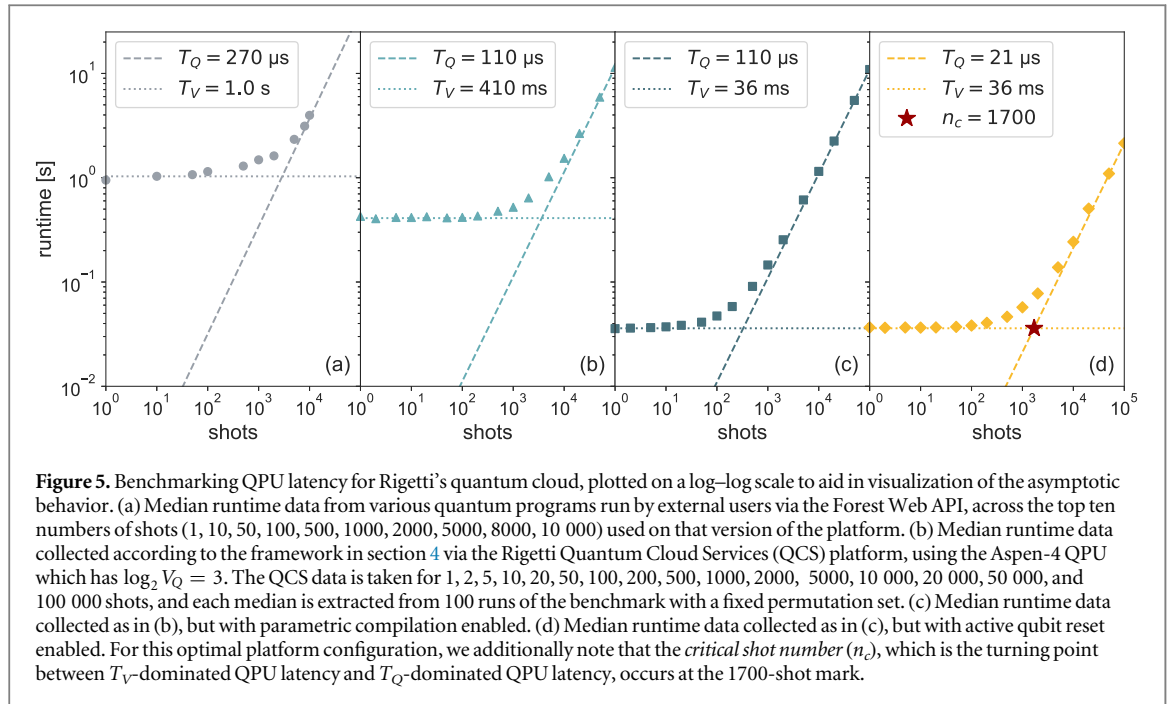


5. QPU latency results on Quantum Cloud Services

With a volumetric framework for benchmarking runtime, we calculate T_V and T_Q for four different versions of the Rigetti quantum cloud platform. To demonstrate the initial performance improvements resulting from simply colocating CPU and QPU, we first analyzed runtime data from over 851 000 quantum programs run by external users on the Acorn and Agave QPUs via the Forest Web API, the initial version of Rigetti's platform (figure 5(a)). The Forest Web API used the first-generation model of quantum cloud access, routing each job through a queue on Amazon Web Services (AWS), which resulted in considerable latencies. Because the Forest Web API programs were run by external users, there is no expectation that the average composition of these programs would match the composition of the benchmark programs described in section 4. However, on average the Forest Web API programs used 3 qubits, 2 CZ gates, and 14 RX gates. A program from RPG(3), upon compilation to the native gateset available on Acorn and Agave, would have more of both gates, and therefore would take longer to execute. This, combined with the fact that the log quantum volume for Aspen-4 is $\log_2 V_Q = 3$, makes it reasonable to compare the data from the Forest Web API and QCS, and in fact skews the comparison in favor of the former. Thus, using median runtime data from Forest Web API, we calculate $T_Q = 270 \mu\text{s}$ and $T_V = 1.0 \text{ s}$. If we instead use the benchmarking framework from section 4 and collect data via QCS's colocated architecture, $T_V = 410 \text{ ms}$ and $T_Q = 110 \mu\text{s}$ for the Aspen-4 QPU (figure 5(b)). Using parametric compilation (figure 5(c)), we can remove the compile step from our runtime calculations, resulting in an improvement for small numbers of shots ($T_V = 36 \text{ ms}$). For higher numbers of shots, passive reset times still dwarf the constant improvement from parametric compilation. Finally, by enabling active qubit reset (figure 5(d)), we get an additional reduction in latency within the quantum execution loop. Thus, in this optimal configuration of the platform, $T_V = 36 \text{ ms}$ and $T_Q = 21 \mu\text{s}$, resulting in greater than $27\times$ and $12\times$ improvements, respectively, over the latencies of the first-generation access model.

6. Conclusions

Quantum Cloud Services may be the first instance of a quantum-classical cloud platform, but this architectural paradigm will become increasingly common as the industry continues to progress toward useful applications of quantum computers. Error rates and qubit count are well-known to be important system benchmarks within the field, but as more and more hardware providers begin to offer access to quantum resources over the cloud, the latencies that govern this access and the resulting application runtimes will also be critical considerations for platform performance. Addressing these latencies requires approaching system bottlenecks with an interdisciplinary *quantum software engineering* mindset, bridging the knowledge bases of classical and quantum computing. We have shown that collocation, parametric compilation, and active qubit reset provide considerable improvements over the first-generation of quantum cloud offerings, but they are just a few of the



many potential platform optimizations for accelerating industry progress and enabling the achievement of quantum advantage.

Acknowledgments

RSS designed the language constructs that support parametric compilation. PJK, ECP, NAT, and RSS built the parametric compilation toolchain. NAT implemented the software for expressing active qubit reset as a control-flow graph. CAR coordinated the integration of active reset into the QCS platform. PJK, ECP, and CAR formulated the benchmark, and PJK collected and analyzed all runtime data. PJK and MPS architected the framework for running VHAs using parametric compilation. RSS supervised the QCS effort. PJK, CAR, and MPS wrote the manuscript and prepared the figures³.

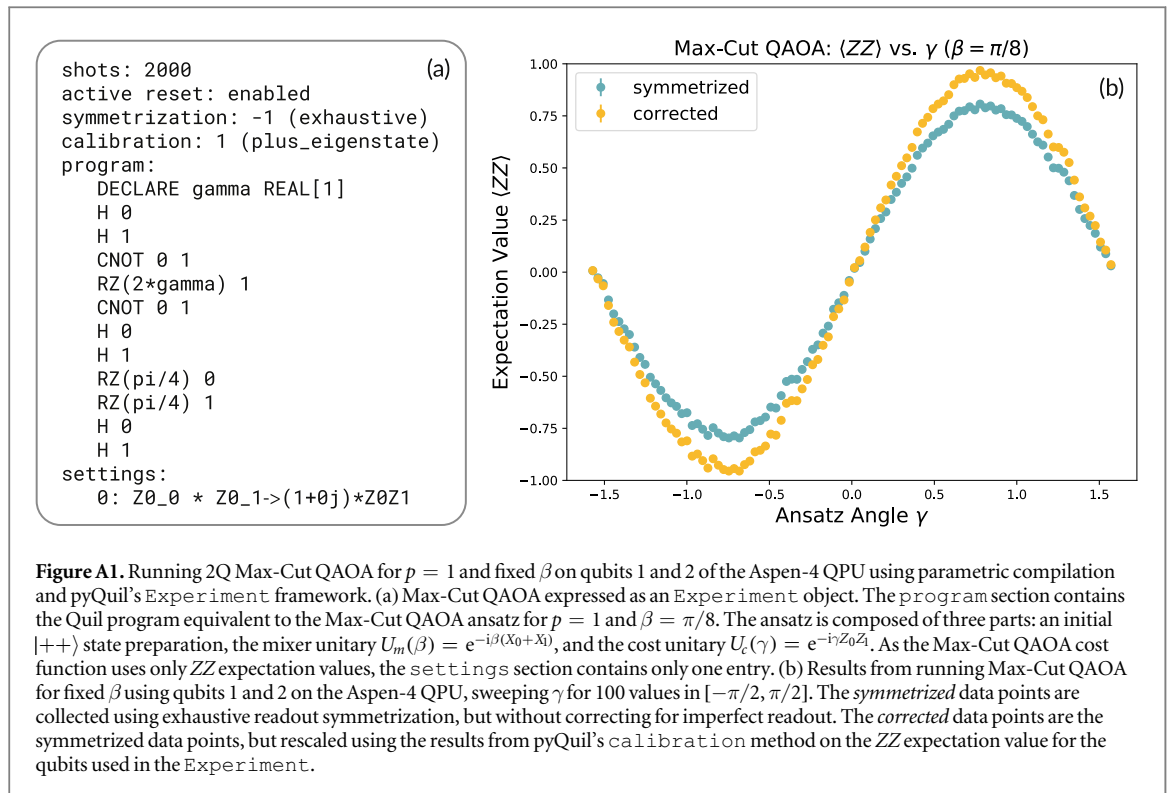
This work was funded by Rigetti & Co Inc., dba Rigetti Computing. We thank the Rigetti quantum software team for providing tooling support, the Rigetti fabrication team for manufacturing the device, the Rigetti technical operations team for fridge maintenance, the Rigetti cryogenic hardware team for providing the chip packaging, the Rigetti control systems and embedded software teams for creating the Rigetti AWG control system, and the Rigetti quantum engineering team for building the infrastructure for automated QPU bringup and recalibration.

In addition, the authors would like to specifically thank Lauren E Capelluto, Steven Heidel, and Anthony M Polloreno for their critical contributions to the QPU compiler toolchain, Glenn E Jones, Rodney F Sinclair, and Blake R Johnson for their work on designing a control system capable of supporting active reset, Alexander D Hill, Joseph A Valery, and Alexa N Staley for their management of QPU performance and deployment, Zachary P Beane, Adam D Lynch, Nicolas J Ochem, and Christopher B Osborn for their orchestration of the QMI compute infrastructure, E Schuyler Fried, Diego Scarabelli, and Prasahnt Sivarajah for their crucial work on experimental bringup of active qubit reset, Joshua Combes, Kyle V Gulshen, and Nicholas C Rubin for building the readout error mitigation tools that enable the implementation of VHAs, Matthew P Harrigan and William J Zeng for their leadership in designing a quantum programming interface to QCS, and Matthew J Reagor for insightful discussions on modeling and benchmarking QPU latency.

Appendix A. Implementing hybrid algorithms with parametric compilation

The structure of quantum programs for variational hybrid algorithms can be segmented into two parts [42]: parameterized ansatz preparation, and measurement in a variety of multi-qubit Pauli bases. Parametric

³ All of the plotting and analysis from the paper can be recreated using the supplementary Jupyter notebooks and datasets [41], which are also available on the notebook hosting service [Binder](#).



compilation handles not only the parameterized ansatz component of a VHA, but can also be used to encapsulate measurement in arbitrary bases and symmetrization of readout error—all in a single quantum program. This allows for any VHA to be expressed via a single parametric binary with RZ arguments that are provided and patched at runtime. To show how this is possible, we walk through three examples using pyQuil [43], the Python library for writing and executing Quil programs.

A.1. Readout error symmetrization and mitigation

We can describe measurement imperfections in terms of a *confusion matrix* \mathbf{M} , which describes probabilities of reported measurement outcomes conditioned on the values of the true state of the qubits. For the case of a single qubit, the confusion matrix can be written as

$$\mathbf{M} = \begin{pmatrix} \Pr(0|0) & \Pr(1|0) \\ \Pr(0|1) & \Pr(1|1) \end{pmatrix} = \begin{pmatrix} 1 - \epsilon_0 & \epsilon_1 \\ \epsilon_0 & 1 - \epsilon_1 \end{pmatrix}. \quad (\text{A.1})$$

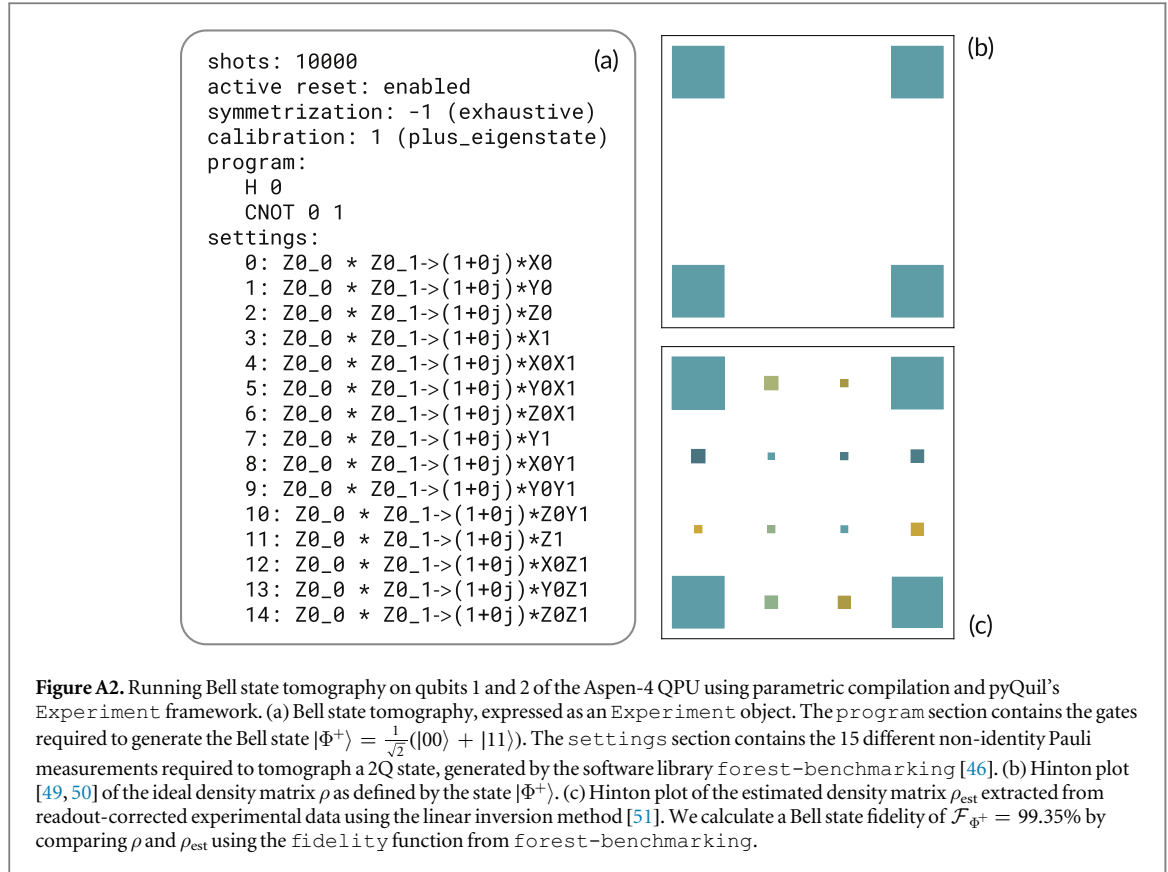
While generally \mathbf{M} will not be symmetric with respect to the exchange of 0s and 1s, one can enforce such symmetry by flipping bits before measurement and subsequently flipping the measurement outcomes. This *symmetrization* procedure corresponds to a simple form of *twirling* [44, 45], and in principle can be performed by randomly choosing which bits to flip independently each time a measurement occurs⁴. When averaging over all the shots, this results in a new *effective* confusion matrix \mathbf{M}' , with ϵ_0 and ϵ_1 replaced by ϵ' , which is the arithmetic mean of the two. This can be implemented in Quil (for a single qubit with index 0) by replacing the measurement section of a program with:

```
DECLARE symmetrization REAL[1]
DECLARE ro BIT[1]
RX(symmetrization[0]) 0
MEASURE 0 ro[0].
```

During the execution of the program, the `symmetrization` memory region assignments 0 and π must be provided, and the results from the π assignment must be flipped (i.e. XORing the results with the bit value of 1).

It is straightforward to check that the effect of symmetric measurement errors is to scale the expectation value of a Pauli observable by factor that depends on the error rates. This immediately suggests how to mitigate the effect of these errors: first characterize the error rates of the symmetrized readout, then rescale the observed expectation values accordingly. The first step is what we call *readout calibration*, while the second step is what we

⁴ More sophisticated approaches to symmetrization can be taken by making use of orthogonal arrays, as implemented in `forest-benchmarking` [46], but a detailed description of those techniques is beyond the scope of this article.



call *readout error mitigation*. Since errors due to the single-qubit rotations associated with mapping one Pauli observable to another are orders of magnitude smaller than measurement errors, readout calibration can focus on the expectation of tensor products of Z observables only. Ideally the expectation of any of these tensor products would be $+1$ for the ground state (the $|00 \dots 0\rangle$ state), but for symmetrized readout error the expectation value will be λ . Simply dividing the expectation value of this same observable for other states by λ corrects for the bias introduced by the measurement errors. Note that, in general, even with symmetrization, a different λ for each different tensor product will be necessary, due to qubit-dependent readout signal-to-noise ratio (SNR) and potential correlations in the readout errors. Moreover, while the bias in the estimation of the expectation value is removed, the uncertainty is increased (as $|\lambda| < 1$), so this procedure is not scalable [47]. That being said, it is remarkably effective for small numbers of qubits.

Rather than handling symmetrization manually, we can use pyQuil's Experiment framework. When defining an Experiment in pyQuil, two arguments are required—a program which defines the main body quantum circuit, and a list of settings which specify the different state preparations and measurements that we want to wrap around the main program. In addition, the Experiment object also specifies the number of shots to take, whether or not to perform active qubit reset, and how to symmetrize and calibrate readout. To show how these features are used, we implement readout symmetrization and correction for a 2Q Max-Cut QAOA program (figure A1).

A.2. Bell state tomography

An arbitrary single-qubit gate U can always be decomposed into what is in essence an Euler angle decomposition [48]

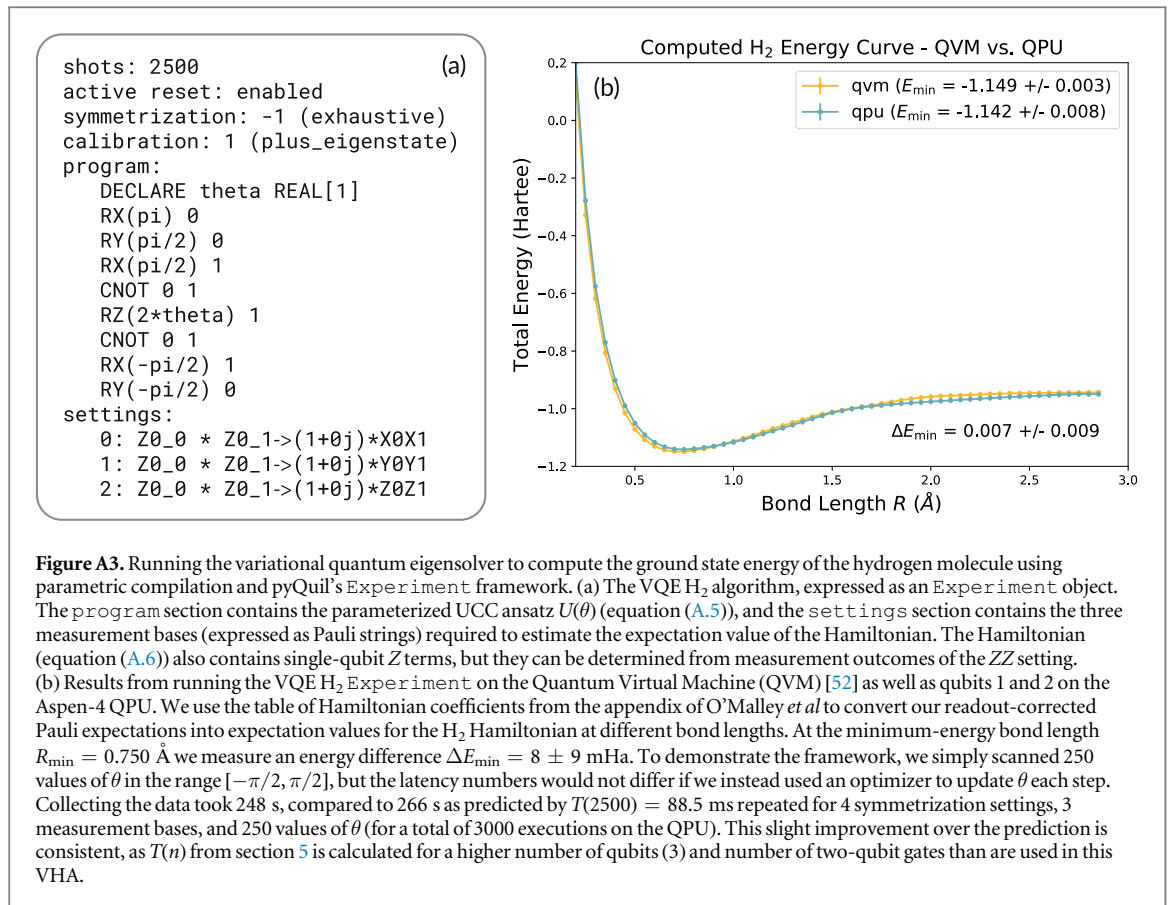
$$U(\alpha, \beta, \gamma) = R_z(\gamma)R_x(-\pi/2)R_z(\beta)R_x(\pi/2)R_z(\alpha). \quad (\text{A.2})$$

Although our QPUs can only perform measurements in the z basis (M_z), performing single-qubit rotations prior to measurement allows us to effectively measure in a different basis

$$M_x = M_z R_y(-\pi/2) = M_z U(0, -\pi/2, 0), \quad (\text{A.3})$$

$$M_y = M_z R_x(\pi/2) = M_z U(\pi/2, \pi/2, -\pi/2). \quad (\text{A.4})$$

Thus, by pre-pending a qubit measurement with an Euler-decomposed single-qubit gate containing R_z arguments that can be changed at runtime, we can perform a collection of different measurements with a single



parametric binary. In Quil, this is accomplished for a single qubit by adding the following snippet before the measurement block:

```
DECLARE measurement_alpha REAL[1]
DECLARE measurement_beta REAL[1]
DECLARE measurement_gamma REAL[1]
RZ(measurement_alpha[0]) 0
RX(pi/2) 0
RZ(measurement_beta[0]) 0
RX(-pi/2) 0
RZ(measurement_gamma[0]) 0.
```

Additionally, this can be combined with readout symmetrization, allowing for any desired observable to be symmetrized, calibrated, and corrected. To show how measurement bases can be changed parametrically, we run Bell state tomography using a single binary (figure A2).

A.3. The variational quantum eigensolver

Finally, we combine the techniques from the previous two examples in order to run a full variational algorithm using a single parametric binary. The *variational quantum eigensolver* (VQE) [15], which is one of the leading VHAs for applications in quantum chemistry, can be used to compute the ground state energy of the hydrogen molecule (H₂). To do so, VQE employs a classical minimization routine in which the objective function is evaluated on the QPU. The procedure begins by preparing a parameterized ansatz wavefunction $|\Psi(\theta)\rangle$ using an initial guess for the variational parameter θ . The parameterized ansatz wavefunction can be chosen to be composed of the unitary coupled cluster (UCC) ansatz $U(\theta)$ and the Hartree–Fock (HF) reference state $|\Phi\rangle$ [14], giving

$$|\Psi(\theta)\rangle = U(\theta)|\Phi\rangle = e^{-i\theta X_0 Y_1}|01\rangle. \quad (\text{A.5})$$

Then, the procedure continues by computing the expectation value of the H₂ Hamiltonian, which can be formulated as a sum of multi-qubit Pauli-basis expectation values with coefficients

$$H = g_0 \mathbb{1} + g_1 Z_0 + g_2 Z_1 + g_3 Z_0 Z_1 + g_4 Y_0 Y_1 + g_5 X_0 X_1. \quad (\text{A.6})$$

Finally, the expectation value $\langle H \rangle$ is fed to the minimizer to choose the variational parameter θ for the next round of iteration, and this repeats until convergence. Following the protocol in O'Malley *et al*, we use parametric compilation and pyQuil's `Experiment` framework to simulate the ground state energy of the H_2 molecule with VQE (figure A3).

ORCID iDs

Peter J Karalekas  <https://orcid.org/0000-0001-6031-4800>

Nikolas A Tezak  <https://orcid.org/0000-0002-5178-2298>

Eric C Peterson  <https://orcid.org/0000-0002-1633-0050>

Colm A Ryan  <https://orcid.org/0000-0001-9923-3354>

Marcus P da Silva  <https://orcid.org/0000-0002-6641-8712>

Robert S Smith  <https://orcid.org/0000-0002-6340-9589>

References

- [1] Chuang I L, Gershenfeld N and Kubinec M 1998 Experimental implementation of fast quantum searching *Phys. Rev. Lett.* **80** 3408–11
- [2] DiCarlo L *et al* 2009 Demonstration of two-qubit algorithms with a superconducting quantum processor *Nature* **460** 240–4
- [3] Gulde S, Riebe M, Lancaster G P T, Becher C, Eschner J, Häffner H, Schmidt-Kaler F, Chuang I L and Blatt R 2003 Implementation of the deutsch-jozsa algorithm on an ion-trap quantum computer *Nature* **421** 48
- [4] Kwiat P G, Mitchell J R, Schwindt P D D and White A G 2000 Grover's search algorithm: an optical approach *J. Mod. Opt.* **47** 257–66
- [5] Arute F *et al* 2019 Quantum supremacy using a programmable superconducting processor *Nature* **574** 505–10
- [6] Ballance C J, Harty T P, Linke N M, Sepiol M A and Lucas D M 2016 High-fidelity quantum logic gates using trapped-ion hyperfine qubits *Phys. Rev. Lett.* **117** 060504
- [7] Hong S S, Papageorge A T, Sivarajah P, Crossman G, Dider N, Polloreno A M, Sete E A, Turkowski S W, da Silva M P and Johnson B R 2020 Demonstration of a parametrically activated entangling gate protected from flux noise *Physical Review A* **101**
- [8] Preskill J 2012 Quantum computing and the entanglement frontier arXiv:1203.5813 [quant-ph]
- [9] Preskill J 2018 Quantum computing in the NISQ era and beyond *Quantum* **2** 79
- [10] Mell P M and Grance T 2011 The NIST definition of cloud computing *Technical Report* SP 800-145 Gaithersburg, MD, United States
- [11] McClean J R, Romero J, Babbush R and Aspuru-Guzik A 2016 The theory of variational hybrid quantum-classical algorithms *New J. Phys.* **18** 023023
- [12] Farhi E, Goldstone J and Gutmann S 2014 A quantum approximate optimization algorithm arXiv:1411.4028 [quant-ph]
- [13] Otterbach J S *et al* 2017 Unsupervised machine learning on a hybrid quantum computer arXiv:1712.05771 [quant-ph]
- [14] O'Malley P J J *et al* 2016 Scalable quantum simulation of molecular energies *Phys. Rev. X* **6** 031007
- [15] Peruzzo A, McClean J, Shadbolt P, Yung M-H, Zhou X-Q, Love P J, Aspuru-Guzik A and O'Brien J L 2014 A variational eigenvalue solver on a photonic quantum processor *Nat. Commun.* **5** 4213
- [16] Wilson C M, Otterbach J S, Tezak N, Smith R S, Polloreno A M, Karalekas P J, Heidel S, Alam M S, Crooks G E and da Silva M P 2018 Quantum kitchen sinks: An algorithm for machine learning on near-term quantum computers arXiv:1806.08321 [quant-ph]
- [17] Bravo-Prieto C, LaRose R, Cerezo M, Subasi Y, Cincio L and Coles P J 2019 Variational quantum linear solver: A hybrid algorithm for linear systems arXiv:1909.05820 [quant-ph]
- [18] Verdon G, Marks J, Nanda S, Leichenauer S and Hidary J 2019 Quantum hamiltonian-based models and the variational quantum thermalizer algorithm arXiv:1910.02071 [quant-ph]
- [19] Kjaergaard M, Schwartz M E, Braumüller J, Krantz P, Wang J I J, Gustavsson S and Oliver W D 2019 Superconducting qubits: Current state of play arXiv:1905.13641 [quant-ph]
- [20] Cross A W, Bishop L S, Smolin J A and Gambetta J M 2017 Open quantum assembly language arXiv:1707.03429 [quant-ph]
- [21] Smith R S, Curtis M J and Zeng W J 2016 A practical quantum instruction set architecture arXiv:1608.03355 [quant-ph]
- [22] McClean J R, Boixo S, Smelyanskiy V N, Babbush R and Neven H 2018 Barren plateaus in quantum neural network training landscapes *Nat. Commun.* **9** 4812
- [23] Smith R S, Peterson E C, Davis E J and Skilbeck M G 2020 *Quilc: An Optimizing Quil Compiler* (Zenodo) (<https://doi.org/10.5281/zenodo.3677536>)
- [24] Gyenis A, Mundada P S, Di Paolo A, Hazard T M, You X, Schuster D I, Koch J, Blais A and Houck A A 2019 Experimental realization of an intrinsically error-protected superconducting qubit arXiv:1910.07542 [quant-ph]
- [25] Nersisyan A *et al* 2019 Manufacturing low dissipation superconducting quantum processors arXiv:1901.08042 [quant-ph]
- [26] Smith R S 2020 *Quil: A Portable Quantum Instruction Language* (Zenodo) (<https://doi.org/10.5281/zenodo.3677540>)
- [27] TIS Committee *et al* 1995 Tool interface standard (TIS) executable and linking format (ELF) specification version 1.2 *TIS Committee*
- [28] IEEE Standard for Information Technology—Portable Operating System Interface (POSIX(R)) Base Specifications, Issue 7, in IEEE Std 1003.1-2017 (Revision of IEEE Std 1003.1-2008), pp.1–3951 (Accessed: 31 Jan. 2018)
- [29] Egger D J, Werninghaus M, Ganzhorn M, Salis G, Fuhrer A, Müller P and Filipp S 2018 Pulsed reset protocol for fixed-frequency superconducting qubits *Phys. Rev. Appl.* **10** 044030
- [30] Magnard P *et al* 2018 Fast and unconditional all-microwave reset of a superconducting qubit *Phys. Rev. Lett.* **121** 060502
- [31] Risté D, Bultink C C, Lehnert K W and DiCarlo L 2012 Feedback control of a solid-state qubit using high-fidelity projective measurement *Phys. Rev. Lett.* **109** 240502
- [32] Ryan C A, Johnson B R, Risté D, Donovan B and Ohki T A 2017 Hardware for dynamic quantum computing *Rev. Sci. Instrum.* **88** 104703
- [33] Aho A V, Sethi R and Ullman J D 1986 *Compilers: Principles, Techniques, and Tools* (Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc.)
- [34] Greenberger D M, Horne M A and Zeilinger A 2007 Going beyond bell's theorem arXiv:0712.0921 [quant-ph]

- [35] Cross A W, Bishop L S, Sheldon S, Nation P D and Gambetta J M 2019 Validating quantum computers using randomized model circuits *Phys. Rev. A* **100** 032328
- [36] Moll N et al 2018 Quantum optimization using variational algorithms on near-term quantum devices *Quantum Sci. Technol.* **3** 030503
- [37] Cowtan A, Dilkes S, Duncan R, Simmons W and Sivarajah S 2019 Phase gadget synthesis for shallow circuits arXiv:1906.01734 [quant-ph]
- [38] Blume-Kohout R and Young K C 2019 A volumetric framework for quantum computer benchmarks arXiv:1904.05546 [quant-ph]
- [39] Peterson E C, Crooks G E and Smith R S 2019 Fixed-depth two-qubit circuits and the monodromy polytope arXiv:1904.10541 [quant-ph]
- [40] Crooks G E 2018 Performance of the quantum approximate optimization algorithm on the maximum cut problem arXiv:1811.08419 [quant-ph]
- [41] Karalekas P J 2020 *Supplementary Notebooks and Datasets for: A Quantum-Classical Cloud Platform Optimized for Variational Hybrid Algorithms* (Zenodo) (<https://doi.org/10.5281/zenodo.3635021>)
- [42] McCaskey A J, Parks Z P, Jakowski J, Moore S V, Morris T D, Humble T S and Pooser R C 2019 Quantum chemistry as a benchmark for near-term quantum computers *npj Quantum Information* **5** 1–8
- [43] Karalekas P J et al 2020 *PyQuil: Quantum Programming in Python* (Zenodo) (<https://doi.org/10.5281/zenodo.3631770>)
- [44] Bartlett S D, Rudolph T and Spekkens R W 2007 Reference frames, superselection rules, and quantum information *Rev. Mod. Phys.* **79** 555–609
- [45] Bennett C H, DiVincenzo D P, Smolin J A and Wootters W K 1996 Mixed-state entanglement and quantum error correction *Phys. Rev. A* **54** 3824–51
- [46] Gulshen K V et al 2019 *Forest Benchmarking: QCVV using PyQuil* (Zenodo) (<https://doi.org/10.5281/zenodo.3455848>)
- [47] Ryan C A, Johnson B R, Gambetta J M, Chow J M, da Silva M P, Dial O E and Ohki T A 2015 Tomography via correlation of noisy measurement records *Phys. Rev. A* **91** 022118
- [48] Nielsen M A and Chuang I L 2011 *Quantum Computation and Quantum Information: 10th Anniversary Edition* 10th edn (New York, NY, USA: Cambridge University Press)
- [49] Hunter J D 2007 Matplotlib: a 2d graphics environment *Comput. Sci. Eng.* **9** 90–5
- [50] Virtanen P et al 2020 SciPy 1.0: fundamental algorithms for scientific computing in Python *Nature Methods* **17** 261–72
- [51] Wood C J 2015 Initialization and characterization of open quantum systems (<http://hdl.handle.net/10012/9557>)
- [52] Smith R S 2020 *Quantum Virtual Machine* (Zenodo) (<https://doi.org/10.5281/zenodo.3677538>)