

Learning from Physics Experiments with Quantum Computers: Applications in Muon Spectroscopy

Sam McArdle 

Department of Materials, University of Oxford, Parks Road, Oxford OX1 3PH, United Kingdom



(Received 16 December 2020; accepted 29 April 2021; published 29 June 2021)

Computational physics is an important tool for analyzing, verifying, and—at times—replacing physical experiments. Nevertheless, simulating quantum systems and analyzing quantum data has so far resisted an efficient classical treatment in full generality. While programmable quantum systems have been developed to address this challenge, the resources required for classically intractable problems still lie beyond our reach. In this work, we consider a new target for quantum-simulation algorithms; analyzing the data arising from physics experiments—specifically, muon-spectroscopy experiments. These experiments can be used to probe the quantum interactions present in condensed-matter systems. However, fully analyzing their results can require classical computational resources scaling exponentially with the simulated system size, which can limit our understanding of the studied system. We show that this task may be a natural fit for the coming generations of quantum computers. We use classical emulations of our quantum algorithm on systems of up to 29 qubits to analyze real experimental data, and to estimate both the near-term and error-corrected resources required for our proposal. We find that our algorithm exhibits good noise resilience, stemming from our desire to extract global parameters from a fitted curve, rather than targeting any individual data point. In some respects, our fault tolerant resource estimates go further than some prior work in quantum simulation, by estimating the resources required to solve a complete task, rather than just to run a given circuit. Taking the overhead of observable measurement and calculating multiple datapoints into account, we find that significant challenges still remain if our algorithm is to become practical for analyzing muon-spectroscopy data.

DOI: [10.1103/PRXQuantum.2.020349](https://doi.org/10.1103/PRXQuantum.2.020349)

At first glance, quantum computing appears to offer remarkable computational power; potentially yielding exponential speedups for simulating quantum systems, or for solving some problems in machine learning. In reality, however, the situation is more nuanced, and there are a number of challenges that quantum algorithms must overcome, in order to become practical for real problems of interest. At the heart of these challenges, is that we must first build a quantum computer that is both sufficiently large, and has low enough error rates to run the calculation (in an error-corrected setting these essentially become the same requirement, as the error rate can be suppressed by adding additional qubits for increased error protection). How large and noiseless the computer must

be is determined by the efficiency of the proposed quantum algorithms, as well as the efficiency of the classical algorithms they aim to surpass.

The most promising use cases for quantum simulation and machine-learning algorithms seem to require a number of features. Firstly, the problem should be hard to solve classically, and should not admit accurate classical approximations. Quantum algorithms for solving the electronic structure problem must overcome this challenge, as sophisticated classical methods can accurately approximate the low-lying energy levels of many small systems of interest, up to around 100 qubits [1,2]. Secondly, the output of the algorithm should be resilient to noise, which will reduce the resources required when performing quantum error correction, or make the calculation amenable to noisy, near-term quantum hardware. In the case of quantum machine-learning algorithms, the problem should also have an inherently quantum structure, to avoid relying on heuristic arguments for quantum advantage. The model should be efficient to train, which is aided by having a good prior for the model parameters and structure [3,4]. Finally, there should be no data input and output limitations, which avoids the need for hypothetical data structures like quantum random-access memory, or the

*sam.mcardle.science@gmail.com

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

ability to directly interface a quantum computer with other quantum systems, which supply the input data state, such as the interactive quantum-likelihood evaluation method [5] or quantum convolutional neural networks [6].

In this work, we introduce a quantum algorithm that sits at the interface of quantum simulation and quantum machine learning, which we apply to analyzing the spectra arising from muon spin rotation, resonance, and relaxation (μ^+ SR) experiments. We seek to optimize our algorithm against the criteria listed above, in order to ensure its practicality for solving real problems of interest. Muon-spectroscopy experiments have been used to analyze a wide range of physical systems and phenomena, such as superconductivity and magnetism. These experiments typically measure the time evolution of the spin polarization of antimuons that have been implanted into a sample of interest [7]. A typical experimental setup is shown in Fig. 1 and discussed fully in Sec. I. As the muon spin interacts with other environmental spins in the system, it oscillates between eigenstates of the system Hamiltonian. Tracking these variations enables us to learn quantitatively about the

interactions at the muon site. Although the inputs to the model are Hamiltonian parameters—classical data—and the outputs of the experiments are classical-data points, accurate analysis can require a fully quantum treatment of the system dynamics. In some cases, this problem appears challenging to solve on classical computers, admitting few simplifications, and often requiring exact diagonalization of the system Hamiltonian [8–10].

In Sec. IV we present an algorithm to solve this problem efficiently using a quantum computer. Our algorithm is relatively simple, and requires modest computational resources. In particular, the finite lifetime of the antimuon ($2.2 \mu\text{s}$ on average) and our desire to extract global system properties appears to set a generous limit on the precision required from our quantum calculations. In Sec. V we carry out classical numerical emulations of our quantum algorithm, to investigate its scaling behavior and noise robustness. We use these emulations to analyze muon-spectroscopy data from a real experiment, finding good agreement with recent state-of-the-art classical analysis [9]. The calculations performed as part of our analysis, involving a Hilbert-space dimension of 2^{29} , are the largest

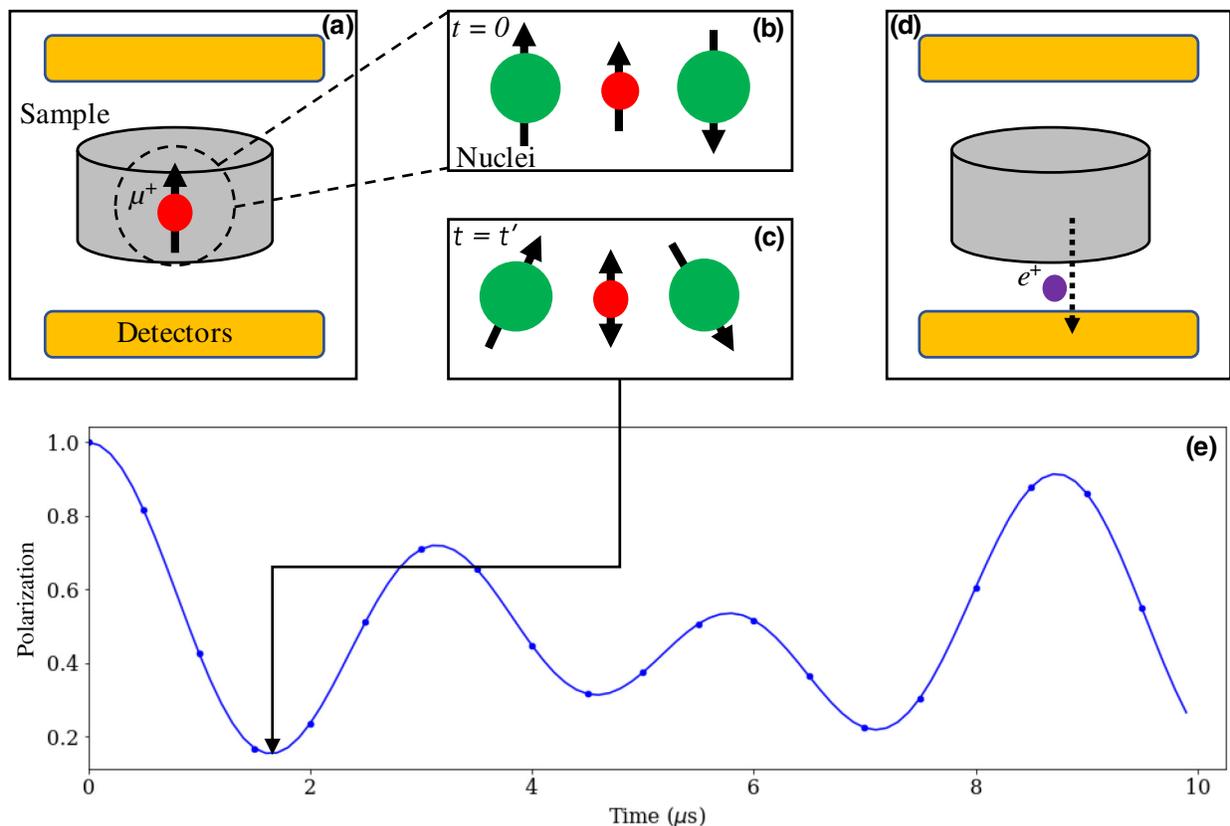


FIG. 1. Schematic of a muon-spectroscopy experiment, as described in the main text. Spin-polarized muons are implanted into a sample of interest (a), whereupon they interact with nearby nuclear spins (b). This causes an oscillation of the muon spin polarization (c). When the muon decays, the positron is preferentially emitted in the direction of the muon spin at the time that it decayed (d). By recording the normalized asymmetry between positron counts registered by the two detectors, we can reconstruct the time evolution of the muon polarization (e).

to date in the muon literature. We use these numerical results as the basis for estimating both the near-term and error-corrected resources required to run our algorithm for problem sizes of interest. At first glance, our algorithm appears to require fewer fault-tolerant resources than solving challenging instances of the electronic structure problem. However, the large number of repetitions required by our algorithm (to estimate observables, calculate multiple datapoints, and repeat the calculation as part of an optimization loop) may result in an impractically long runtime. We argue in Sec. VI that this is not necessarily a limitation of our algorithm, but a challenge facing many quantum algorithms.

Our approach is similar in nature to the proposals of Refs. [11,12], which also consider using quantum computers to analyze the outputs of quantum experiments (inelastic neutron scattering on simplified models of magnetic molecules, and linear response nuclear magnetic resonance experiments, respectively). All three methods are related to quantum-likelihood evaluation [5], whereby the data from one untrusted quantum source (the experiment) is fitted to simulated data from a trusted quantum source (the quantum computer) in order to extract system parameters of interest.

I. MUON SPECTROSCOPY

Muon spectroscopy, more commonly known as muon spin rotation, relaxation, and resonance (μ^+ SR), emerged as an experimental technique in the 1970s [13], following the theoretical proposal by Garwin *et al.* [14]. The technique is closely related to other spin-based methods for probing magnetic interactions, such as nuclear magnetic resonance (NMR) and electron spin resonance (ESR). μ^+ SR experiments can provide quantitative information on a range of physical phenomena, including magnetic ordering and phase transitions in materials, the diffusion of light interstitial defects in semiconductors, vortex formation in superconductors, and low-dimensional magnetism. In this section, we provide an introduction to μ^+ SR. We direct the interested reader to the review article by Blundell [15] and the textbook by Yaouanc and De Reotier [7] for more detailed discussions of the method.

As mentioned in the Introduction, μ^+ SR uses spin-polarized beams of positive (anti)muons (hereafter referred to as muons, as the ordinary muon is rarely used in muon-spectroscopy experiments, and so is not discussed in this work) to probe the interactions in a sample of interest. A typical μ^+ SR experiment proceeds as follows.

- (a) The μ^+ is a positively charged, spin-1/2 particle, with a mass approximately $1/9^{\text{th}}$ that of the proton. It can be produced at particle accelerators, from the decay of pions: $\pi^+ \rightarrow \mu^+ + \nu_\mu$. Working in the rest frame of the pion, and conserving linear and

angular momentum, we find that both the muon-emission direction and spin polarization (the spin projection along a specified axis) must be opposite to those of the neutrino. This decay proceeds via the weak interaction, and so violates parity. This causes the neutrino to have its spin polarization orientated antiparallel to its momentum—and thus, the same happens to the muon. By selecting muons arising from pions decaying at rest, spin-polarized muon beams are produced. In this work, we choose the initial spin polarization of the muon beam to be along the positive Z axis. As a result, we can consider the muons to initially be in state $|0\rangle \langle 0|_\mu$, where $|0\rangle$ is the $+1$ eigenstate of the Pauli- Z matrix.

- (b) The beam is directed into the sample of interest. The muons are brought to rest by electrostatic interactions, which do not depolarize the beam [15].
- (c) The muons interact with their local environment via spin-spin interactions. As the initial polarization of the muon spin is unlikely to be an eigenstate of the muon-environment Hamiltonian, the muon spin polarization will evolve according to the Schrödinger equation.
- (d) The muon decays with a half-life of $2.2 \mu\text{s}$ into a positron, electron neutrino, and muon antineutrino. Once again, as a consequence of the weak interaction, the positron is emitted preferentially in the direction of the muon spin polarization.
- (e) The emitted positrons exit the sample, and are detected by detectors placed forwards and backwards along the direction of the muons' initial spin polarization. The detection of positrons gives an asymmetry function

$$A(t) = \frac{N_+ - \alpha N_-}{N_+ + \alpha N_-}, \quad (1)$$

which records the normalized difference between the forward and backward detector counts ($N_{+/-}$, respectively), where α accounts for detection inefficiencies and imbalances. This value is converted to a polarization function, using

$$P(t) = \frac{A(t) - A_{\text{bg}}}{A_0}, \quad (2)$$

where A_{bg} gives the background positron count, and $A_0 = A(0) - A_{\text{bg}}$. The polarization function is normalized to between ± 1 , and corresponds to the spin polarization of the muon beam at time t . Mathematically, the polarization function is given by

$$\begin{aligned} P(t) &= \text{Tr}[Z_\mu \rho_s(t)], \\ &= \text{Tr}\{Z_\mu e^{-iH_s t} [|0\rangle \langle 0|_\mu \otimes \rho_e(0)] e^{iH_s t} \}, \end{aligned} \quad (3)$$

where Z_μ is the Pauli-Z matrix acting on the muon, H_s is the system Hamiltonian, $\rho_s(t)$ is the state of the muon-environment system at time t , and $\rho_e(0)$ is the initial state of the environment.

- (f) Measuring the time evolution of the polarization function enables us to infer the interactions felt by the muon at its rest site (given a model for the initial state of the environment), and so learn quantitatively about the sample.

An understanding of μ^+ SR experiments can be gleaned from the following semiclassical example. A spin-polarized muon beam is directed into a sample of interest, where it interacts with a transverse magnetic field \vec{B} . The muon spin polarization will precess at a frequency $\omega = B\gamma_\mu$, where $\gamma_\mu = 2\pi \times 135.5 \text{ MHz T}^{-1}$ is the gyromagnetic ratio of the muon. After converting the positron counts into the polarization function, we observe that the polarization function is given by $P(t) = \cos(B\gamma_\mu t)$. This enables us to infer the strength of the magnetic field at the muon site. This example is similar in spirit to many of the early μ^+ SR experiments performed. Due to their large magnetic moment, muons can be used as sensitive probes of both local magnetic fields and spin-spin interactions.

There are two types of muon beams: continuous wave, and pulsed. Muons in a continuous beam are implanted one by one; a timer is started when a muon enters the sample, and stopped when a positron is detected. If the previous muon has not decayed before the next muon arrives, then these events are discarded. In order to keep the count rate high, the polarization function is only recorded to around $10 \mu\text{s}$. Continuous beams have a time resolution of around 1 ns. In contrast, pulsed beams implant hundreds of muons into the sample in a single pulse. The arrival time between pulses is large compared to the muon lifetime, which enables polarization functions to be measured to around $30 \mu\text{s}$. The resolution of pulsed beams is limited to around $0.1 \mu\text{s}$ by the finite width of the pulses. As a result, continuous and pulsed beams are better suited for studying fast and slow dynamics, respectively.

Experiments can be carried out at high pressures (approximately 1.5 GPa), low temperatures (less than 1 K), with high strength transverse or longitudinal magnetic fields (approximately 8 T), or with the application of time-dependent radio-frequency pulses [7]. Each of these techniques enables the investigation of certain phenomena more closely. For example, if we are investigating relaxation of the muon spin polarization due to either T_1 relaxation (arising from magnetic field fluctuations in time, which cause the muon to exchange energy with the environment) or T_2 relaxation (caused by a spatially varying static field distribution, which causes the muons to precess at different frequencies, dephasing the beam) then a strong longitudinal magnetic field can be used. This “locks” the muon polarization along its initial direction, which then

makes it easier to measure differences in the polarization function arising from relaxation processes.

As described above, μ^+ SR experiments have been used to investigate a range of physical phenomena. For example, we can measure the temperature dependence of oscillations in the polarization function (or the lack thereof) in order to investigate phase transitions in magnetic materials, such as low-dimensional spin chains [16]. Other experiments have measured the polarization function of muons in semiconductors, at a range of temperatures, in order to investigate the diffusion of muons within the sample. The muon acts as a light proton (often capturing an electron in semiconductors to form “muonium”), so these experiments can be used to examine the effects of hydrogenic defect diffusion in semiconductors [17]. Similar μ^+ SR experiments have examined Li^+ ion diffusion in lithium battery materials [18–20]. In superconducting systems, μ^+ SR experiments have been used to measure the superconducting electron density, determine phase diagrams, and characterize vortex lattices [7]. μ^+ SR experiments have also been applied to biological and chemical systems; for example, to investigate oxygen-dependent effects in the radiation treatment of cancer [21].

Muon spectroscopy is a versatile technique, that has provided insights on a range of physical systems. The technique is still undergoing active development, including the introduction of lower-energy muon beams, which can be used to probe surface effects [7]. However, there are also theoretical challenges for μ^+ SR that are desirable to address. While some systems can be analyzed using a mean-field, or semi-classical, approach, others appear to require a fully quantum treatment [8,10,22,23]. In the following section, we discuss the simulation and analysis of muon polarization functions in more detail.

II. MUON POLARIZATION FUNCTIONS

In order to analyze the polarization function arising from a given μ^+ SR experiment, we can compare it to a theoretical polarization function obtained from a physical model of the studied system. The accuracy of these theoretical polarization functions is determined by both the level of detail included in the model, and the method used to simulate the model. This can be likened to the field of computational chemistry, where the accuracy of a simulation is determined by both the physical effects included in the model (e.g., the Born-Oppenheimer approximation, or relativistic effects) and the method used to simulate the system (e.g., mean-field approaches, or exact diagonalization). We focus first on the different methods used to obtain the theoretical polarization function, before considering other physical effects that can be incorporated into the model.

On a most basic level, one can use a semiclassical, mean-field approach, which considers how the muon

polarization evolves given a model for the surrounding magnetic field distribution. For example, a time-independent magnetic field with Gaussian distributed field strengths leads to the Kubo-Toyabe polarization function [7,24]. This analytically derivable formula, and generalizations of it, are widely applied within the μ^+ SR field, and are acceptably accurate in many circumstances.

If greater accuracy is required, techniques have been developed that treat the muon interactions with its local spin environment semiclassically [25,26]. In some circumstances, these methods can yield high accuracy at a modest computational cost. However, they cannot always account for strong spin-spin interactions, or the effect of quadrupole interactions [7].

Finally, the highest level of accuracy for a given model can be obtained by using a quantum-mechanical analysis. The calculations consider a quantum Hamiltonian between the muon and its local spin environment, and evolve the muon polarization in time according to the Schrödinger equation. We discuss these calculations in more detail below. However, we note here that the cost of these calculations is believed to scale exponentially with the size of the system simulated, due to the computational complexity of storing highly entangled quantum states.

In this work, we focus on systems that require a fully quantum treatment in order to obtain accurate polarization functions. Two techniques have been developed by the muon community for exactly simulating the polarization function. The first relies on exact diagonalization of the muon-environment Hamiltonian. Because the thermal energies encountered in μ^+ SR experiments are typically much larger than the nuclear energy levels, the environment is normally assumed to be in the maximally mixed state $\rho_e(0) = I_e/D_e$ where D_e is the Hilbert-space dimension of the environment. The polarization function can then be obtained by

$$\begin{aligned} P(t) &= \frac{1}{D_s} \text{Tr}([Z_\mu \otimes I_e] e^{-iH_s t} [|0\rangle \langle 0|_\mu \otimes I_e] e^{iH_s t}), \\ &= \frac{1}{2D_s} \text{Tr}\{[Z_\mu \otimes I_e] e^{-iH_s t} [(I_\mu + Z_\mu) \otimes I_e] e^{iH_s t}\}, \\ &= \frac{1}{2D_s} \text{Tr}([Z_\mu \otimes I_e] e^{-iH_s t} [Z_\mu \otimes I_e] e^{iH_s t}). \end{aligned} \quad (4)$$

We then use a resolution of the identity in terms of eigenstates of the system Hamiltonian; $|m\rangle, |n\rangle$

$$\begin{aligned} P(t) &= \frac{1}{2D_s} \sum_m \langle m| Z_\mu e^{-iH_s t} Z_\mu e^{iH_s t} |m\rangle, \\ &= \frac{1}{2D_s} \sum_m \sum_n \langle m| Z_\mu e^{-iH_s t} |n\rangle \langle n| Z_\mu e^{iH_s t} |m\rangle, \\ &= \frac{1}{2D_s} \sum_m \sum_n |\langle m| Z_\mu \otimes I_e |n\rangle|^2 e^{i(E_m - E_n)t}, \end{aligned} \quad (5)$$

where E_m is the eigenvalue of eigenstate $|m\rangle$. As a result, we can see that performing an exact diagonalization of the system Hamiltonian can be used to calculate the polarization function. However, the dimension of the system Hamiltonian will scale exponentially with the number of spins considered, which limits the size of these calculations. These exact calculations have been performed for a range of systems [7–10,23,27–29], with Hilbert-space dimensions up to around 2048 [9,10]. Recently, a method was developed to scale the interactions between the muon and more distant nuclei, in order to act as a proxy for the remaining nuclei in the sample, which showed promising results for muon experiments on CaF₂ and NaF [9].

An alternative method, developed by Celio [30], reduces the memory cost of the calculation, at the expense of introducing statistical uncertainty into the measured result, which can be reduced through sampling. This method uses a first-order product-formula-based approach, whereby the time-evolution operator is divided into a product of operators acting on subsystems of the system (this is also referred to as ‘‘Trotterization’’ by the quantum-computing community). When the environment is initially in a pure state $|\phi(0)\rangle_e$, the wavefunction of the system at time t is given by

$$\begin{aligned} |\psi_s(t)\rangle &= e^{-iH_s t} |0\rangle_\mu |\phi(0)\rangle_e, \\ &\approx \left(\prod_\alpha e^{-iH_s^\alpha t/n} \right)^n |0\rangle_\mu |\phi(0)\rangle_e, \\ &\equiv U_1(t) [|0\rangle_\mu |\phi(0)\rangle_e], \end{aligned} \quad (6)$$

where H_s^α are subterms in the Hamiltonian (for example, the interaction between a muon and a single nearest-neighbor spin), and n is referred to as the number of Trotter steps used. Equality is recovered in the limit that $n \rightarrow \infty$. Celio [30] made use of a random-phase approximation-inspired method, which sets the wave function at time t is as

$$|\psi_s(t)\rangle = U_1(t) \left[\frac{1}{\sqrt{D_e}} \sum_j e^{i\theta_j} |0\rangle_\mu |j\rangle_e \right], \quad (7)$$

where θ_j are randomly chosen on each sample of the algorithm, and $|j\rangle_e$ denote unentangled basis states of the environment. In this case, an approximation for the polarization function can be obtained from

$$\begin{aligned} P(t) &\approx \langle \psi_s(t) | Z_\mu \otimes I_e | \psi_s(t) \rangle, \\ &= \sum_j \frac{1}{D_e} \langle 0|_\mu \langle j|_e U_1(t)^\dagger Z_\mu U_1(t) |0\rangle_\mu |j\rangle_e, \\ &\quad + \sum_{j \neq k} \frac{1}{D_e} e^{i(\theta_k - \theta_j)} \langle 0|_\mu \langle j|_e U_1(t)^\dagger Z_\mu U_1(t) |0\rangle_\mu |k\rangle_e. \end{aligned} \quad (8)$$

The first term is equal to the polarization function that we wish to measure (up to an error induced by the Trotterization of the time-evolution operator). As the phases in the second term are chosen randomly, many of these terms cancel, leading to a small error on the polarization function. This error can be reduced by repeating the method with independently randomly generated values of θ_j . The error also decreases as the Hilbert-space dimension of the system increases (see Appendix C 2). The Trotter error can be reduced by using higher-order product formulae, or by increasing the number of Trotter steps used. We discuss this source of error in more detail in Sec. IV. While the exact diagonalization method requires manipulating matrices of dimension $D_s \times D_s$, this product-formula-based approach requires only storing the wave function, which has dimension D_s . This has enabled much larger calculations to be performed using this method, including simulations with Hilbert-space dimension approximately 2^{17} [31] and 2^{26} [32]. This latter calculation is, to the best of our knowledge, the largest μ^+ SR calculation performed to date with an exact method.

Exact methods for calculating the polarization function have predominantly been used for two purposes: locating the muon rest site, and studying muon diffusion. Determining the muon rest site is an important challenge in μ^+ SR experiments [33]. If we are able to accurately simulate the muon polarization function, then the polarization functions arising from a number of candidate sites can be generated, and compared to the experimental data, in order to determine the most likely muon location. This technique can be employed in conjunction with density-functional theory [34–38] and experimental methods [39] to give greater certainty on the muon location. Exact simulation has repeatedly been employed when studying fluorinated materials, due to the strong dipolar interaction between the spin-1/2 fluorine nuclei and the muon, and the large electronegativity of the fluorine ion, which “traps” the muon [40]. This effect has been used to locate the muon rest site in fluorinated polymers [29,41], molecular magnets [28], and ionic crystals [9,42]. Similar calculations were used to identify the muon rest sites in the high-temperature superconductor $\text{La}_{2-x}\text{Sr}_x\text{CuO}_4$ [32]. The calculations have often included heuristic terms in the polarization function, to compensate for the limited system sizes that can be simulated using these costly techniques.

Polarization functions calculated from quantum models have also been used to study muon diffusion. As the muon is a light particle, its diffusion between lattice sites is an inherently quantum phenomenon, involving quantum tunneling through the potential barrier between sites. Calculating the hopping amplitudes from first principles would be an exceedingly costly calculation, requiring an accurate description of the electronic structure of the muon-environment system. Fortunately, diffusion of the muon between lattice sites leads to a measurable change in

the polarization function. We can calculate the muon diffusion rate by comparing the polarization function obtained in the absence of diffusion (which can be measured at low temperatures) to the polarization function when diffusion is present. The effect of muon diffusion is typically incorporated into theoretical calculations using the “strong collision model” (SCM). The SCM assumes Markovian, stochastic hopping of the muon between lattice sites. The muon is considered to probabilistically jump to a new site, where the surrounding nuclei are in the maximally mixed state. This leads to a damping of the polarization function, the strength of which depends on the hopping rate. Dalmás de Réotier *et al.* [31] showed that using the mean-field Kubo-Toyabe function in conjunction with the SCM can lead to inaccurate results for the muon hopping strength. More accurate results were obtained by first computing a static polarization function using the product-formula method described above, and then augmenting this with the SCM. Kadono *et al.* [43] and Luke *et al.* [44] used the product-formula-based method in conjunction with the SCM to analyze the diffusion of muons in copper. They noted that using the Kubo-Toyabe function in conjunction with the SCM would lead to an overestimation of the muon diffusion rate.

It is known that the strong collision model is of limited accuracy. The assumption that the muon moves to a new site with unpolarized nuclei means that it cannot move to a nearest-neighbor lattice site (as the muon would have altered the polarization of the nuclei shared between its previous and new sites), or hop back to its previous site. However, quantum tunneling is a non-Markovian process, and is strongly affected by the presence of phonons in the system. These phonons can lead to a degeneracy between the energies of different sites, increasing the transition probability. In particular, this effect increases the likelihood of the muon returning to its previous position [7]. A more accurate model for muon diffusion was developed by Celio [22]. A rigorous equation for the dynamic polarization function with a single jump is given by

$$\begin{aligned} G(t) &= G_0(t)e^{-\nu t} + \nu \int_0^t dt' e^{-\nu t'} F(t, t'), \\ G_0(t) &= \text{Tr}[e^{-iH_0 t} \rho_s(0) e^{iH_0 t} Z_\mu], \\ F(t, t') &= \text{Tr}[e^{-iH_1(t-t')} e^{-iH_0 t'} \rho_s(0) e^{iH_0 t'} e^{iH_1(t-t')} Z_\mu] \end{aligned} \quad (9)$$

where H_0 is the Hamiltonian at the initial muon site, H_1 is the Hamiltonian at the new site, t' denotes the time of the hop, and ν is the hopping rate. The trace is over all nuclear sites involved in both Hamiltonians. This equation can be generalized to a larger number of jumps. Celio [22] showed that the inferred hopping rate can differ by around 20% when using this rigorous approach, compared to using the SCM. To

the best of our knowledge, this model has never been employed when analyzing experimental data, likely due to the large computational cost of considering many different sites.

Having established the scenarios where quantum models are typically utilized in muon-spectroscopy analysis, we now consider the interactions present in such systems. The muon-system Hamiltonian typically includes contributions from dipolar interactions between spins, quadrupole interactions between nuclear spins and electric field gradients in the sample, and the coupling of spins to time-dependent or independent magnetic fields. The dipolar contribution is given by

$$H_D = \frac{1}{2} \sum_{ij} \frac{\hbar^2 \mu_0 \gamma_i \gamma_j}{4\pi r_{ij}^3} \left[\vec{S}_i \cdot \vec{S}_j - 3(\vec{S}_i \cdot \hat{r}_{ij})(\vec{S}_j \cdot \hat{r}_{ij}) \right], \quad (10)$$

where μ_0 is the permeability of free space, γ_i is the gyromagnetic ratio of spin i , \vec{r}_{ij} is the vector connecting spins i and j , and \vec{S}_i are the vectorized generalized spin matrices for a particle with spin quantum number s . We move the factor of \hbar from the spin matrices to the coefficient of the sum. For example, for a spin-1/2 particle we have $\vec{S}_i = (1/2)(X_i, Y_i, Z_i)$, where X_i, Y_i, Z_i are the Pauli matrices acting on spin i . These can be generalized to spins of higher dimension, which we discuss in more detail in Sec. IV A. This Hamiltonian contains $O(N^2)$ terms, where N is the number of spins considered in the simulation.

There is also an interaction between the quadrupole moment of nuclei with $s > (1/2)$, and any nonzero electric field gradients in the sample. Such electric field gradients can often be induced by the presence of the muon. The quadrupole interaction Hamiltonian is given by [9]

$$H_Q = \sum_{i \in \mathcal{Q}} \frac{\hbar e Q_i (1 + \Gamma_i)}{2s_i(2s_i - 1)} \left(\vec{S}_i^T \cdot \mathbf{G}(\vec{r}_i) \cdot \vec{S}_i \right), \quad (11)$$

where \mathcal{Q} is the set of nuclei in the simulation with quadrupole moments, e is the electron charge, Q_i and Γ_i are the quadrupole coupling factor and antishielding factor (respectively) of the i th spin, and $\mathbf{G}(\vec{r}_i)$ is the electric field gradient tensor at position \vec{r}_i . The elements of $\mathbf{G}(\vec{r}_i)$ are $G_{\alpha\beta}(\vec{r}_i) = [\partial^2 V(\vec{r}_i) / \partial r_\alpha \partial r_\beta]$, where $V(\vec{r}_i)$ is the Coulomb potential at position \vec{r}_i . This Hamiltonian contains $O(N_{\mathcal{Q}})$ terms, each acting on a single spin (where $N_{\mathcal{Q}}$ is the number of nuclei with quadrupole moments).

The Zeeman interactions of each spin with an applied magnetic field are given by

$$H_M(t) = \sum_i \hbar \gamma_i \vec{S}_i \cdot \vec{B}(t), \quad (12)$$

where $\vec{B}(t)$ is the magnetic field.

Some previous calculations have neglected the dipolar interactions between the nuclear spins, as they are often much smaller than those between the muon and the nuclear spins [30,32]. Neglecting these interactions reduces the number of terms in the Hamiltonian to $O(N)$.

In this section, we discuss theoretical methods to generate simulated polarization functions. The accuracy of these polarization functions depends on both the model for the system (e.g., whether effects like muon diffusion are taken into account), and the method used to solve the model. The most accurate calculations require a fully quantum model of the system. Unfortunately, the cost of classically simulating the dynamics of a quantum system is expected to scale exponentially with the size of the simulated system. This high cost has restricted the simulation of μ^+ SR experiments to small systems, with Hilbert-space dimensions of less than 2^{30} . In Sec. IV we show how this exponential cost can be circumvented by using quantum hardware as the simulation platform.

III. QUANTUM COMPUTING

Quantum computing leverages quantum-mechanical effects in order to perform certain information processing tasks more efficiently than appears possible with classical computers. Quantum computers were originally proposed as efficient simulators of other quantum systems [45], and were later formalized as a more general computational device [46]. Since these initial proposals, a number of quantum algorithms have been developed, which appear to asymptotically outperform known classical algorithms. These include algorithms for factoring numbers [47], searching databases [48], and simulating quantum systems [49]. There has also been significant progress in developing platforms on which to run these algorithms. Physical qubits can be created in a number of different systems [50]. These include using the two lowest energy levels of superconducting circuit resonators [51,52], path or polarization degrees of freedom in linear optical photonic systems [53], or a pair of energy levels in ions [54,55]. Quantum computing is currently referred to as being in the “noisy, intermediate-scale quantum” (NISQ) era [56]. Current quantum computers possess up to around 50 qubits, and exhibit gate error rates of approximately 10^{-3} , at best [57]. These limited computational resources mean that we have so far been unable to run classically challenging instances of the algorithms referenced above. In this section, we provide an introduction to quantum computing. We refer the interested reader to the textbook by Nielsen and Chuang [58] for more information.

In this work, we focus on the circuit model of quantum computation. This formalism abstracts away the physical details of the hardware implementation, and treats the qubits as generic two-level systems. The computational

basis states of the qubit Hilbert space are taken to be

$$|0\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad |1\rangle = \begin{bmatrix} 0 \\ 1 \end{bmatrix}. \quad (13)$$

A general single qubit state is described by

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle, \quad (14)$$

where α, β are complex amplitudes, and the state is normalized to one. Qubits are manipulated by applying quantum logic gates to the system. These gates are defined by unitary matrices acting on the qubit wave function. Typical single-qubit gates include the Pauli gates

$$X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad Y = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, \quad Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \quad (15)$$

the single-qubit rotation gates

$$R_x(\theta) = e^{\left(\frac{-i\theta X}{2}\right)}, \quad R_y(\theta) = e^{\left(\frac{-i\theta Y}{2}\right)}, \quad R_z(\theta) = e^{\left(\frac{-i\theta Z}{2}\right)} \quad (16)$$

and the Hadamard and T gates

$$H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \quad T = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{bmatrix}. \quad (17)$$

A system of Q qubits is described by a normalized wave function in the Hilbert space of dimension 2^Q . We can construct wave functions that cannot be decomposed into tensor products of individual qubits by using multiqubit entangling gates, such as the controlled NOT (CNOT) gate

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}. \quad (18)$$

A quantum circuit consists of a sequence of unitary gates, applied to a well-defined initial state, such as $|\bar{0}\rangle = |0\rangle \cdots |0\rangle$. Qubits can be measured in the computational basis, either at an intermediate stage of the computation, or at the end of the calculation. When a qubit is measured in the computational basis, it will “collapse” to state $|0\rangle$ or state $|1\rangle$, with probability $|\alpha|^2$ and $|\beta|^2$, respectively.

While some quantum algorithms (such as Ref. [48]) extract their solution from a single measurement of the qubits, other algorithms require measurements of observables, O , which are represented by Hermitian matrices. For algorithms in the latter category, we typically seek the average value of an observable over many measurements, $\bar{O} = \langle \psi | O | \psi \rangle$. In this work, we are predominantly concerned with measuring the spin polarization of a qubit representing the muon. To measure the polarization along

the Z axis, we can simply prepare the desired state, and measure the qubit representing the muon in the computational basis. We assign a value of $+1$ to the outcome $|0\rangle$, and -1 to the outcome $|1\rangle$. We repeat this process a number of times, and average the results. To measure the polarization along a different axis we use a single-qubit rotation to rotate the measurement axis to the Z axis, and then apply the procedure described above. For example, when measuring X , we use the H (Hadamard) gate to rotate the basis. In Sec. IV D we discuss a more complex, but asymptotically more efficient way of measuring the expectation value of Hermitian observables.

The formalism introduced above can be used to construct and describe quantum algorithms, by specifying the actions applied to individual qubits. These abstract instructions are then converted into physical controls applied to the quantum hardware, such as applying laser pulses to excite a trapped ion qubit. These physical controls are imperfect, introducing noise to the system. Moreover, the quantum states constructed are sensitive to decoherence caused by interaction with the surrounding environment. These processes introduce noise into the system, which can be modeled in a number of ways. The presence of noise can be overcome by using quantum-error-correction codes. These codes construct a single logical qubit from highly entangled states of a number of physical qubits. It is important to note the high overheads introduced by quantum error correction; a common rule of thumb is that around $10^3 - 10^4$ physical qubits per logical qubit may be required to solve classically intractable problem instances [59–61]. We elaborate more on this resource estimation in Sec. V C. We refer the interested reader to Refs. [62–65] for a more comprehensive discussion of quantum error correction.

The limited number of qubits in current NISQ devices precludes our ability to perform quantum error correction. Instead, alternative algorithmic error mitigation strategies have been developed, which effectively seek to identify the noiseless signal from an increased number of noisy experimental repetitions. We refer the interested reader to Ref. [66] for a more detailed discussion of error-mitigation techniques. In this work, we place particular focus on the extrapolation method of error mitigation. Error extrapolation uses expectation values obtained at multiple different physical noise levels to infer the noiseless expectation value [67–71]. We can increase the noise level in the quantum processor in a number of ways, including “stretching” the duration of gates [72], or by converting the noise into a Pauli channel and then introducing additional Pauli errors with the appropriate probabilities [67].

IV. QUANTUM SIMULATION OF μ^+ SR

In this section, we illustrate how quantum computers can be used to analyze muon-spectroscopy data.

We show that a quantum computer is able to simulate muon polarization functions using resources scaling polynomially with the size of the simulated system. This

is in contrast to the classical methods for simulating muon polarization functions described in Sec. II, which require exponentially scaling resources. Our algorithm

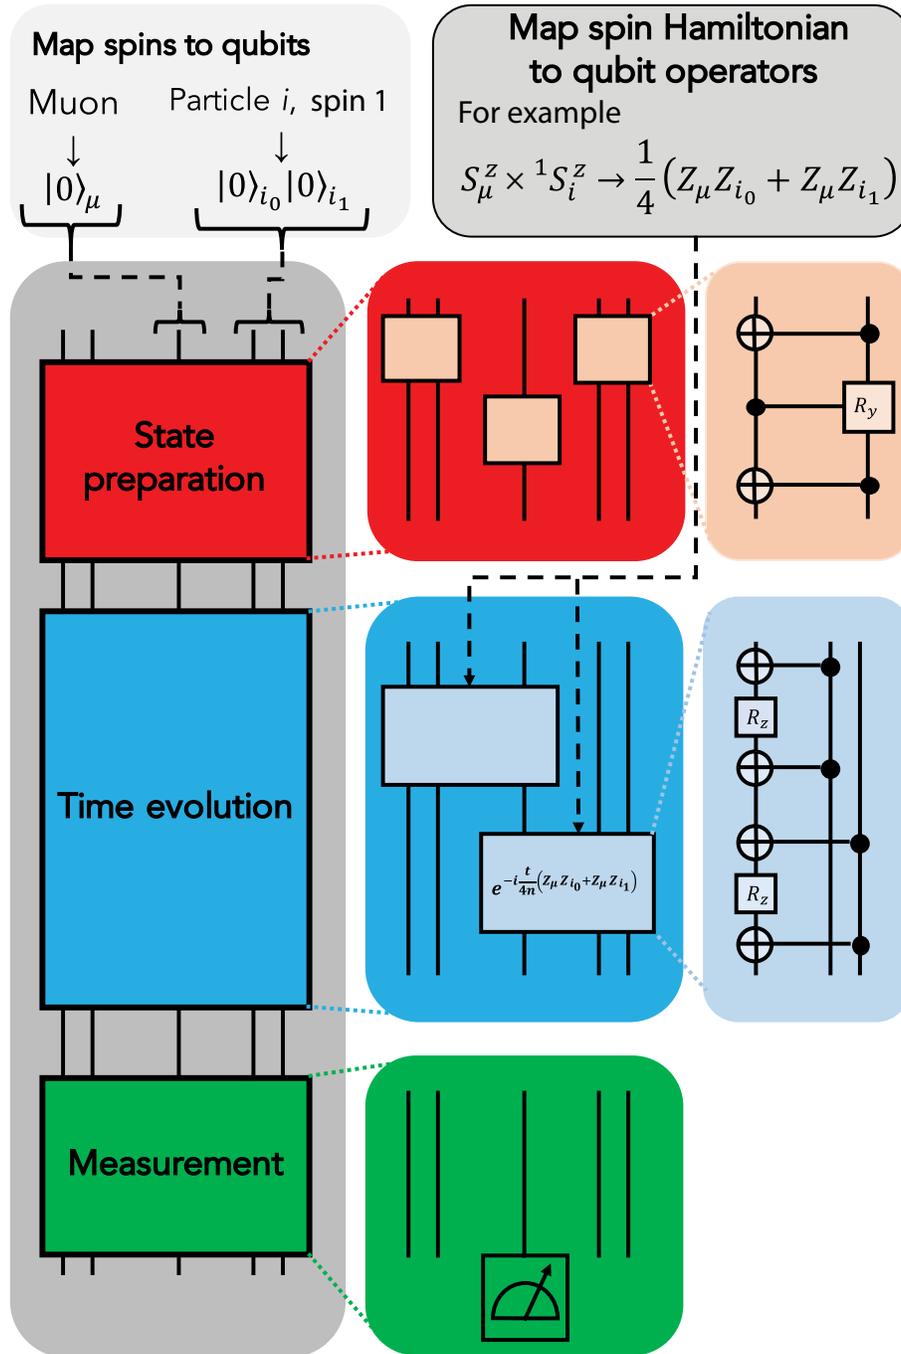


FIG. 2. Outline of the proposed algorithm for simulating muon polarization functions on quantum computers. We first map the spins onto qubits via the mapping in Sec. IV A, which also enables us to map the spin Hamiltonian to qubit operators. We prepare the desired initial state on these qubit registers using the techniques described in Sec. IV B. We can then evolve this state in time, using the methods in Sec. IV C. We can measure the polarization value of the muon qubit, as described in Sec. IV D. This diagram shows the most simple form of the algorithm. This circuit must be repeated many times in order to estimate the Z expectation value of the muon at a single time value. We must repeat this procedure for each desired time value. This will generate a simulated polarization function, which can be incorporated into, for example, an optimization routine to determine the muon rest site, by fitting the generated data to experimental data.

is illustrated in Fig. 2, and can be summarized as follows:

1. Map the spin system of interest to qubits.
2. Prepare the quantum registers in the desired initial state.
3. Evolve the system in time for the desired duration, t .
4. Measure the muon Z expectation value, $P(t)$.

We can repeat steps 2–4 of this process at many different t values in order to obtain a simulated version of the polarization function for the given system. Below, we discuss in more detail how each of these steps can be implemented.

A. Mapping the spin system to qubits

We first need to map the spin system under consideration onto the quantum computer, such that the spin states of the system correspond to valid states of the quantum register. For the case of spin-1/2 particles, this mapping takes a simple form

$$\begin{aligned} \left| \frac{1}{2} \right\rangle &:= |0\rangle, \\ \left| -\frac{1}{2} \right\rangle &:= |1\rangle. \end{aligned} \quad (19)$$

This leads to a natural mapping of the spin operators of particle i

$$\begin{aligned} S_i^x &= \frac{1}{2} X_i, \\ S_i^y &= \frac{1}{2} Y_i, \\ S_i^z &= \frac{1}{2} Z_i. \end{aligned} \quad (20)$$

When considering particles with higher spin s , there are a number of possible mappings that we could use. One approach is to store the spin of the particle in a register with $\log_2(2s+1)$ qubits [73]. However, this compact mapping can lead to more complicated expressions for the spin operators, and may be problematic if the spin multiplicity is not a power of 2. In this work, we employ a spin-to-qubit mapping based on symmetric quantum states of a given Hamming weight (known as Dicke states) [74]. For example, a spin-1 particle is stored as

$$\begin{aligned} |1\rangle &:= |00\rangle, \\ |0\rangle &:= \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle), \\ |-1\rangle &:= |11\rangle. \end{aligned} \quad (21)$$

This mapping can be derived by considering the joint Hilbert space of two spin-1/2 particles, which can be partitioned into a spin-1 triplet space, and a spin-0 singlet space. This mapping is an embedding of the three-dimensional triplet space into the four-dimensional two-qubit Hilbert space. The spin operators for particle i are given by

$$\begin{aligned} S_i^x &= \frac{1}{2}(X_{i_0} + X_{i_1}), \\ S_i^y &= \frac{1}{2}(Y_{i_0} + Y_{i_1}), \\ S_i^z &= \frac{1}{2}(Z_{i_0} + Z_{i_1}). \end{aligned} \quad (22)$$

We can generalize this mapping for a spin- s particle by assigning each of the states $|j\rangle$, $j \in [-s, \dots, s]$ to the symmetric superposition of states with Hamming weight h such that $h = s - j$. For example, a spin-3/2 particle can be represented by

$$\begin{aligned} \left| \frac{3}{2} \right\rangle &:= |000\rangle, \\ \left| \frac{1}{2} \right\rangle &:= \frac{1}{\sqrt{3}}(|001\rangle + |010\rangle + |100\rangle), \\ \left| -\frac{1}{2} \right\rangle &:= \frac{1}{\sqrt{3}}(|011\rangle + |101\rangle + |110\rangle), \\ \left| -\frac{3}{2} \right\rangle &:= |111\rangle. \end{aligned} \quad (23)$$

The spin operator for a spin- s particle i is given by

$${}^s S_i^\alpha = \frac{1}{2} \sum_{j=0}^{2s-1} P_{i_j}^\alpha, \quad (24)$$

where $\alpha \in [x, y, z]$ denotes which of the Pauli matrices $P^\alpha \in [X, Y, Z]$ are used, and j denotes the qubits in register i . For example, the S^x operator on particle $i = 1$ with spin $s = (3/2)$ is given by

$$\frac{3}{2} S_1^x = \frac{1}{2}(X_{1_0} + X_{1_1} + X_{1_2}). \quad (25)$$

The total number of qubits required is given by $\sum_{i=1}^N 2s_i$, where s_i is the spin of the i th particle, and N is the number of particles in the system. For example, for six spin-3/2 particles, plus a muon, we need $[6 \times 2 \times (3/2)] + [1 \times 2 \times (1/2)] = 19$ qubits. We can use this mapping to obtain qubit representations of the system Hamiltonians given in Eqs. (10)–(12). The dipole interaction is the dominant contribution to the number of terms in the Hamiltonian. Each term ${}^s S_i^\alpha \otimes {}^{s'} S_j^\beta$ is mapped to $4ss'$ two-qubit Pauli terms. As a result, the Hamiltonian contains up to $\mathcal{O}(N^2 s_{\max}^2)$ two-qubit Pauli terms, where s_{\max} is the largest spin value in the system.

B. Preparing the initial state of the system

As discussed in Sec. II, the initial state of the system is typically taken as $\rho_0 = |0\rangle\langle 0|_\mu \otimes (I_e/D_e)$. There are two routes to effectively time evolve this mixed initial state. The first approach is to emulate nature; we can prepare the environment register in a state $|k\rangle_e = \bigotimes_{i \in e} |j_i\rangle_i$ that is the tensor product of each environment spin in a randomly selected s_z basis state. Each state $|k\rangle_e$ is chosen with probability $p_k = (1/D_e)$. By repeating the simulation many times, we obtain

$$\begin{aligned} & \sum_k p_k \text{Tr} \left[(Z_\mu \otimes I_e) e^{-iHt} (|0\rangle\langle 0|_\mu \otimes |k\rangle\langle k|_e) e^{iHt} \right] \\ &= \text{Tr} \left[(Z_\mu \otimes I_e) e^{-iHt} \left(|0\rangle\langle 0|_\mu \otimes \frac{I_e}{D_e} \right) e^{iHt} \right], \end{aligned} \quad (26)$$

as required. This method converges as $1/\sqrt{\omega}$, where ω is the number of samples taken.

The second approach to effectively sample from the time-evolved initial mixed state is a modified version of the random-phase-approximation method [30] discussed in Eqs. (7) and (8). We first initialize the environment register in an equal superposition of all possible $|k\rangle_e$ states. For the case of spin-1/2 particles, this can be accomplished by applying a Hadamard gate to each of the qubits. For the case of higher-spin particles, we discuss below how to construct this superposition. These steps will prepare the state

$$|\psi\rangle = \frac{1}{\sqrt{D_e}} \sum_k |0\rangle_\mu |k\rangle_e, \quad (27)$$

which can be compared with the state given in Eq. (7). In order to generate the desired random phase for each basis state $|k\rangle_e$, we apply random single-particle $R_z(\theta)$ rotations to each particle. This procedure will not generate a state with completely independent phases for each basis state. If it is necessary to further randomize the state, we can apply layers of R_z and controlled- R_z rotations between the different particles. We average the results of several simulations (each with a different set of randomly chosen θ_i) to obtain the polarization function, as shown by Eq. (8). While the asymptotic convergence properties of the two methods are the same, we expect the latter method to yield a smaller error for a given number of samples. As shown in Appendix C, the variance of the first scheme is upper bounded by 1, and examples can be found that saturate this bound. In contrast, the variance of the second scheme is upper bounded by $1/D_e$, and so decreases exponentially with the number of environment spins considered. The smaller sampling error of the second method is shown numerically in Fig. 7.

For both of the methods discussed above, we must construct the states $|k\rangle_e$, either alone, or in superposition.

These states are tensor products of the environment spins each in a arbitrary state $|j\rangle$, that is an s_z eigenstate. As discussed in Sec. IV A, we map these spin states onto qubit Dicke states. Efficiently constructing Dicke states (or superpositions of Dicke states) on quantum computers has remained challenging for a number of years [75–78], but has recently been made possible with the elegant inductive solution of Bartschi and Eidenbenz [79]. Their algorithm can be summarized as follows, and is explained in more detail in Appendix B. We first define a Dicke state with Hamming weight h , on n qubits, as $|D_h^n\rangle$. For example, the $|-1/2\rangle$ state in Eq. (23) equates to $|D_2^3\rangle$. We can then observe that

$$|D_h^n\rangle = \sqrt{\frac{h}{n}} |D_{h-1}^{n-1}\rangle \otimes |1\rangle + \sqrt{\frac{n-h}{n}} |D_h^{n-1}\rangle \otimes |0\rangle. \quad (28)$$

We assume the existence of a unitary operator $U_{n,k}$ such that $U_{n,k} |0\rangle^{\otimes n-h} |1\rangle^{\otimes h} = |D_h^n\rangle$ for all $h \leq k$. Through induction, it can be shown that this unitary operator exists, and how to construct it from typical single- and two-qubit gates. For example, note that

$$|D_h^n\rangle = U_{n,k} |0\rangle^{\otimes n-h} |1\rangle^{\otimes h}, \quad (29)$$

and

$$\begin{aligned} |D_h^n\rangle &= \sqrt{\frac{h}{n}} |D_{h-1}^{n-1}\rangle \otimes |1\rangle + \sqrt{\frac{n-h}{n}} |D_h^{n-1}\rangle \otimes |0\rangle, \\ &= U_{n-1,k} \left[\sqrt{\frac{h}{n}} |0\rangle^{\otimes n-h} |1\rangle^{\otimes h} \right. \\ &\quad \left. + \sqrt{\frac{n-h}{n}} |0\rangle^{\otimes n-1-h} |1\rangle^{\otimes h} |0\rangle \right], \\ &= U_{n-1,k} V_{n,k} |0\rangle^{\otimes n-h} |1\rangle^{\otimes h}, \end{aligned} \quad (30)$$

implying that

$$U_{n,k} = U_{n-1,k} V_{n,k}. \quad (31)$$

As shown in Appendix B, these relations can be repeated recursively, to obtain a circuit purely in terms of the $V_{n,k}$ -type gates. Bartschi and Eidenbenz [79] show how these gates can be constructed from CNOT gates and single-qubit rotations (see also Appendix B). The entire state-preparation circuit has a depth of $\mathcal{O}(n)$, and requires $\mathcal{O}(kn)$ gates, even on a one-dimensional (1D) linear chain of qubits. The algorithm also requires no additional ancillary qubits. Because the unitary $U_{n,k}$ that creates the Dicke states was defined to work for all input states with $h \leq k$, we can use unitary $U_{n,n}$ to create all of the Dicke states

$|D_{h_{\leq n}}^n\rangle$. As a result, if we input the superposition

$$\sqrt{\frac{1}{n+1}} \sum_{k=0}^n |0\rangle^{\otimes n-k} |1\rangle^{\otimes k} \quad (32)$$

then we obtain an equally weighted superposition of all of the possible Dicke states. If we apply this approach to each environment spin, then this will generate the state shown in Eq. (27). It is straightforward to generate the state in Eq. (32), using a ladder of control- R_y gates, as shown in Ref. [79]. As a result, a similar number of gates are required for both of the initial state-preparation methods described in this section.

C. Evolving the state in time

Once we generate the desired initial state, we can then evolve it in time for the desired duration. There are a number of approaches that can be used for performing time evolution on quantum computers [67,80–83]. In this work, we focus on product-formula-based approaches (also referred to as “Trotterization”), which are observed to lead to lower gate counts than some asymptotically more efficient algorithms when considering the time evolution of spin systems [84]. Trotter-based methods implement time evolution by dividing the propagator into a product of time evolutions under subterms in the Hamiltonian, that have known decompositions into single- and two-qubit gates. Detailed derivations of the error resulting from Trotterization are given in Ref. [85], and we report some of that work’s key results below. We decompose the Hamiltonian into a sum of Pauli strings, as $H = \sum_{\alpha} h_{\alpha} H_{\alpha}$, where h_{α} are real coefficients. A first-order Trotter decomposition of the time evolution operator is given by

$$U_1(t) := \left(\prod_j e^{-\frac{it}{n} h_j H_j} \right)^n, \quad (33)$$

where n is the number of Trotter steps used. The error in the first-order product formula is upper bounded by

$$\|U(t) - U_1(t)\| := \epsilon \sim \mathcal{O}\left(\frac{(L\Lambda t)^2}{n}\right), \quad (34)$$

where $U(t)$ is the true time-evolution operator, $\|\cdot\|$ denotes the spectral norm, L is the number of terms in H , and $\Lambda = \max_{\alpha} \|h_{\alpha}\|$. A tighter bound on the error due to first-order Trotterization is given by

$$\epsilon \sim \frac{t^2}{2n} \sum_i \left\| \sum_{j>i} [h_j H_j, h_i H_i] \right\|, \quad (35)$$

which is known to be tight, up to an application of the triangle inequality. We can also consider higher-order product formulae, which yield improved error scaling. The

second-order product formula is given by

$$U_2(t) := \left(\prod_{j=1}^L e^{-\frac{it}{2n} h_j H_j} \prod_{k=L}^1 e^{-\frac{it}{2n} h_k H_k} \right)^n. \quad (36)$$

An upper bound on the second-order Trotter error is given by

$$\|U(t) - U_2(t)\| := \epsilon \sim \mathcal{O}\left(\frac{(L\Lambda t)^3}{n^2}\right), \quad (37)$$

and a tighter bound (tight up to an application of the triangle inequality) is given by

$$\begin{aligned} \epsilon \sim \frac{t^3}{n^2} & \left[\frac{1}{12} \left(\sum_i \left\| \sum_{j>i} \sum_{k>i} [h_k H_k, [h_j H_j, h_i H_i]] \right\| \right) \right. \\ & \left. + \frac{1}{24} \left(\sum_i \left\| \sum_{j>i} [h_i H_i, [h_i H_i, h_j H_j]] \right\| \right) \right]. \quad (38) \end{aligned}$$

It has also been shown that introducing aspects of randomized compilation into these algorithms can lower the gate counts required to obtain results of a given accuracy. For example, it has been shown that randomly permuting the ordering of the product-formula terms in each step can reduce the error scaling obtained [86]. An alternative randomized procedure, known as qDRIFT [83], probabilistically selects a number of terms from the Hamiltonian according to their strength, and then evolves under these terms (each for an equal duration in time). A single step of qDRIFT is equivalent to implementing the channel

$$\mathcal{E}(\rho) = \sum_j \frac{h_j}{\lambda} e^{-i\frac{\lambda}{N} H_j} \rho e^{i\frac{\lambda}{N} H_j}, \quad (39)$$

where $\lambda = \sum_j h_j$ is the 1 norm of the Hamiltonian (when applying qDRIFT, we shift the signs from the h_j coefficients to the H_j operators, such that h_j are all real and positive), and N is the number of qDRIFT steps used in this simulation. The error scaling of qDRIFT is given by

$$\epsilon \sim \frac{\lambda^2 t^2}{N}. \quad (40)$$

Most notably, the error scaling of qDRIFT is independent of the number of terms in the Hamiltonian, making it an interesting candidate for systems with a large number of weakly interacting terms. Given the rapid power-law decay of the dipolar interactions studied in this work, qDRIFT may be an interesting candidate for simulating muon spectroscopy on quantum computers. While a thorough numerical comparison of these approaches beyond the scope of this work, it would be an interesting area for future study—especially if methods that

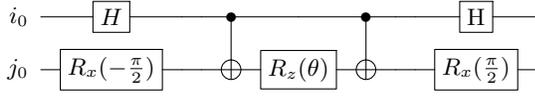


FIG. 3. A quantum circuit implementing the term $\exp(-i\frac{\theta}{2}X_{i_0}Y_{j_0})$.

interpolate between Trotterization and qDRIFT are also considered [87].

The aim of our simulation algorithm is to obtain an accurate estimate for the value of the muon polarization function at a given time. As a result, we are not interested in the Trotter error ϵ directly, but in the error it induces on $\text{Tr}[Z_\mu \rho(t)]$. We expect that the Trotter error will provide a loose upper bound for this error, as has been seen previously [85,88].

Each of the operators of the form $\exp(-ih_\alpha H_\alpha t)$ can be decomposed into a sequence of single- and two-qubit gates. We illustrate this by considering a term in the dipolar Hamiltonian for spin-1 particles. A single term, such as ${}^1S_i^x \otimes {}^1S_j^y$ can be mapped to

$$\begin{aligned} e^{-i\theta} {}^1S_i^x \otimes {}^1S_j^y &= e^{-i\frac{\theta}{2}(X_{i_0}+X_{i_1}) \otimes (Y_{j_0}+Y_{j_1})}, \\ &= e^{-i\frac{\theta}{2}X_{i_0}Y_{j_0}} e^{-i\frac{\theta}{2}X_{i_0}Y_{j_1}} e^{-i\frac{\theta}{2}X_{i_1}Y_{j_0}} e^{-i\frac{\theta}{2}X_{i_1}Y_{j_1}}. \end{aligned} \quad (41)$$

There is no Trotter error arising from this decomposition, as all of the terms in the exponential commute with each other. In Fig. 3 we show a quantum circuit that performs time evolution under one of the exponentials in this term.

As discussed above, the dipolar Hamiltonian can be decomposed into up to $L \sim \mathcal{O}(N^2 s_{\max}^2)$ two-qubit Pauli terms, where s_{\max} is the largest spin value in the system. Considering a second-order product formula, the number of Trotter steps required to obtain Trotter-error ϵ is upper bounded by

$$n \sim \frac{N^3 s_{\max}^3 \Lambda^{\frac{3}{2}} t^{\frac{3}{2}}}{\epsilon^{\frac{1}{2}}}. \quad (42)$$

Each second-order Trotter step requires $\mathcal{O}(L)$ gates to implement, resulting in a total gate count that scales at worst as

$$G \sim \frac{N^5 s_{\max}^5 \Lambda^{\frac{3}{2}} t^{\frac{3}{2}}}{\epsilon^{\frac{1}{2}}}. \quad (43)$$

We note that this scaling is obtained without considering any possible compilations, and by using the loose second-order Trotter error bounds. As a result, it is likely that this scaling estimate could be tightened significantly when considering a real system of interest. Similar results were obtained in the context of quantum chemistry calculations. While initial estimates suggested a scaling of $\mathcal{O}(N^{11})$ [89], more careful analysis reduced the asymptotic scaling to $\mathcal{O}(N^{5.5})$ [90].

D. Measuring the polarization

The final stage of the algorithm consists of measuring the Z expectation value of the muon qubit. The most straightforward way to measure this value is described in Sec. III. We repeatedly prepare the desired state at time t , measure the muon qubit in the computational basis, and average the results. The standard error in the estimate obtained is given by

$$\alpha = \frac{\sqrt{1 - \bar{Z}_\mu^2}}{\sqrt{M}}, \quad (44)$$

where M is the number of samples taken. In order to obtain an error rate of 10^{-3} with this approach, we require up to 10^6 samples. The repetition rate of a quantum processor can depend on a number of factors, including the depth of the circuit, and the speed of executing quantum gates. The speed of implementing logic gates in a quantum computer depends on the hardware considered, and can range from 10–100 MHz in superconducting qubits [91] to 10 KHz–1 MHz in trapped ion qubits [92]. As a result, even if a gate depth of only 10^3 was required for the circuit, obtaining an estimate of \bar{Z}_μ to a precision of 10^{-3} would take at least 10 s to calculate on a superconducting qubit processor (in this estimate we neglect the time taken for qubit readout and initialization, which can often be longer than the gate time in superconducting qubits [91]). These estimates become even more costly if the overhead of quantum error correction is taken into account. When performing quantum error correction, the logical gate speed of the quantum computer depends on the time taken to measure and decode the error syndromes of all of the physical qubits. This has previously been assumed to be on the order of $10 \mu\text{s}$ [93], which would lead to an estimate of around 2.8 h to measure \bar{Z}_μ to a precision of 10^{-3} using the method described above. This is too slow for our purpose of analyzing polarization curves consisting of hundreds of data points. While this direct sampling method has a time cost of $\mathcal{O}(1/\epsilon^2)$, there are alternative quantum algorithms, which can reduce the time cost to $\mathcal{O}(1/\epsilon)$. These techniques rely on a combination of quantum amplitude amplification and phase estimation [94,95]. These approaches use a constant number of samples, but require a circuit depth of $\mathcal{O}(d_U/\epsilon)$, where d_U is the circuit depth required to implement the time-evolution circuit $U(t)$. This is achieved by repeatedly evolving the register under a controlled version of the operator

$$\begin{aligned} \Omega &:= \Omega_1 \Omega_2, \\ \Omega_1 &= U(t)[I_{e,\mu} - 2|\Phi\rangle\langle\Phi|]U^\dagger(t), \\ \Omega_2 &= Z_\mu \Omega_1 Z_\mu, \end{aligned} \quad (45)$$

where $|\Phi\rangle$ is the initial state of the system. This conditional evolution is controlled by the state of an auxiliary register.

Controlling the evolution on an auxiliary register enables us to perform quantum phase estimation on the unitary Ω , which we can use to extract the value of \bar{Z}_μ to the desired precision. A more detailed discussion of this approach is given in Refs. [94,95].

The steps outlined in this section can be used to obtain a simulated polarization function for the system of interest. In order to use this function to analyze experimental data, the quantum-simulation routine can be incorporated into an optimization loop. As an example, we consider the problem of locating the muon rest site. We first specify the positions of each particle in the system, and use this information to generate the spin and qubit Hamiltonians. We can then use the quantum-simulation routine outlined above to calculate the simulated polarization function. We can quantify the agreement between the experimental polarization data and the simulated polarization function using a suitable loss function, such as χ^2 (the sum of the normalized squared residuals). An optimization loop can then be used to minimize the value of the loss function, by updating the positions of the particles in the system.

V. RESULTS

In order to investigate the practicality of the algorithm discussed in Sec. IV, we perform numerical simulations of systems with up to 29 qubits. We focus on the dipolar interactions between an implanted muon, and spin-1/2 fluorine nuclei in a sample of interest. As discussed in Sec. II, the electronegativity of the fluorine ion (F^-) acts as a “trap” for the positively charged muon. We can use the “fingerprint” left by the F^- - μ^+ dipolar interaction on the polarization function to determine the muon rest site. This technique has been used to locate the muon in a range of systems, including the ionic crystals considered in this work [9,41]. We investigate applying our quantum algorithm to analyze the spectra arising from μ^+ SR experiments on the ionic crystal calcium fluoride (CaF_2), obtained in Refs. [9,96]. CaF_2 provides an ideal test system for the methods introduced in this work for two main reasons. Firstly, the calcium nuclei have spin 0 with an abundance of approximately 99.9%, so we need only to consider the dipolar interaction between the muon and the spin-1/2 fluorine nuclei [9]. Secondly, the recent exact simulation results of Wilkinson and Blundell [9] for this system provide a useful benchmark for our results.

The geometry studied consists of a simple cubic lattice of F^- ions, with a lattice constant of 2.72 Å. The muon implantation site is taken to be between two adjacent fluorines, as shown in Fig. 4. As the system is composed of only spin-1/2 particles, we can map each particle to a single qubit, as described in Sec. IV A. We consider a dipolar interaction between all particles; between the muon and the fluorines, and between the fluorines themselves. The

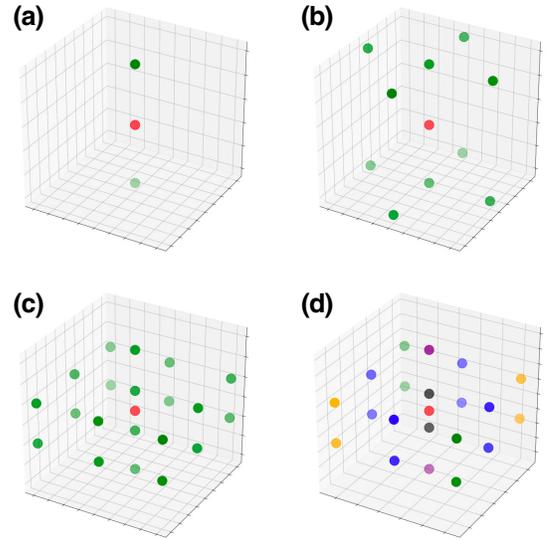


FIG. 4. A selection of geometries investigated for the $CaF_2 + \mu^+$ system considered in this work. As discussed in the main text, the calcium ions are spin 0 with an abundance of around 99.9%, so we include only the fluorine ions [green in (a)–(c)] and the muon (red) in our simulations. (a) The three spin $F-\mu^+-F$ system. (b) The 11-spin $F-\mu^+-F$ system with eight next-nearest-neighbor fluorines. (c) The 21-spin system, which includes the next-next-nearest ten fluorines to the muon. (d) The 21-spin system, with fluorines grouped by color. The distances between the muon and each fluorine in a given group are the same. When fitting simulated data for the 21-qubit system to experimental data, we separately parameterized the $F-\mu$ distances according to these groups.

Hamiltonian is given by

$$H_D = \frac{1}{2} \sum_{ij} \frac{\hbar^2 \mu_0 \gamma_i \gamma_j}{4\pi r_{ij}^2} \left[\vec{S}_i \cdot \vec{S}_j - 3(\vec{S}_i \cdot \hat{r}_{ij})(\vec{S}_j \cdot \hat{r}_{ij}) \right], \quad (46)$$

where μ_0 is the permeability of free space, γ_i is the gyromagnetic ratio of spin i ($\gamma_\mu = 2\pi \times 1.355 \times 10^8 \text{ Hz T}^{-1}$, $\gamma_F = 2\pi \times 4.006 \times 10^7 \text{ Hz T}^{-1}$), \vec{r}_{ij} is the vector connecting spins i and j , and $\vec{S}_i = (1/2)(X_i, Y_i, Z_i)$, where X_i, Y_i, Z_i are the Pauli matrices acting on qubit i .

When calculating the polarization function for this polycrystalline system, we must perform an angular average, because there are a number of equivalent sites located between two fluorines where the muon could implant. We must therefore average the polarization function over all possible orientations of the crystal axes. To calculate $\langle \text{Tr}[\hat{n} \cdot \vec{\sigma}_\mu \rho(t)] \rangle_{\hat{n}}$, we note that $\hat{n} \cdot \vec{\sigma}_\mu = xX_\mu + yY_\mu + zZ_\mu$, and $\rho(t) = e^{-iH_S t} [(1/2)(I_\mu + \hat{n} \cdot \vec{\sigma}_\mu) \otimes \rho_e(0)] e^{iH_S t}$. In the absence of an applied magnetic field, the Hamiltonian does not depend on \hat{n} . Using that $\rho_e(0) = (I_e/D_e)$, and that

$\text{Tr}(\sigma_j) = 0$, we can then expand

$$\text{Tr}[\hat{n} \cdot \vec{\sigma}_\mu \rho(t)] = \frac{1}{2D_e} \sum_{j,k \in \{x,y,z\}} jk \text{Tr}(\sigma_j e^{-iH_s t} \sigma_k e^{iH_s t}). \quad (47)$$

Using that $\text{Tr}[\sigma_j e^{-iH_s t} (I_\mu/2) \otimes (I_e/D_e) e^{iH_s t}] = 0$, this expression can be rewritten as

$$\begin{aligned} & \text{Tr}[\hat{n} \cdot \vec{\sigma}_\mu \rho(t)] \\ &= \sum_{j,k \in \{x,y,z\}} jk \text{Tr} \left\{ \sigma_j e^{-iH_s t} \left[\frac{1}{2} (I_\mu + \sigma_k) \otimes \rho_e(0) \right] e^{iH_s t} \right\}. \end{aligned} \quad (48)$$

Using spherical polar coordinates $x = \cos(\phi)\sin(\theta)$, $y = \sin(\phi)\sin(\theta)$, $z = \cos(\theta)$, and performing the integrals $\int_{\phi=0}^{2\pi} \int_{\theta=0}^{\pi} jk \sin(\theta) d\theta d\phi$, we find that the cross terms vanish, while the diagonal terms give $(4\pi/3)$. Dividing by the area element 4π yields the following formula for the angular averaged polarization function:

$$\begin{aligned} \langle P(t) \rangle &= \frac{1}{3} \left(\text{Tr} \{ Z_\mu e^{-iH_s t} [|0\rangle \langle 0|_\mu \otimes \rho_e(0)] e^{iH_s t} \} \right. \\ &\quad + \text{Tr} \{ X_\mu e^{-iH_s t} [|+ \rangle \langle +|_\mu \otimes \rho_e(0)] e^{iH_s t} \} \\ &\quad \left. + \text{Tr} \{ Y_\mu e^{-iH_s t} [|+Y \rangle \langle +Y|_\mu \otimes \rho_e(0)] e^{iH_s t} \} \right) \end{aligned} \quad (49)$$

where $|+\rangle = (1/\sqrt{2})(|0\rangle + |1\rangle)$, $|+Y\rangle = (1/\sqrt{2})(|0\rangle + i|1\rangle)$, are the $+1$ eigenstates of X and Y , respectively. If an external magnetic field were to be applied, then the angular average would need to be performed explicitly.

We can obtain an increasingly accurate description of the $\text{CaF}_2 + \mu^+$ system by growing the number of fluorines considered in the simulation. The smallest system considered requires three qubits, representing the muon and its two nearest-neighbor (NN) fluorines. Additional fluorine nuclei are then added in “shells” determined by their distance from the muon. In the “next-nearest-neighbor” (NNN) shell, there are an additional eight nuclei. The “next-next-nearest-neighbor” (NNNN) shell contributes an additional ten nuclei. The NNN-NN shell adds eight nuclei. As a result, we consider system sizes of 3 qubits ($\mu + 2$ NN F’s), 11 qubits ($+8$ NNN F’s), 21 qubits ($+10$ NNNN F’s), and 29 qubits ($+8$ NNN-NN F’s).

We utilize a range of simulation techniques, in order to investigate a number of different properties of our proposed algorithm. Simulations of the random-phase-approximation-based approach described in Sec. II are carried out using the QuEST package for emulating quantum circuits [97]. These simulations manually generated the initial state shown in Eq. (7), and then carried out quantum circuit emulations of first and second-order Trotter

decompositions of the time-evolution operator. QuEST is implemented in the C language, and can be efficiently parallelized using OpenMP or MPI. This efficiency enables us to perform calculations on system sizes of up to 29 qubits, surpassing the largest calculations performed previously in the μ^+ SR literature [32] (to the best of our knowledge). Because these simulations used the random-phase-approximation approach, sampling noise is present in the results. We also perform quantum circuit-level simulations of running the algorithm on a quantum processor, using the density-matrix simulator included in Cirq [98], a PYTHON package for the simulation of NISQ hardware. Due to the increased computational cost of storing and manipulating the density matrix, these calculations are restricted to systems of up to 11 qubits. However, these simulations enable us to initialize the environment in a maximally mixed state, eliminating the sampling error present in wave-function-based approaches. They also provide a more efficient way to investigate the effects of circuit-level noise on the algorithm. We also perform exact diagonalization of the Hamiltonians of systems with up to 11 qubits. This provides exact results, which can be used to quantify the error introduced by Trotterizing the time-evolution operator. The Hamiltonians in this work are generated using OpenFermion, a PYTHON package for generating qubit-mapped Hamiltonians of physical systems, such as molecular electronic structure Hamiltonians [99].

A. Noiseless simulations

As discussed in Sec. II, there are two sources of algorithmic error in the wave-function-based method introduced by Celio [30] for simulating muon polarization functions; error introduced by Trotterizing the time-evolution operator, and sampling errors arising from working with a random-phase augmented wave function, rather than with an actual mixed state. In this section, we discuss numerical simulations that quantify the magnitude and scaling of these sources of error, for the $\text{CaF}_2 + \mu^+$ system investigated in this work. Given the similarities between Celio’s method and the quantum algorithm introduced in Sec. IV, these simulations enable us to better quantify the quantum resources required to run our algorithm.

1. Quantifying errors

We first investigate the Trotter error present in the algorithm for the 3- and 11-qubit systems. We are able to isolate the Trotter error from the sampling error by simulating the Trotterized time evolution of the density matrix $(I_e/D_e) \otimes |0\rangle \langle 0|_\mu$. We time evolve this state using both first- and second-order Trotter formulae, with a “magnitude ordering” of the terms. By this, we mean that the strongest Hamiltonian terms are placed first in the product-formula sequence. We then calculate the polarization value of the muon qubit at a given time, and compare this to the

value obtained from an exact diagonalization of the muon-environment Hamiltonian [see Eq. (5)]. We compare the numerically simulated Trotter error to two different bounds for the Trotter error. The first bound is given by Eq. (34) for a first-order Trotterization, and Eq. (37) for second-order Trotter, and represents a loose bound on the error between the exact time-evolution operator and the unitary specified by the product formula. The second bound is given by Eq. (35) for a first-order Trotterization, and Eq. (38) for second-order Trotter (the commutators in the sums are evaluated, the resulting operator is simplified by collecting like terms, and then the triangle inequality is applied, so that the spectral norm is bounded by the 1 norm of the resulting error operator), and represents a tighter bound on the error between the exact time-evolution operator and the unitary specified by the product formula. The improved “tightness” of the second bound stems from it taking into account commutativity between different terms in the Hamiltonian. These quantities upper bound the error in any observable measured after time evolution, and so will be strictly larger than the error in the polarization value obtained from our numerical simulations. The Trotter errors for the 3- and 11-qubit systems are shown in Figs. 5 and 6, respectively. We plot the Trotter errors obtained at time values of 5, 10, and 15 μs .

We see from both plots that in all cases the “loose” error bounds are orders of magnitude larger than the “tight” error bounds, which in turn are at least an order of magnitude larger than the numerically observed Trotter errors. As discussed above, this can be partially explained by the fact that the analytic bounds give errors in the unitary evolution, while the numerical results give the error in a specific observable. In addition, despite the improved bounds given by the “tight” formulae, they are still known to not be completely tight to numerical results [85]. However, the fact that the tight bounds are still at least an order of magnitude larger than the numerical results highlights the value in work to bound the error in specific observables, rather than existing worst-case bounds. While initial steps have been taken in this direction [85,88], it would be interesting to consider if tighter bounds are possible for the case of muon spectroscopy, given the simple form of the observable measured.

Another interesting observation is that both the first- and second-order product formulae appear to give similar asymptotic scaling for the numerically observed Trotter error. This is in contrast to the expected behavior that is evident in the analytic bounds; that the second-order formula should show improved asymptotic scaling as a function of the number of Trotter steps taken. As implementing the second-order formula on a quantum computer can require twice the number of quantum gates needed for the first-order formula, this suggests that there may be scenarios in which it is preferable to use the first-order approach. For example, the three-qubit data at 15 μs

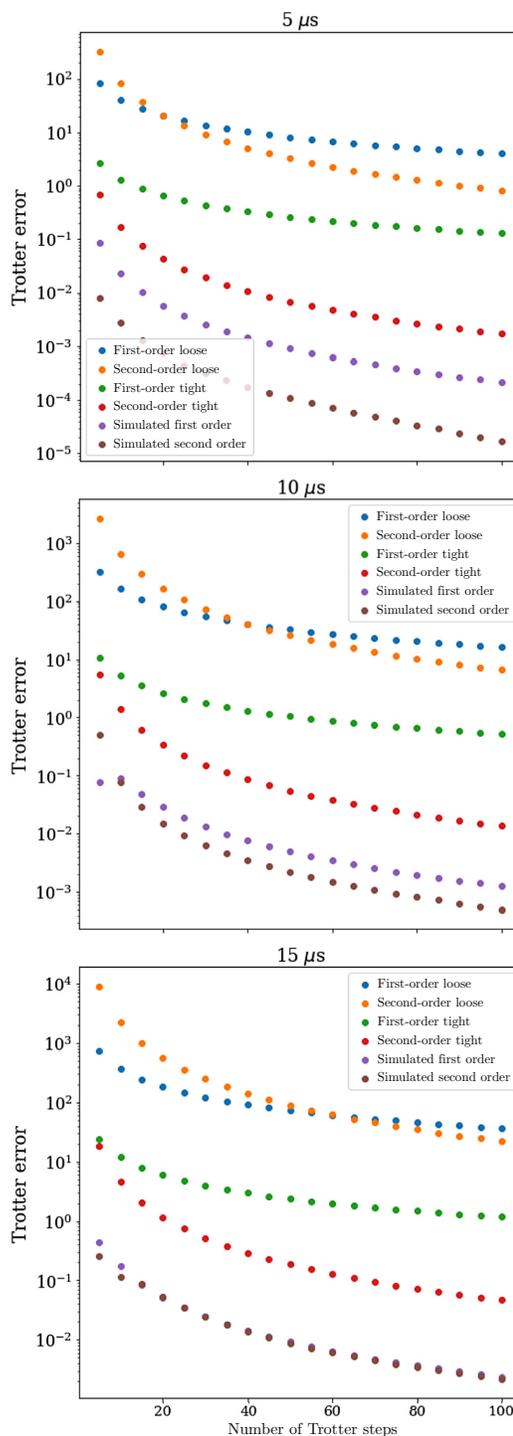


FIG. 5. The Trotter error in the three-qubit $F-\mu^+-F$ system. The error bounds used are described in the main text. In all of these results, we use a magnitude ordering of Hamiltonian terms in the product formula; the terms with the largest coefficients are placed first in the Trotter formula.

would suggest using a first-order decomposition, over a second-order approach. In most cases however, the second-order formula appears to offer improved constant factor

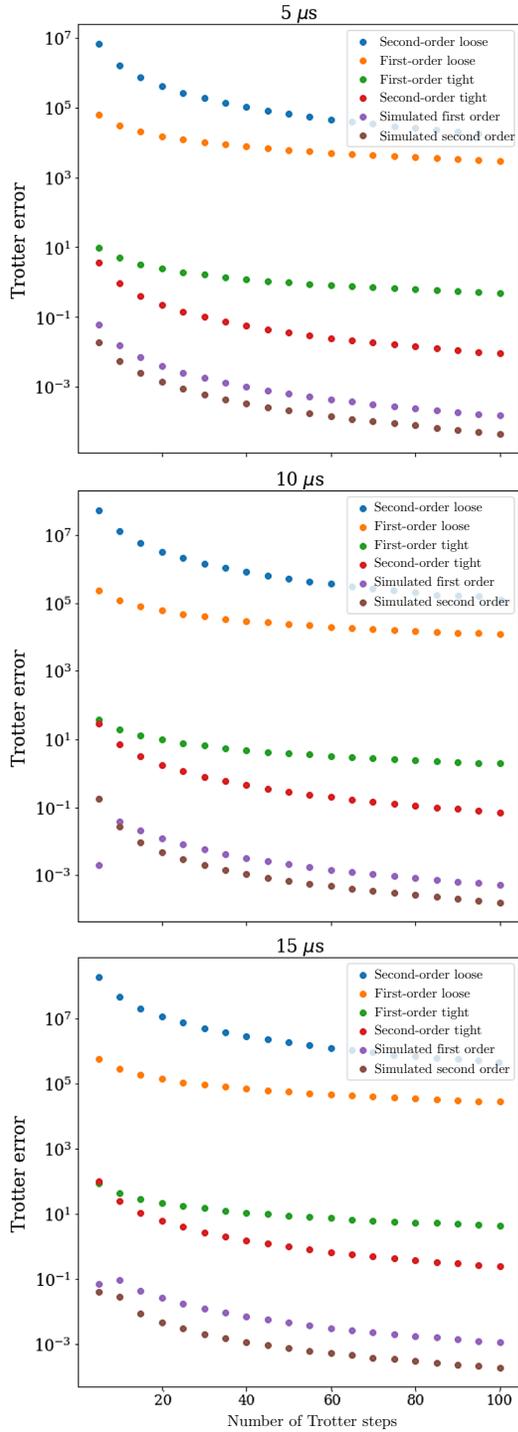


FIG. 6. The Trotter error in the 11-qubit system. The error bounds used are described in the main text. In all of these results, we use a magnitude ordering of Hamiltonian terms in the product formula; the terms with the largest coefficients are placed first in the Trotter formula.

scaling that makes up for the increased gate depth required for implementation. It would be interesting to investigate whether it is possible to obtain further improvements using higher-order product formulae, as has previously been

observed in the simulation of other spin-1/2 systems [84]. We can see that for the 11-qubit system, approximately 40 second-order Trotter steps are sufficient to obtain an accuracy of 10^{-3} in the polarization function at times less than $15 \mu s$. Interestingly, the Trotter error for the 11-qubit system does not seem to worsen significantly as the simulated time is increased.

We also investigate the error that arises from sampling a wave-function-based simulation, rather than time evolving the exact mixed state that describes the system. We consider sampling environmental basis states $|j\rangle_e$ chosen with equal probability, as well as two possible sampling schemes given by Eq. (7), which we restate here

$$|\psi_s(t)\rangle = U(t) \left[\frac{1}{\sqrt{D_e}} \sum_j e^{i\theta_j} |0\rangle_\mu |j\rangle_e \right], \quad (50)$$

The first scheme, referred to as the “random-phase-approximation (RPA) method” is exactly the same as Celio’s method, and considers θ_j chosen uniformly at random in the range $[0, 2\pi)$. The second, referred to as the “dephasing method” considers θ_j chosen at random from the discrete set $\{0, \pi\}$. We refer to this approach as the dephasing method, as it can be understood from a quantum-computing perspective as first applying a Hadamard gate to each environment qubit to obtain the state

$$|\psi_s\rangle = \left[\frac{1}{\sqrt{D_e}} \sum_j |0\rangle_\mu |j\rangle_e \right], \quad (51)$$

and then passing this state through an N_e -qubit dephasing channel, with an equal probability of all errors

$$\rho \rightarrow \rho' = \frac{1}{2^{N_e}} \left[\rho + \sum_{i=0}^{N_e-1} Z_i \rho Z_i + \dots + \bigotimes_{j=0}^{N_e-1} Z_j \rho Z_j \right]. \quad (52)$$

This channel can easily be sampled on a quantum computer by applying any of the possible Z strings acting on N_e qubits with equal probability. This approach can be generalized to particles with spin greater than $1/2$ using the method discussed in Sec. IV B.

In order to isolate the sampling error, we fix the number of Trotter steps, and compare the results obtained from a calculation on the 11-qubit system using both sampling and Trotterization, to those from a calculation that uses Trotterized time evolution of the mixed initial state $(I_e/D_e) \otimes |0\rangle\langle 0|_\mu$. Without loss of generality, we fix the number of Trotter steps to 40, and take the average of the absolute values of the sampling error at 20 time values, spaced equally in the interval $[0, 10) \mu s$. In Fig. 7 we observe that the sampling error arising from

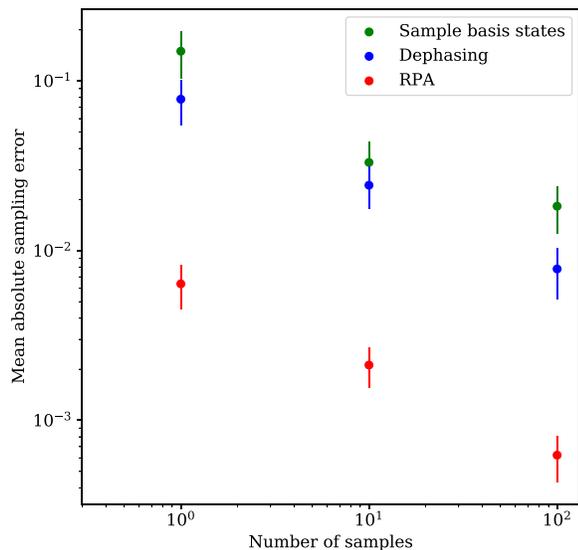


FIG. 7. The sampling error present in the random-phase-approximation, dephasing, and basis state sampling methods described in the main text, for the 11-qubit system with 40 second-order Trotter steps. Each point is the average of 20 values taken in the interval $[0, 10)$ μs .

the RPA method is around an order of magnitude smaller than that of the dephasing method, due to the increased randomization of the former technique.

As shown in Eq. (8), the error in the sampling-based methods decreases as the size of the simulated system increases. We investigate this by considering the sampling error in a larger, 21-qubit simulation. Due to the increased size of this simulation, we are unable to carry out simulations utilizing the “exact” mixed initial state for comparison. In lieu of this, we investigate the convergence of the obtained polarization function as the number of samples is increased. In Fig. 8 we investigate this behavior for both the RPA and dephasing methods. We observe that the RPA method rapidly converges, and exhibits small fluctuations around its converged value. In particular, we find that using the RPA method, only 20 samples suffice to converge the polarization function to within 10^{-4} of the value obtained with 9600 samples.

In a similar vein, we can investigate the convergence of Trotter errors for systems too large to be exactly simulated. In Fig. 9 we show how the polarization function converges as the number of Trotter steps is increased, for both the 21- and 29-qubit systems. These simulations are performed using the RPA method, with 48 samples for the 21-qubit simulation, and a single sample for the 29-qubit simulation. We compare the polarization function at a range of Trotter steps to that obtained with 30 Trotter steps. We note that this metric does not provide a bound on the Trotter error at 30 Trotter steps. However, given that the convergence error in both cases is monotonically decreasing, and less than 10^{-2} by 20 Trotter steps, this

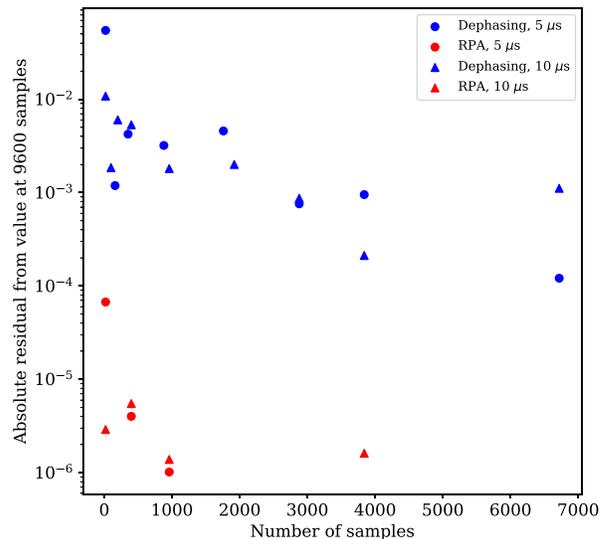


FIG. 8. We compare the sampling error obtained using both the random-phase-approximation method and the dephasing method, at a given number of samples, to that obtained at 9600 samples. We see that the RPA method rapidly converges to the value obtained at 9600 samples. While we are unable to bound the sampling error at 9600 samples, we see from Fig. 7 that the RPA method can obtain an accuracy of less than 10^{-3} for the 11-qubit system with 100 samples. Given that the sampling error decreases with both system size, and the number of samples, we can infer that the error is likely small for even a modest number of samples with the RPA method for the 21-qubit system.

may be taken as an indication that the polarization function is rapidly converging as the number of Trotter steps is increased.

2. Analyzing $\text{CaF}_2 + \mu^+$ spectra

We can use the results discussed above to determine how many samples and Trotter steps to use in order to quantitatively investigate the $\text{CaF}_2 + \mu^+$ system. In Fig. 10 we plot the angular averaged polarization functions obtained for the 3-, 11-, 21-, and 29-qubit $\text{CaF}_2 + \mu^+$ systems for the first 9.5 μs of evolution. The simulated data points for the 3- and 11-qubit systems are obtained by time evolving the initial state $(I_e/D_e) \otimes |0\rangle\langle 0|_{\mu}$, using 30 second-order Trotter steps, which suffices to measure the polarization function to an accuracy of less than 10^{-2} . These data points are plotted with the polarization functions obtained from exact diagonalization of the corresponding system Hamiltonian. The simulated data points for the 21- and 29-qubit systems are obtained using the RPA method, with 48 samples used for the 21-qubit system, and a single sample used for the 29-qubit system. Both of these simulations also use 30 second-order Trotter steps. We see from Fig. 10 that introducing additional fluorine nuclei causes a damping effect on the polarization function, which appears well

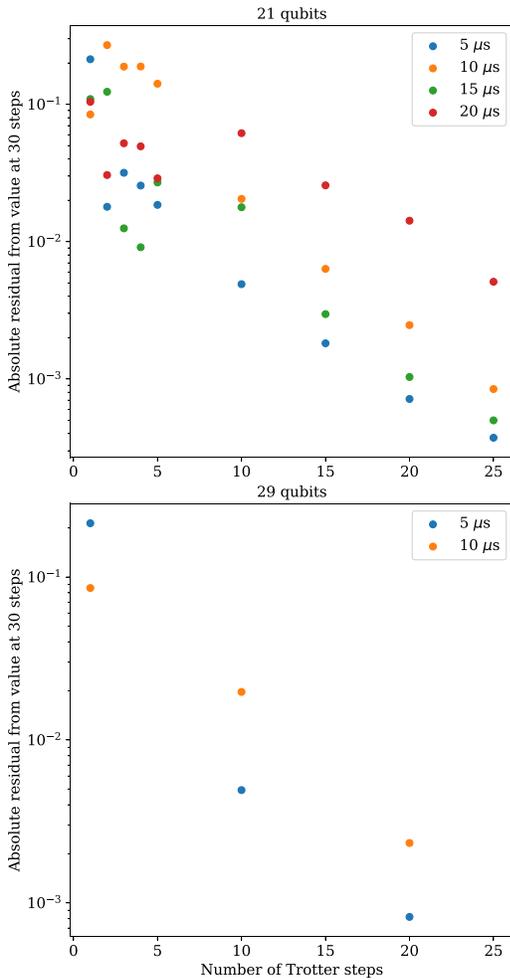


FIG. 9. Convergence of the Trotter error for both the 21- and 29-qubit systems, as a function of the number of Trotter steps used. As these systems are too large to exactly simulate classically, we are unable to calculate the Trotter error in the polarization function. As a result, we plot how the polarization converges towards the value at 30 Trotter steps (which may still be far from the true value).

converged at 29 qubits. This observation is in good agreement with the recent numerical results of Wilkinson and Blundell [9], who showed that the additional environmental spins act as a source of decoherence, as polarization leaks from the muon to the environment.

Having confirmed that our algorithm produces the qualitative results expected, we can use it to locate the muon rest site in CaF_2 . This is achieved by parametrizing the muon-fluorine distances, and generating a polarization function for a given geometry. We can then compute the value of a loss function between the simulated data and the experimental data (in this case, the reduced- χ^2 value), and use an optimization algorithm to generate new values of the parameters that describe the geometry of the system. In our numerical simulations, we use the Nelder-Mead

algorithm. It is necessary to use a derivative-free optimization method because of the sampling noise present in the RPA method, which is larger than the finite-difference steps used to calculate the gradient in many black-box optimization algorithms. We use this approach to optimize the geometry of the 21-qubit system (composed of the muon, the two nearest-neighbor fluorines, eight next-nearest-neighbor, and ten next-next-nearest-neighbor fluorines). Five geometric fitting parameters are used, which we describe via the colors used in Fig. 4(d):

1. The distance between the muon (red) and the nearest-neighbor fluorines (both colored black).
2. The distance between the muon and the next-nearest-neighbor fluorines (all colored blue).
3. The distance between the muon and the green next-next-nearest-neighbor fluorines.
4. The distance between the muon and the purple next-next-nearest-neighbor fluorines.
5. The distance between the muon and the orange next-next-nearest-neighbor fluorines.

along with two parameters to fit the asymmetry

$$A(t) = A_0 P(\theta_1, \dots, \theta_5; t) + A_{\text{bg}}. \quad (53)$$

We use the RPA method to generate the polarization function, with 40 second-order Trotter steps, and a single sample for each data point.

We plot the fit to the experimental data in Fig. 11. The fitted data shows excellent qualitative agreement with the experimental data, particularly at early times. The fit obtains a reduced χ^2 value of 2.13, which suggests an incomplete fit to the data. We attribute this to a combination of limited experimental data at times beyond 10 μs , as well as the ineffective nature of the Nelder-Mead optimization algorithm, which is liable to becoming trapped in local minima. Our fitting procedure causes the nearest-neighbor fluorines to move towards the muon by 0.188 Å. This is in excellent agreement with the calculation of Wilkinson and Blundell [9], which yielded a value of 0.190(1) Å for the same quantity. We note that the results of Ref. [9] were obtained by fitting an 11-spin system to the experimental data, and considering two physical parameters; the muon–nearest-neighbor fluorines distance, and a factor that scaled the strength of the next-nearest-neighbor interactions to act as a proxy for more distant nuclei. Our fitting procedure also caused the next-nearest-neighbor fluorines to move towards the muon by 0.206 Å. While Ref. [9] did not explicitly consider the effect of perturbing the positions of the more distant nuclei, they carried out density-functional-theory (DFT) calculations suggesting that the next-nearest-neighbor fluorines move only towards the muon by around 0.03 Å. We carry out additional numerical simulations with this geometry for

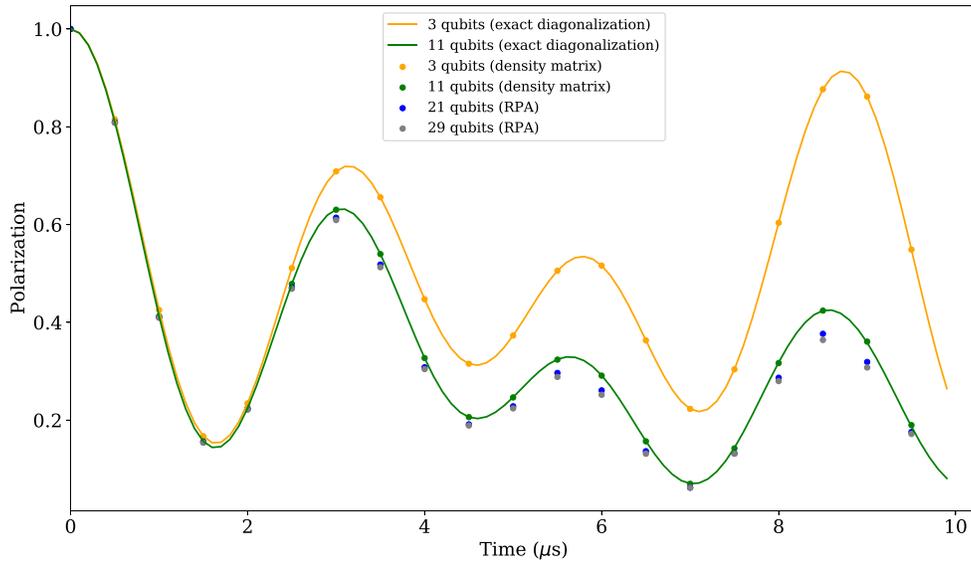


FIG. 10. Polarization functions for the unperturbed 3-, 11-, 21-, and 29-qubit $\text{CaF}_2 + \mu^+$ systems. All simulations use 30 second-order Trotter steps to simulate time evolution. The 3- and 11-qubit results are obtained using a density-matrix simulation, that eliminates sampling error. The 21- and 29-qubit simulations are performed using the random-phase-approximation method, with 48 and 1 samples, respectively.

the 21-spin system (shown in Fig. 16 in Appendix A), and found a worse fit for the experimental data, with a χ^2_ν value of 4.44. This implies that either:

1. Our system size of 21 spins is still not large enough to fully capture the continuum extrapolation of Ref. [9], and larger simulations are required (which are impractical to perform on classical hardware).
2. The fit obtained by our simulation may suggest an inaccuracy in the DFT results. However, from a physical perspective, it would be surprising if the more distant nuclei are more strongly attracted towards the muon than the nearest neighbors.

The most likely explanation may be a combination of these factors, together with the fact that a number of similarly good local minima are present in the χ^2 surface for our parameter space. This highlights the value in utilizing complementary techniques to analyze muon-spectroscopy data, as well as the benefit provided by having access to as large a simulation of the system as possible.

B. Noisy simulations

As discussed in Sec. III, existing quantum computers have far higher error rates than their classical counterparts. The presence of physical noise can corrupt the results of calculations on these devices, rendering any quantum speedup offered moot. As such, it is essential to investigate the effects of noise on our algorithm. A simple model of the noise present in near-term quantum-computing devices

is the single-qubit depolarizing noise channel

$$\rho \rightarrow \rho' = (1 - p)\rho + \frac{p}{3}(X\rho X + Y\rho Y + Z\rho Z), \quad (54)$$

where ρ and ρ' are the density matrix of the qubit before and after the noise channel (respectively). This channel effectively applies an X , Y , or Z error to the qubit, with probability $p/3$, and leaves it unchanged with probability $(1 - p)$. We assume that this noise model is applied after each gate in the circuit, acting separately on each qubit involved in the gate. In Fig. 12 we plot the effects of this depolarizing noise model on an exact density-matrix simulation of the three-qubit F- μ^+ -F system. We use 20 second-order Trotter steps for the simulation, which suffices to obtain an accuracy of less than 10^{-3} in the polarization function at all values within the first 5 μs . We merge adjacent single-qubit gates together, to yield a circuit with 900 single-qubit gates, and 680 two-qubit gates. In this error model, we are assuming that noise is uncorrelated between the qubits involved in a two-qubit gate, and that the error rate for single-qubit gates is the same as that for two-qubit gates. We note that the circuit depth is the same for all of the points calculated. This means that the Trotter error is likely smaller in points taken at earlier times than at later times, while the effective physical noise rate of the circuit will be similar at all points. In reality, we would likely choose to fix the Trotter error along the polarization function, which would enable us to use a shorter depth circuit to simulate points at earlier times—thus reducing the physical noise in those datapoints. We see from Fig. 12 that even with a depolarizing noise rate of $p = 10^{-4}$ (which

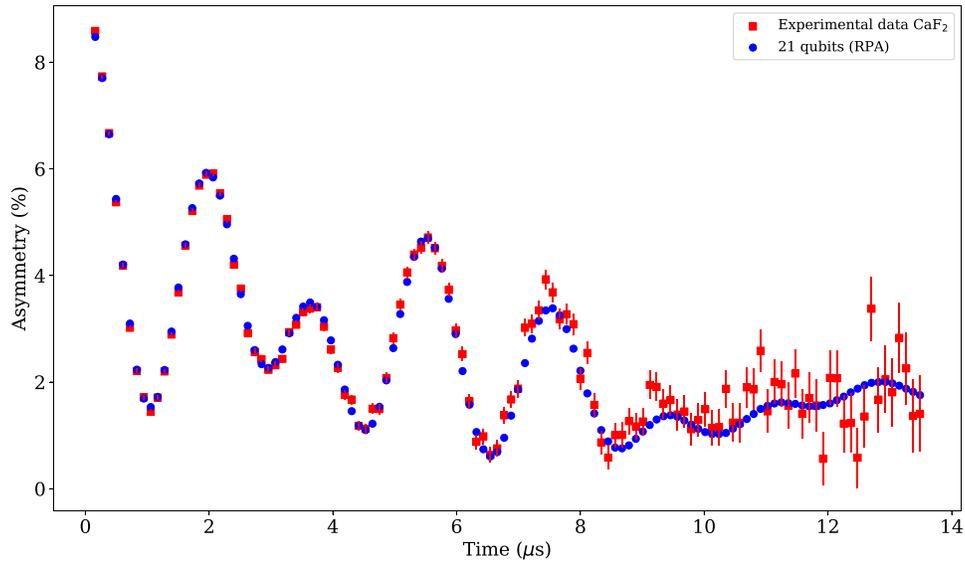


FIG. 11. Fit of the 21-qubit simulated data (using the random-phase-approximation method with one sample per data point, and 40 second-order Trotter steps) to the experimental data obtained in Ref. [9]. The fit is performed using the gradient-free Nelder-Mead algorithm. The fitting parameters are the muon-fluorine distances described in the main text, and a scaling factor and offset to convert the simulated polarization value to an asymmetry value.

is an order of magnitude lower than the best error rates observed to date in quantum hardware [100,101]), there is still a significant decay in the value of the simulated polarization function. This noise rate corresponds to an expected number of errors per circuit (defined as the error rate multiplied by the number of gates in the circuit) of around 0.16.

It has been previously observed that when the expected number of errors in the circuit is less than around unity, the error-mitigation techniques introduced in Sec. III can be successful in recovering the noiseless value of the circuit [68,71]. In this work, we investigate the use of exponential extrapolation to mitigate the effects of noise. We consider the same three-qubit system discussed above, with a baseline noise rate of $p = 5 \times 10^{-4}$. We then artificially “boost” the error rate by a heuristically chosen factor of $\lambda = 1.1$, and calculate the extrapolated expectation value as

$$P_0 = \left(\frac{P_\epsilon^\lambda}{P_{\lambda\epsilon}} \right)^{\frac{1}{\lambda-1}}, \quad (55)$$

where P_ϵ is the polarization value calculated with the baseline noise rate, and $P_{\lambda\epsilon}$ is the polarization value calculated with the boosted noise rate. We see from Fig. 13 that this exponential extrapolation is able to recover almost noiseless results, despite the large damping of the polarization function in the unmitigated case. The noise strength of $p = 5 \times 10^{-4}$ corresponds to an expected number of errors per circuit of 0.79. The mean absolute error in polarization function after exponential extrapolation is 0.011.

In this work, we investigate a quantum algorithm for analyzing muon-spectroscopy data. As such, it is interesting to ask whether the noise inevitably present in μ^+ SR experimental results means that physical errors in the quantum computer are less damaging than they would be in

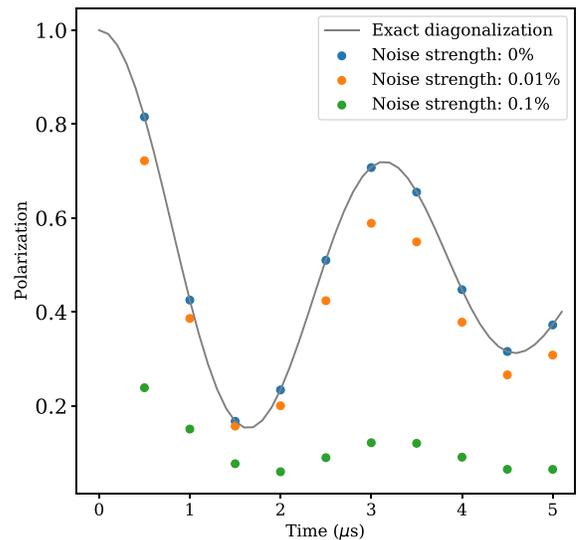


FIG. 12. Density-matrix simulations of the three qubit F- μ^+ -F system, with single-qubit depolarizing noise applied after each gate. The noise channel is applied independently to each qubit involved in the gate. The simulation used 20 second-order Trotter steps, which yield a circuit with 900 single-qubit gates, and 680 two-qubit gates.

other applications (such as trying to calculate the ground-state energies of chemical systems). As shown in Fig. 11, muon-spectroscopy data can exhibit errors arising from a number of sources—most notably, the statistical uncertainties arising from recording a finite number of muon decay events. In order to investigate the noise robustness of our algorithm for analyzing μ^+ SR data, we try to locate the muon rest site in the presence of noise in the simulated data for the 11-qubit system. Rather than considering the circuit-level noise models discussed above, we effectively engineer noise in the polarization function through undersampling in the RPA method. As shown in Fig. 7, we can generate mean errors of 0.0243, 0.0068, and 0.0022 in the polarization function by using the dephasing method with ten samples, the RPA method with one sample, and the RPA method with ten samples, respectively. We note that there are two important assumptions present in this noise model. Firstly, we are assuming that physical noise would cause the data points to become normally distributed around the noiseless value, as happens for the sampling noise. As is evident from Fig. 12, this is not necessarily the case; many sources of noise will simply cause a decay in the polarization function. However, we can see from Fig. 13 that performing extrapolation can change how the data points are distributed around the noiseless result. When shot noise in the quantum algorithm is taken into account, this may enable us to approach the normal distribution of noise considered in these simulations. Another path to recover normally distributed results could be to use an alternative error-mitigation method, known as the “quasiprobability technique” [68,69,102,103]. The

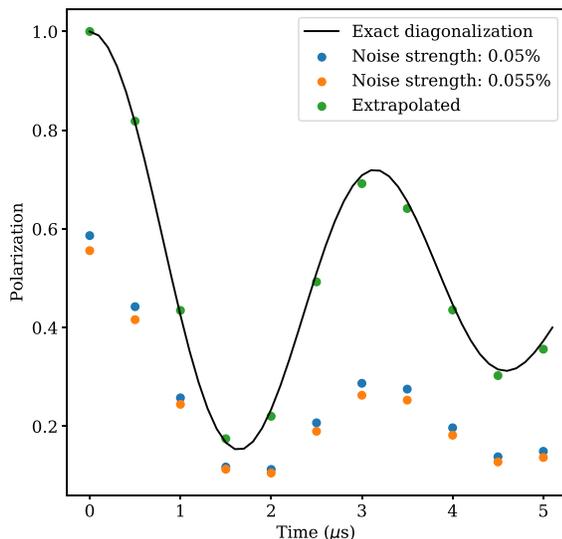


FIG. 13. Using exponential extrapolation of data obtained at two different noise rates to infer the noise-free polarization function. The extrapolation is performed using Eq. (55). The data points are obtained in the same way as described in the caption of Fig. 12. The expected number of errors in the circuit is 0.79.

quasiprobability technique attempts to invert the physical noise channel of each gate, by sampling from an increased number of carefully modified, noisy circuits, and combining the results in postprocessing. This approach has previously been observed to give results that are approximately normally distributed around the true values [68]. Secondly, we assume that the magnitude of the noise is the same for all data points. In reality, as discussed above, we would fix the Trotter error along the polarization function, and thus consider shorter depth circuits at earlier times. This would reduce the noise in earlier data points, mimicking the error bars present in the experimental results (in the experimental case, the muon half life is around $2.2 \mu\text{s}$, so a significant number of decay events must be observed to record sufficient statistics at later time values). In our simulations we use 40 second-order Trotter steps, and consider two geometric fitting parameters:

1. The muon–nearest-neighbor fluorine distance.
2. The muon–next-nearest-neighbor fluorine distance.

Along with two parameters to fit the asymmetry

$$A(t) = A_0 P(\theta_1, \theta_2; t) + A_{\text{bg}}. \quad (56)$$

Once again, we use the Nelder-Mead algorithm to minimize the error-weighted least-squares residuals between the experimental data and the simulated data.

In Fig. 14 we plot examples of the fitted data obtained at the three noise values listed above. We see that even at high noise rates of around 2.5% in each data point, the fitting procedure is able to approximately capture the overall shape of the polarization function. As the strength of the noise is reduced, we obtain increasingly better fits to the experimental data.

In order to better quantify the degree of noise robustness present in the algorithm, we plot the error that noise induces on the noiseless values of the two geometric fitting parameters in Fig. 15. The noiseless values of the fitted parameters are obtained by fitting the 11-qubit polarization function, generated by exact diagonalization, to the experimental data, using the Nelder-Mead algorithm. The upper plot in Fig. 15 shows that as the average noise in each individual data point is reduced, we obtain an improved estimate of the noiseless muon–nearest-neighbor fluorine distance. Interestingly, the fractional error in the parameter is around an order of magnitude smaller than minimum fractional error in the polarization value (the maximum value of the polarization function is 1, so the minimum fractional error is equivalent to the noise strength of the simulation). We attribute this noise resilience to the fact that the fitting procedure is extracting a global property of the system (the muon-fluorine distance) rather than the value of any individual data point. The lower plot in Fig. 15 suggests that the muon–next-nearest-neighbor fluorine distance is less resilient to noise, as the fractional error in

the parameter is approximately the same order of magnitude as (but larger than) the magnitude of the noise in the polarization function. Nevertheless, we are optimistic that the algorithm could be made more noise resilient by (1) reducing the noise strength at earlier times by adapting the number of Trotter steps used (as discussed above), and (2) using an optimization routine known to be more resilient to noise than the Nelder-Mead algorithm, such as Bayesian optimization.

Linking back to our earlier simulations of circuit-level noise, we note the following:

1. Our fitting routine obtained accurate parameter values when the error in each simulated polarization data point is less than around 0.01 (for the 11-qubit system).

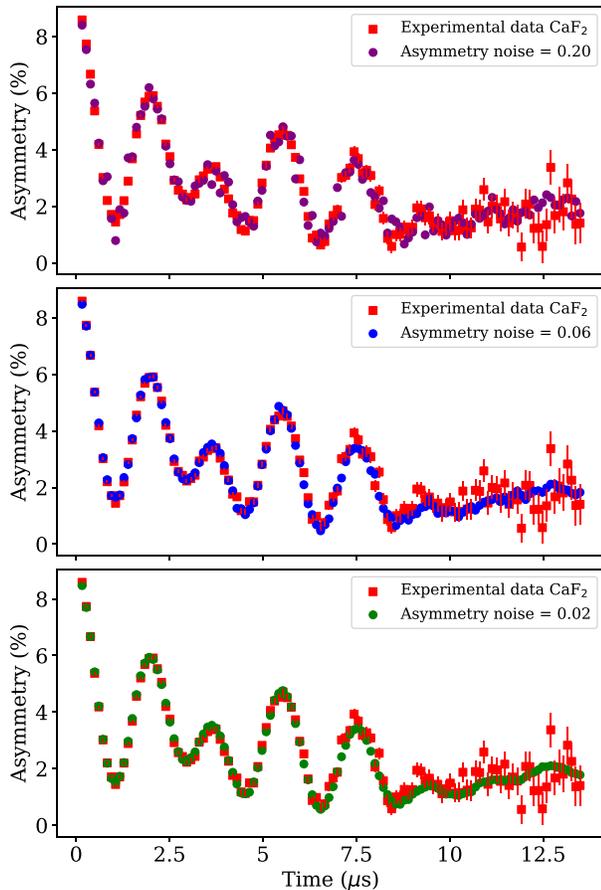


FIG. 14. Fits of the 11-qubit simulated data with 40 second-order Trotter steps to the experimental data obtained in Ref. [9]. The noise in each simulated datapoint is varied by controlling the number of samples used. The fit is performed using the gradient-free Nelder-Mead algorithm. The fitting parameters are the muon-fluorine distances described in the main text, and a scaling factor and offset to convert the simulated polarization value to an asymmetry value.

2. Using exponential extrapolation, we obtain a mean absolute error in the polarization function of 0.011 when the expected number of errors in the circuit is around 0.8 (for the three-qubit system).

Assuming that the circuit-level result holds for larger system sizes (we would expect the results to be better for larger system sizes, as when errors do occur, in larger system sizes they are less likely to occur on the crucial muon qubit), we infer that our data-analysis algorithm is able to tolerate an expected error rate of around 0.8 errors per circuit, on average. In the 11-qubit case, this would enable us to learn the muon–nearest-neighbor fluorine distance to an accuracy better than 0.1%, and the muon–next-nearest-neighbor fluorine distance to an accuracy of around 5%.

It would be interesting to consider whether this fitting procedure could be shown to be provably noise robust. If so, we would expect these results to be useful beyond this work, earmarking experimental data analysis as a

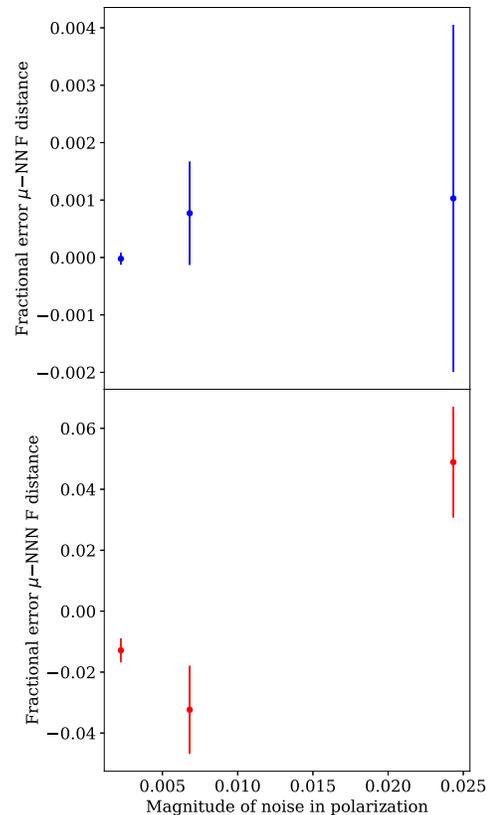


FIG. 15. Fractional error in the geometric fitting parameters used to fit the simulated muon-spectroscopy data to the experimental data in Fig. 14. The noiseless values of the parameters are obtained through exact diagonalization of the 11-qubit Hamiltonian to obtain a polarization function, which is fitted using the Nelder-Mead algorithm.

promising application area for quantum computers, even in the presence of noise.

C. Resource estimates

We now consider the quantum resources required to perform the calculations described above. We consider different cost metrics for NISQ simulations and error-corrected simulations. When performing NISQ calculations, we aim to minimize both the circuit depth and the total number of gates, as these will determine the number of errors that are likely to occur in the circuit. In particular, we focus on two-qubit gates, as in many systems these have error rates at least an order of magnitude worse than those of the single-qubit gates. As NISQ devices typically lack the resources required for quantum error correction, the number of gates that can be applied in the circuit is limited by the magnitude of the noise in the system.

In order to determine an appropriate fault-tolerant cost metric, we first need to consider which error-correcting code we are considering using for the simulation. One possible choice of error-correcting code is the surface code [104]. The surface code has received considerable attention in recent years, owing to its comparatively high code threshold of around 1%, and its compatibility with two-dimensional (2D) nearest-neighbor connectivity architectures [59]—which are realistic architectures for solid-state qubit systems. When considering error-corrected calculations using the surface code, we seek to minimize the number of T gates. While T gates (or an alternative “non-Clifford group” gate) are required for universality in quantum computation, these gates cannot be natively implemented in a fault-tolerant manner in the surface code. They are typically implemented via processes known as “magic state distillation” and “magic state injection”, and are often the dominant cost in fault-tolerant resource estimates. Arbitrary angle single-qubit rotations can be synthesized from these T gates. The number of T gates required per rotation depends logarithmically on the inverse of the synthesis error [105]. A reasonable assumption for calculations of the size considered in this work, is that around 100 T gates per single-qubit rotation will suffice [106].

1. NISQ resource estimates

As shown in Fig. 6, we require approximately 40 second-order Trotter steps to reduce the Trotter error in the observable to 10^{-3} when calculating $P(t = 15 \mu\text{s})$ for the 11-qubit system. We perform a basic compilation of the circuit, by merging adjacent single-qubit gates. This results in a gate count of 5×10^4 single-qubit gates, and 3.9×10^4 two-qubit gates. If we assume that single-qubit gates have an error rate 10 times lower than two-qubit gates, then this circuit is roughly equivalent to implementing 4.4×10^4 two-qubit gates. As discussed in Sec. VB,

we observe that for the three-qubit system, exponential extrapolation is able to recover acceptably accurate results when there are around 0.8 expected errors per circuit. In order to achieve this circuit error rate for the 11-qubit system with 4.4×10^4 two-qubit gates, we would require a two-qubit gate error rate of $p = 2 \times 10^{-5}$. This is approximately 2 orders of magnitude lower than the current lowest two-qubit gate error rates [57,100,101]. Simulations of larger, classically challenging system sizes will require an even larger number of gates. The most promising avenue for realizing these calculations before the advent of quantum error correction is to significantly reduce the gate count required. One possible route towards this goal would be to consider Trotter orderings that reduce the Trotter error, or that enable a larger number of gates to be cancelled. We could alternatively consider neglecting the dipolar interactions between the nuclei, retaining only the muon-fluorine dipolar interactions (as has previously been done in classical simulations [32]). Even if these optimizations can be incorporated, it appears challenging for NISQ devices to surpass classical μ^+ SR simulation capabilities. This is not necessarily reflective of the performance of the algorithm described in this work, and instead serves to highlight both the quality of classical simulations of quantum systems, and the challenges inherent in many NISQ algorithms, resulting from high noise rates in current quantum processors.

2. Error-corrected resource estimates

Given the challenges faced by NISQ machines, it is natural to investigate the resources required to implement the algorithm described in this work in an error-corrected setting. As discussed above, we consider error correction within the surface code. We initially consider running our algorithm on a “small” fault-tolerant quantum computer, meaning that we use the smallest footprint possible to implement the algorithm. Nevertheless, this device will likely still contain many thousands of physical qubits. We follow the approach to surface-code resource estimation taken in Ref. [107], which considers surface-code operations implemented using lattice surgery, complemented with magic state distillation for implementing T gates. We use a “compact” data block of physical qubits to store each logical qubit. For a system with Q logical qubits, using a compact block results in needing $\lceil 1.5Q + 3 \rceil$ surface-code tiles. For a distance d surface code, each surface-code tile consists of $2d^2$ physical qubits (d^2 data qubits and d^2 syndrome qubits).

For the 11-qubit system with 40 second-order Trotter steps, we require 1.96×10^4 non-Clifford single-qubit rotations, which we assume requires 1.96×10^6 T gates. We use a 15-1 distillation block for magic state distillation, which, for a physical noise rate of $p = 10^{-3}$, produces output magic states with an error rate of at most 3.5×10^{-8} ,

TABLE I. Error-corrected resource estimates for simulations of the 11-qubit $\text{CaF}_2 + \mu^+$ system with 40 second-order Trotter steps. Here, p is the physical error rate, ϵ is the expected number of errors per circuit, and d is the surface-code distance used.

p	ϵ	d	Physical qubits	Circuit time (s)
10^{-3}	0.01	22	30 008	474
10^{-4}	0.01	10	6200	216
10^{-3}	0.8	18	20 088	388
10^{-4}	0.8	8	3968	172

and requires an additional 11 surface-code tiles. For our 11-qubit system, this results in $B = 31$ surface-code tiles in total. With this setup, we can consume a magic state every $11d$ surface-code cycles. Our calculation therefore must “survive” for $T = 11d \times (1.96 \times 10^6)$ code cycles. The size of the surface code required to perform our calculation can be determined using [107]

$$B \times T \times 0.1(100p)^{(d+1)/2} < \epsilon, \quad (57)$$

where p is the physical error rate, and ϵ is the target circuit error rate. The fault-tolerant resources required are shown in Table I, for a number of scenarios. We consider resource estimates with a realistic two-qubit gate error rate of $p = 10^{-3}$, as well as a more optimistic error rate of $p = 10^{-4}$. We also consider two target noise suppression levels. In the first case, we consider suppressing the error rate such that there are $\epsilon = 0.01$ errors in the circuit, on average. In the second case, we assume that we are able to tolerate a noise rate of $\epsilon = 0.8$ errors in the circuit, on average. This choice of target circuit error rate is motivated by the success of exponential extrapolation in obtaining accurate results for the three-qubit system in the presence of noise of this magnitude. We assume a surface-code decoding cycle time of $1 \mu\text{s}$. We can see from Table I that while the error robustness of the algorithm can reduce the resources required by a factor of around 1.5, this is overshadowed by the exponential improvements arising if the physical error rate can be reduced.

We can also estimate the resources required for a larger system, such as the 29-qubit system. We assume that 50 second-order Trotter steps are needed, which leads to a circuit involving around 2.3×10^7 T gates. We use the same surface-code setup as described above, and consider a physical noise rate of $p = 10^{-3}$, and a target circuit error rate of $\epsilon = 0.8$. We find that a distance 21 surface code is needed, requiring around 51 000 physical qubits. The calculation takes 5234 s to run. On first inspection, this is a modest number of physical qubits, compared to existing resource estimates for solving classically challenging chemistry calculations with quantum hardware. At $p = 10^{-3}$ physical error rates, it has previously been estimated that around 3×10^5 physical qubits would be required to find the ground-state energies of classically challenging

instances of the Fermi-Hubbard model [93], and that over a million physical qubits running for several days may be needed to find the ground state of small molecules relevant for catalysis [108,109]. We note that these resource estimates are upper bounds that may be loose (particularly in the case of the Fermi-Hubbard resource estimates, which use loose analytic bounds on the Trotter error), while our estimates use tighter error bounds obtained from numerical simulations. Moreover, the calculations considered in these other works are likely more challenging than the 29-qubit simulations discussed herein. Nevertheless, these mitigating factors are unlikely to fully account for the reduced resources needed for our approach. We attribute some of these resource savings to the fact that we are simulating spin systems, which may require less overhead than fermionic systems when mapping to qubits. Moreover, we are able to exploit the noise robustness exhibited by our algorithm to further reduce the surface-code distance required.

Unfortunately, there are limitations to the approach taken in our resource estimates above, that cause us to reconsider how these calculations would be performed on error-corrected quantum computers. As seen in Table I, performing a single iteration of even the 11-qubit simulation would take several minutes, assuming a surface-code decoding cycle time of $1 \mu\text{s}$ (which may be optimistic for slower systems, such as trapped ion quantum computers). If we estimate the polarization value through direct sampling of the wave function, then obtaining a precision of 10^{-2} in the polarization would take around 10 000 repetitions. We must then repeat the calculation at a number of simulated time values to perform a single step of our optimization subroutine. In turn, the optimization routine may require hundreds of iterations in order to obtain a good fit for the experimental data. As a result, if the methods discussed in this work are to prove useful for analyzing experimental data, we must find ways to reduce the runtime of the algorithm.

As discussed in Sec. IV D, we can obtain a measurement of the polarization using a combination of amplitude amplification and quantum phase estimation. This approach increases the circuit depth required by a factor of $1/\epsilon_m$ (where ϵ_m is the measurement precision desired), but reduces the number of repetitions required to a constant value. While a full resource estimate of using this technique is beyond the scope of this work, we can perform a rough estimate of how this approach would compare to those discussed above. We assume a value of $\epsilon_m = 10^{-2}$, increasing the circuit depth required by a factor of 100. For the 29-qubit system, this increases our T count estimate to 2.3×10^9 T gates. We neglect the resources required for magic state distillation, and focus on the resources required to reduce the expected number of errors in the logical qubits to less than $\epsilon = 0.01$. Assuming that distilling a magic state still takes $11d$ surface-code cycles, we

need to solve

$$\begin{aligned} & \lceil (1.5 \times 29) + 3 \rceil \times (11d \times 2.3 \times 10^9) \\ & \times 0.1(100p)^{(d+1)/2} < 0.01. \end{aligned} \quad (58)$$

This expression is satisfied for $d \geq 29$. The compact data block for the 29 logical qubits corresponds to around 80 000 physical qubits. With this minimal setup, the calculation would take at least 8 days to complete. As discussed in Ref. [107], we can reduce the runtime of the calculation by exchanging spatial resources for time resources. Adding more qubits enables us to distill and consume magic states more quickly, until we reach the point at which a magic state can be consumed every d surface-code clock cycles. The addition of further qubits allows us to parallelize the implementation of T gates using quantum teleportation. If the T gates required by our algorithm can be fully parallelized into around 8×10^7 layers, and one layer can be implemented per surface-code cycle, a calculation of the polarization value along a single axis, at a single point in time, would finish in 80 s. Building a quantum computer with these capabilities would require billions of qubits [107]—considerably larger than any machines that are planned within the coming decades.

Given the challenges associated with building even a small error-corrected quantum computer, it is natural to ask whether there are other ways to reduce the runtime of the algorithm. First and foremost, we should search for algorithmic improvements. As in the NISQ case, better compilation routines, or the use of more efficient time-evolution algorithms, could reduce the number of gates required to execute our proposal. In particular, we could consider alternative Hamiltonian decompositions that reduce the Trotter error by grouping commuting terms [110], or that exploit the locality of power law interactions [111]. Similarly, we could utilize techniques that exploit symmetries in the Hamiltonian coefficients to reduce the total number of T gates required for the circuit [93,112].

Given the embarrassingly parallel nature of our algorithm, it may have to wait until quantum hardware becomes as cheap and ubiquitous as classical computing is today. This would enable us to distribute each data point and repetition across a number of different processors—in much the same way as the classical emulations in this work are performed. Finally, we note that the challenges discussed in this section are not unique to our proposal, but will be faced by all quantum algorithms that require a large number of repetitions to fulfil their purpose. Perhaps such “multirepetition” algorithms will be impractical, unless we are able to develop improved error-correcting codes, more reliable physical qubits, or reduced code cycle times—which would all improve the effective clock speed of the quantum processor.

VI. CONCLUSION

In this work, we introduce and investigate a quantum algorithm for analyzing the data arising from muon spectroscopy experiments. As discussed in Sec. I, muon spectroscopy is a versatile experimental technique, that has been used to investigate a wide range of physical systems and phenomena. In some cases, accurately analyzing the spectra produced by these experiments requires comparison to data generated by a quantum model for the system. The cost of simulating these quantum models scales exponentially with the size of the simulated system, using all known classical methods. These techniques are particularly necessary when trying to locate the muon rest site, or when studying the effects of muon diffusion.

The quantum algorithm introduced in this work is quite simple, and resembles the (classical) random-phase-approximation-based method introduced to the muon community by Celio [30]. Our algorithm constructs suitable initial states on the quantum computer, evolves them in time, and measures the Z expectation value of the muon qubit, in order to construct simulated polarization functions. The number of gates required by our algorithm scales polynomially with the size of the system, at worst as $\mathcal{O}(N^5)$, as shown in Sec. IV C.

Numerical emulations of our quantum algorithm on classical hardware enable us to bound the error from both Trotterization, and finite sampling of the initial mixed state. We observe that the error in the polarization function resulting from Trotterization is orders of magnitude smaller than existing analytic bounds. We apply these numerical emulations to analyze the data from a muon-spectroscopy experiment on CaF_2 , and observe good agreement with the recent analysis of Wilkinson and Blundell [9] for the same dataset, which uses state-of-the-art classical simulation methods. In the process, we perform the largest simulation employed in muon-spectroscopy analysis to date, with a Hilbert-space size of 2^{29} .

By considering the impact of noise in the quantum computer on our algorithm, we are able to estimate the quantum resources required to perform classically challenging instances of muon spectroscopy analysis. In particular, we observe significant noise robustness in our algorithm, stemming from our aim of extracting a global parameter from the fitted polarization function, rather than targeting any individual data point. These results may find use beyond this work, and suggest that analyzing noisy experimental data may be a good target for future quantum computers.

Nevertheless, our resource estimates highlight the challenges faced by both the algorithm introduced in this work, and many other quantum algorithms. In the context of NISQ simulations, we observe that the gate counts produced by our algorithm are likely too large to simulate on devices with realistic noise rates, even if error-mitigation techniques are utilized. Similar challenges have been

noted previously for NISQ electronic structure calculations [113]. While noise can ultimately be overcome using quantum error correction, challenges also remain in this arena. At first glance, our algorithm appeared to require fewer fault-tolerant resources than solving challenging instances of the electronic structure problem. However, the large number of repetitions required by our algorithm (to estimate observables, calculate multiple datapoints, and repeat the calculation as part of an optimization loop) result in an impractically long runtime. While this runtime can be reduced by increasing the size of the quantum computer, or parallelizing the calculation across multiple quantum machines, these both come at significant cost. We argue that this is not necessarily a limitation of our algorithm, but a challenge facing many quantum algorithms. For example, while Ref. [84] performed a more rigorous analysis of the gate counts required to perform classically intractable simulations of spin-1/2 systems than this work, it stopped short of estimating the resources required to estimate a given observable. Their gate counts, on the order of 10^8 T gates, are similar to those found in this work. As a result, they will run into the same problems that we have, when taking into account the resources required to estimate a given observable, or to sample a range of conditions. Similarly, while existing estimates for solving the electronic structure problem on small fault-tolerant quantum processors [93,108,109,114,115] show that we can obtain the energy of a classically intractable system in hours or days, these works typically stop short of considering what problem that actually solves. In order to optimize a molecular geometry, estimate other observables on the ground state, or elucidate a phase diagram, these calculations will likely have to be repeated many times—resulting in what may be a prohibitively long calculation time. As such, we stress that it is important to consider the total resources required to solve a given problem with a quantum computer, and not just the quantum resources required to run the corresponding circuits. The embarrassingly parallel nature of many classical algorithms (including the classical emulations in this work) and the low cost and ubiquity of classical hardware, will place stringent requirements on the performance of future quantum algorithms. Similar arguments were recently made by Babbush *et al.* [116], in the context of whether quadratic quantum speedups will be sufficient to show quantum advantage on realistic problem instances.

There are a number of optimizations that could be introduced to our algorithm, which would help to make it more practical for muon-data analysis. In particular, better compilation of our time-evolution circuits, or the use of more efficient time-evolution algorithms, could dramatically reduce the number of gates required to achieve a given accuracy. It will also be interesting to consider which real-world systems would be most interesting to apply our algorithm to. As discussed in Sec. II, many systems can be

satisfactorily analyzed using mean-field-type methods. As such, it is important to engage with the muon community to elucidate which systems appear difficult to analyze with classical techniques. One possible avenue for exploration, would be the application of our algorithm to the quantum model of muon diffusion discussed in Sec. II [see Eq. (9)], which has not yet been employed for data analysis, due to the large Hilbert space of the resulting simulation. Liaising with the muon community will also enable the introduction of problem-specific optimizations of the algorithm. For example, one could consider incorporating the ideas of Ref. [9], which scaled the interaction strengths of more distant nuclei to act as a proxy for the rest of the sample.

A related question is whether the quantum algorithm introduced in this work can be simulated efficiently using a classical computer. The environment being in a mixed initial state may motivate the belief that approximate classical methods may be able to efficiently simulate this problem. While answering this question is beyond the scope of this work, our approach has clear links to the “one clean qubit” (DQC1) model of quantum computing [117], which has so far resisted efficient classical simulation [118–121].

Finally, it is interesting to ask if other, more complex quantum algorithms could be applied to analyzing muon-spectroscopy data. One possibility could be to use quantum read-only memory (QROM) [114] to load the experimental data values into our quantum computer. We could then attempt to compute all of the simulated data in superposition [the state would resemble $\sum_t |t\rangle |P(t)\rangle$, where t denotes the simulated time], and use existing quantum machine learning algorithms to extract fitting parameters of interest. While the details of this approach would likely be more complicated than outlined here, we note that this is a slightly different approach than is typically considered in quantum machine-learning algorithms. Here, we exploit the exponential speedup in calculating $P(t)$ (and potentially polynomial speedups in parameter fitting), and do not mind that there is no speedup for loading in the data. We would seek to load in a number of datapoints that is generally constant for a given type of muon-spectroscopy experiment (as it is defined by the characteristics of the beam type used). Given the numerous possible avenues for exploration, we believe that both muon spectroscopy, and the analysis of data arising from other experiments underpinned by quantum mechanics, are promising targets for future quantum computers.

ACKNOWLEDGMENTS

This work is supported by the EPSRC National Quantum Technology Hub in Networked Quantum Information Technology (EP/M013243/1) and the QCS Hub (EP/T001062/1). We thank S. Blundell and J. Wilkinson for kindly providing the μ^+ SR experimental data analyzed in this work, and for helpful discussions on this project,

paper, and μ^+ SR experiments. We are grateful to S. Benjamin and T. Jones for useful discussions on this work, and for the implementation of functions for Trotterized time evolution in QuEST. We thank Z. Cai for useful feedback and suggestions on this paper. We also thank X. Yuan, Y. Su, D. Litinski, J. McClean, B. Koczor, O. Higgott, and E. Campbell for helpful conversations on various aspects of this work. We acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out this work.

APPENDIX A: SIMULATION DETAILS AND ADDITIONAL NUMERICAL RESULTS

The simulations presented in this work are performed using Cirq [98] and QuEST [97], as described in the main text. QuEST simulations are run on CPU nodes containing two Intel Xeon 8268 processors, each with 24 cores, with a total of 384 GB of RAM. For simulations with over 21 qubits, OpenMP is used to parallelize over the 48 available cores. Each datapoint is run on a separate node, to parallelize the calculation of the polarization function. The largest calculations, utilizing 29 qubits, 30 second-order Trotter steps, and a single sample, takes approximately 4 days to run.

In Fig. 16 we present the alternative fit to the experimental data with the 21-qubit system, discussed in the main text.

APPENDIX B: PREPARING DICKE STATES ON QUANTUM COMPUTERS

In this section, we discuss the method developed by Bärtschi and Eidenbenz [79] to prepare Dicke states on quantum computers. As discussed in the main text, to prepare a Dicke state with Hamming weight h , acting on n qubits, the algorithm requires $\mathcal{O}(kn)$ gates, $\mathcal{O}(n)$ depth, and n qubits. Here, we present a slightly less rigorous, but more pedagogical overview of the algorithm introduced in Ref. [79], and refer the reader to the original reference for more information.

The state-preparation algorithm proceeds recursively, making use of the following expression for Dicke states:

$$|D_h^n\rangle = \sqrt{\frac{h}{n}} |D_{h-1}^{n-1}\rangle \otimes |1\rangle + \sqrt{\frac{n-h}{n}} |D_h^{n-1}\rangle \otimes |0\rangle. \quad (\text{B1})$$

We assume the existence of a unitary operator $U_{n,k}$ such that $U_{n,k} |0\rangle^{\otimes n-h} |1\rangle^{\otimes h} = |D_h^n\rangle$ for all $h \leq k$. As we show below, this operator exists, and can be constructed from typical single- and two-qubit gates.

As a first step, we note that

$$|D_h^n\rangle = U_{n,k} |0\rangle^{\otimes n-h} |1\rangle^{\otimes h}, \quad (\text{B2})$$

and

$$\begin{aligned} |D_h^n\rangle &= \sqrt{\frac{h}{n}} |D_{h-1}^{n-1}\rangle \otimes |1\rangle + \sqrt{\frac{n-h}{n}} |D_h^{n-1}\rangle \otimes |0\rangle, \\ &= \left(U_{n-1,k} \otimes I \right) \left[\sqrt{\frac{h}{n}} |0\rangle^{\otimes n-h} |1\rangle^{\otimes h} \right. \\ &\quad \left. + \sqrt{\frac{n-h}{n}} |0\rangle^{\otimes n-1-h} |1\rangle^{\otimes h} |0\rangle \right], \\ &= \left(U_{n-1,k} \otimes I \right) \cdot \left(I_{n-k-1} \otimes V_{n,k} \right) |0\rangle^{\otimes n-h} |1\rangle^{\otimes h}, \end{aligned} \quad (\text{B3})$$

with

$$\begin{aligned} &\left(I_{n-k-1} \otimes V_{n,k} \right) |0\rangle^{\otimes n-h} |1\rangle^{\otimes h} \\ &= \sqrt{\frac{h}{n}} |0\rangle^{\otimes n-h} |1\rangle^{\otimes h} \\ &\quad + \sqrt{\frac{n-h}{n}} |0\rangle^{\otimes n-1-h} |1\rangle^{\otimes h} |0\rangle. \end{aligned} \quad (\text{B4})$$

These relations imply that

$$U_{n,k} = \left(U_{n-1,k} \otimes I \right) \cdot \left(I_{n-k-1} \otimes V_{n,k} \right). \quad (\text{B5})$$

Here, $U_{n-1,k}$ acts on the leftmost $n-1$ qubits. $V_{n,k}$ is defined on n qubits, but acts trivially on the first $n-k-1$ qubits, so can be considered to act only on the final $k+1$ qubits. We can then recurse this relationship

$$U_{n-1,k} = \left(U_{n-2,k} \otimes I \right) \cdot \left(I_{n-k-2} \otimes V_{n-1,k} \right), \quad (\text{B6})$$

where $V_{n-1,k}$ acts on the leftmost $k+1$ of the final $k+2$ qubits in the state. This enables us to write that

$$\begin{aligned} U_{n,k} &= \left(U_{k+1,k} \otimes I_{n-k-1} \right) \cdots \left(I_{n-k-2} \otimes V_{n-1,k} \otimes I \right) \\ &\quad \cdot \left(I_{n-k-1} \otimes V_{n,k} \right). \end{aligned} \quad (\text{B7})$$

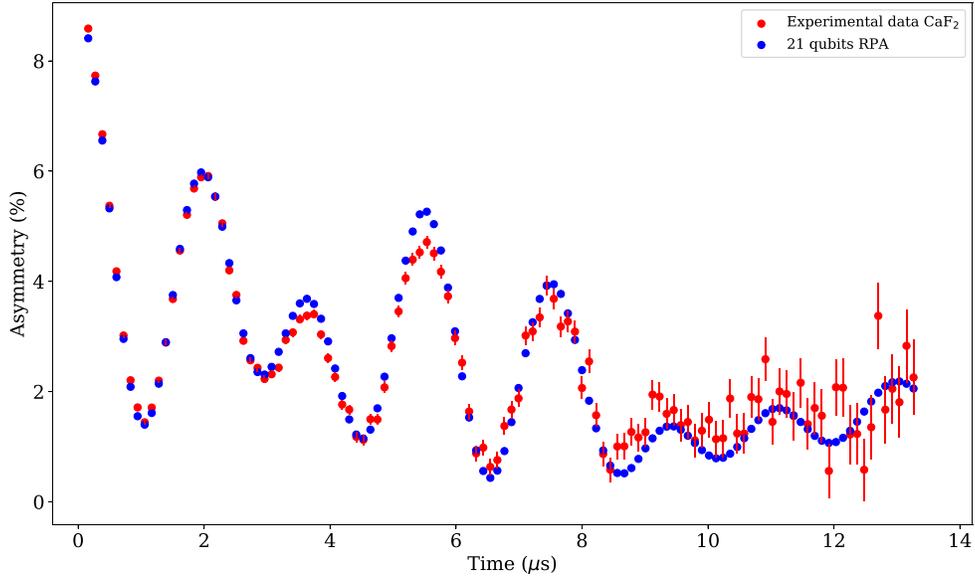


FIG. 16. The $\text{CaF}_2 + \mu^+$ geometry is fixed in this simulation, and given by the density-functional-theory results of Ref. [9]. The nearest-neighbor fluorines are moved towards the muon by 0.19 \AA , and the next-nearest-neighbor fluorines are moved towards the muon by 0.027 \AA . The NNNN fluorines are at their equilibrium positions. We use a 21-qubit simulation with the random-phase-approximation method (with 48 samples per data point, and 40 second-order Trotter steps) to generate a polarization function for this geometry, $P(t)$. This polarization function is fitted to the experimental asymmetry data obtained in Ref. [9], using $A(t) = A_0 P(t) + A_{\text{bg}}$, where A_0 and A_{bg} are fitting parameters. The fit is performed using the Levenberg-Marquardt algorithm to minimize the normalized square residuals. The reduced χ^2 value of the generated data is 4.44.

We can repeat our analysis above for $U_{k+1,k}$, considering just the $k+1$ qubits acted upon

$$\begin{aligned}
 U_{k+1,k} |0\rangle |1\rangle^{\otimes k} &= |D_k^{k+1}\rangle, = \sqrt{\frac{k}{k+1}} |D_{k-1}^k\rangle \otimes |1\rangle + \sqrt{\frac{1}{k+1}} |1\rangle^{\otimes k} \otimes |0\rangle, \\
 &= \left(U_{k,k-1} \otimes I \right) \left[\sqrt{\frac{k}{k+1}} |0\rangle |1\rangle^{\otimes k-1} \otimes |1\rangle + \sqrt{\frac{1}{k+1}} |1\rangle^{\otimes k} \otimes |0\rangle \right], \\
 &= \left(U_{k,k-1} \otimes I \right) \cdot \left(V_{k+1,k} \right) |0\rangle |1\rangle^{\otimes k-1} \otimes |1\rangle, \tag{B8}
 \end{aligned}$$

where $V_{k+1,k}$ acts on all of the $k+1$ qubits. We can then continue the recursion above

$$\begin{aligned}
 U_{n,k} &= \left(U_{k+1,k} \otimes I_{n-k-1} \right) \cdots \left(I_{n-k-2} \otimes V_{n-1,k} \otimes I \right) \times \left(I_{n-k-1} \otimes V_{n,k} \right) = \left(U_{k,k-1} \otimes I_{n-k} \right) \times \left(V_{k+1,k} \otimes I_{n-k-1} \right) \cdots \\
 &\quad \times \left(I_{n-k-2} \otimes V_{n-1,k} \otimes I \right) \times \left(I_{n-k-1} \otimes V_{n,k} \right) = \cdots \\
 &= \left(V_{2,1} \otimes I_{n-2} \right) \times \left(V_{3,2} \otimes I_{n-3} \right) \cdots \left(V_{k,k-1} \otimes I_{n-k} \right) \\
 &\quad \times \left(V_{k+1,k} \otimes I_{n-k-1} \right) \times \left(I \otimes V_{k+2,k} \otimes I_{n-k-2} \right) \cdots \left(I_{n-k-2} \otimes V_{n-1,k} \otimes I \right) \times \left(I_{n-k-1} \otimes V_{n,k} \right). \tag{B9}
 \end{aligned}$$

We can express this more concisely via the expression in Lemma 2 of Ref. [79]:

$$U_{n,k} = \prod_{l=2}^k \left(V_{l,l-1} \otimes I_{n-l} \right) \cdot \prod_{l=k+1}^n \left(I_{l-k-1} \otimes V_{l,k} \otimes I_{n-l} \right). \quad (\text{B10})$$

This confirms the existence of the gate $U_{n,k}$, if we are able to construct the unitary $V_{n,k}$ for arbitrary n, k . As discussed in Ref. [79], we can construct $V_{n,k}$ from standard single-, two-, and three-qubit gates as follows.

We have that

$$\begin{aligned} \left(V_{n,k} \right) |0\rangle^{\otimes k+1-h} |1\rangle^{\otimes h} &= \sqrt{\frac{h}{n}} |0\rangle^{\otimes k+1-h} |1\rangle^{\otimes h} \\ &+ \sqrt{\frac{n-h}{n}} |0\rangle^{\otimes k-h} |1\rangle^{\otimes h} |0\rangle. \end{aligned} \quad (\text{B11})$$

This unitary only changes the value of the zeroth qubit from the right, and the h th qubit from the right. First, consider the circuit shown in Fig. 17. If we set $\theta = 2\cos^{-1}(\sqrt{1/n})$, then the circuit applies the unitary

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \sqrt{\frac{1}{n}} & -\sqrt{\frac{n-1}{n}} & 0 \\ 0 & \sqrt{\frac{n-1}{n}} & \sqrt{\frac{1}{n}} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (\text{B12})$$

Similarly, the circuit in Fig. 18 with $\theta = 2\cos^{-1}(\sqrt{\alpha/n})$ acts trivially on all input states, except for

$$\begin{aligned} |011\rangle &\rightarrow \sqrt{\frac{\alpha}{n}} |011\rangle + \sqrt{\frac{n-\alpha}{n}} |110\rangle, \\ |110\rangle &\rightarrow -\sqrt{\frac{n-\alpha}{n}} |011\rangle + \sqrt{\frac{\alpha}{n}} |110\rangle. \end{aligned} \quad (\text{B13})$$

We can now use these building blocks to construct the unitary $V_{n,k}$. We first apply the gate in Fig. 17 to the final two qubits in the register, and then repeatedly apply the gate W_n^α shown in Fig. 18, incrementing the value of α each time. As we are applying these gates on the state $|0\rangle^{\otimes n-h} |1\rangle^{\otimes h}$, the first h qubits encountered are all in the $|1\rangle$ state, so these gates act trivially. After these h qubits, the three-qubit W_n^h gate encounters the state $|0\rangle_{n-h} |1\rangle_{n-h+1} |1\rangle_n$, which it transforms into $\sqrt{(h/n)} |0\rangle_{n-h} |1\rangle_{n-h+1} |1\rangle_n + \sqrt{(n-h/n)} |1\rangle_{n-h} |1\rangle_{n-h+1} |0\rangle_n$. The subsequent $W_n^{\alpha>h}$ gates act trivially on both branches of this superposition. Taken together, these gates carry out the transform in Eq. (B11), for all $h \leq k$. It is shown in Ref. [79] how the three-qubit W gate can be decomposed into single- and two-qubit gates, and how the circuit can be implemented with $\mathcal{O}(kn)$ gates, $\mathcal{O}(n)$ depth, and n qubits on a linear array of qubits, with nearest-neighbor connectivity.

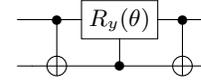


FIG. 17. A circuit that yields the unitary operator given by Eq. (B12).

APPENDIX C: BOUNDING THE VARIANCE OF THE SAMPLING SCHEMES

As discussed in Sec. IV B, there are two approaches that can be used to effectively sample from the maximally mixed environment state. The first, which we refer to in this Appendix as scheme A, samples environment basis vectors $|k\rangle_e$ with equal probability $(1/D_e)$. The second approach, referred to here as scheme B, is the random-phase-approximation method introduced in Ref. [30], and discussed in Eq. (8). In this Appendix, we show that the expectation values of these two schemes are equivalent to sampling from the maximally mixed state, and calculate the variance of both schemes.

1. Scheme A

We sample a random variable $\langle 0|_\mu \langle j|_e Z(t) |0\rangle_\mu |j\rangle_e := \alpha_{jj}(j)$ (that depends on the discrete random variable $j \in [0, D_e)$, with probability $p(j) = (1/D_e)$). We define $Z(t) = e^{iHt} Z_\mu e^{-iHt}$. The expectation value of the random variable is

$$\begin{aligned} \mathbb{E}[\alpha_{jj}(j)] &= \frac{1}{D_e} \sum_j \langle 0|_\mu \langle j|_e Z(t) |0\rangle_\mu |j\rangle_e, \\ &= \text{Tr} \left[e^{-iHt} \left(|0\rangle \langle 0|_\mu \otimes \frac{I_e}{D_e} \right) e^{iHt} Z_\mu \right], \\ &= P(t), \end{aligned} \quad (\text{C1})$$

as required. The variance of the random variable is given by

$$\begin{aligned} \text{Var}[\alpha_{jj}(j)] &= \mathbb{E}[\alpha_{jj}(j)^2] - \mathbb{E}[\alpha_{jj}(j)]^2, \\ &= \left(\frac{1}{D_e} \sum_j \alpha_{jj}^2 \right) - \left(\frac{1}{D_e^2} \sum_{j,k} \alpha_{jj} \alpha_{kk} \right). \end{aligned} \quad (\text{C2})$$

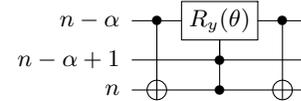


FIG. 18. A circuit that yields the unitary operator, which we refer to as W_n^α . This operator acts trivially on all input states, except those noted in Eq. (B13).

This can be bounded as

$$\text{Var}[\alpha_{jj}(j)] = \frac{1}{D_e} \sum_j \alpha_{jj}^2 - P(t)^2 \leq \frac{1}{D_e} \sum_j \alpha_{jj}^2 \leq 1. \quad (\text{C3})$$

If half of $\alpha_{jj} = 1$, and half are -1 , then this bound is saturated.

2. Scheme B

We first define $\alpha_{jk} := \langle 0 |_\mu \langle k |_e Z(t) | 0 \rangle_\mu | j \rangle_e$. Unlike in scheme A, α_{jk} is not a random variable. We then define a random variable $X(\vec{\theta}) := (1/D_e) \sum_{j,k} e^{i\theta_j} e^{-i\theta_k} \alpha_{jk}$, that is a function of the random variables $\vec{\theta} = \{\theta_1, \dots, \theta_i, \dots, \theta_{D_e}\}$ that are chosen uniformly at random in the range $[0, 2\pi)$. We assume that these random variables are independent here, which will be a reasonable assumption when the circuit used to prepare the initial state is sufficiently

randomizing. The expectation value of the random variable X is

$$\begin{aligned} \mathbb{E}[X(\vec{\theta})] &= \frac{1}{D_e} \mathbb{E} \left[\sum_{j,k} e^{i\theta_j} e^{-i\theta_k} \alpha_{jk} \right], \\ &= \frac{1}{D_e} \sum_j \alpha_{jj} + \frac{1}{D_e} \sum_{j \neq k} \mathbb{E}[e^{i\theta_j} e^{-i\theta_k}] \alpha_{jk}, \\ &= \frac{1}{D_e} \sum_j \alpha_{jj} + \frac{1}{D_e} \left(\frac{1}{2\pi} \int_{\theta_j=0}^{2\pi} e^{i\theta_j} d\theta_j \right) \\ &\quad \times \left(\frac{1}{2\pi} \int_{\theta_k=0}^{2\pi} e^{-i\theta_k} d\theta_k \right) \alpha_{jk}, \\ &= \frac{1}{D_e} \sum_j \alpha_{jj} = \frac{1}{D_e} \sum_j \langle 0 |_\mu \langle j |_e Z(t) | 0 \rangle_\mu | j \rangle_e \\ &= P(t), \end{aligned} \quad (\text{C4})$$

as required. The variance is given by

$$\begin{aligned} \text{Var}[X(\vec{\theta})] &= \mathbb{E}[|X(\vec{\theta})|^2] - |\mathbb{E}[X(\vec{\theta})]|^2 \\ &= \frac{1}{D_e^2} \sum_{j,k,l,m} \left(\frac{1}{(2\pi)^4} \int_{\theta_j, \theta_k, \theta_l, \theta_m=0}^{2\pi} e^{i\theta_j} e^{-i\theta_k} e^{-i\theta_l} e^{i\theta_m} d\theta_j d\theta_k d\theta_l d\theta_m \right) \alpha_{jk} \alpha_{lm}^* - P(t)^2. \end{aligned} \quad (\text{C5})$$

The integral evaluates to zero unless $j = k$ and $l = m$ or $j = l$ and $k = m$ (but $j \neq k$). Thus,

$$\begin{aligned} \text{Var}[X(\vec{\theta})] &= \left(\frac{1}{D_e} \sum_j \alpha_{jj} \right) \left(\frac{1}{D_e} \sum_l \alpha_{ll} \right) + \frac{1}{D_e^2} \sum_{j \neq k} \alpha_{jk} \alpha_{jk}^* - P(t)^2, \\ &= \frac{1}{D_e^2} \sum_{j \neq k} |\alpha_{jk}|^2. \end{aligned} \quad (\text{C6})$$

We can bound this as [defining $Z(t) | 0 \rangle_\mu | k \rangle_e := |\phi_k\rangle$]

$$\begin{aligned} \frac{1}{D_e^2} \sum_{j \neq k} |\alpha_{jk}|^2 &\leq \frac{1}{D_e^2} \sum_{j,k} |\alpha_{jk}|^2, \\ &= \frac{1}{D_e^2} \sum_k \langle 0 |_\mu \langle k |_e Z(t) \left(| 0 \rangle_\mu \langle 0 |_\mu \otimes \sum_j | j \rangle_e \langle j |_e \right) Z(t) | 0 \rangle_\mu | k \rangle_e, \\ &= \frac{1}{D_e^2} \sum_k \langle \phi_k | (| 0 \rangle_\mu \langle 0 |_\mu \otimes I_e) | \phi_k \rangle, \\ &\leq \frac{1}{D_e}. \end{aligned} \quad (\text{C7})$$

This bound is independent of the values α_{jk} and shows that the variance is suppressed exponentially in the number of environment spins.

- [1] J. P. F. LeBlanc *et al.*, (Simons Collaboration on the Many-Electron Problem), Solutions of the Two-Dimensional Hubbard Model: Benchmarks and Results from a Wide Range of Numerical Algorithms, *Phys. Rev. X* **5**, 041041 (2015).
- [2] K. T. Williams *et al.*, (Simons Collaboration on the Many-Electron Problem), Direct Comparison of Many-Body Methods for Realistic Electronic Hamiltonians, *Phys. Rev. X* **10**, 011041 (2020).
- [3] J. McClean, S. Boixo, V. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9**, 4812 (2018).
- [4] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, Cost-function-dependent barren plateaus in shallow quantum neural networks, [arXiv:2001.00550](https://arxiv.org/abs/2001.00550) (2020).
- [5] N. Wiebe, C. Granade, C. Ferrie, and D. G. Cory, Hamiltonian Learning and Certification Using Quantum Resources, *Phys. Rev. Lett.* **112**, 190501 (2014).
- [6] I. Cong, S. Choi, and M. D. Lukin, Quantum convolutional neural networks, *Nat. Phys.* **15**, 1273 (2019).
- [7] A. Yaouanc and P. D. De Reotier, *Muon Spin Rotation, Relaxation, and Resonance: Applications to Condensed Matter* (Oxford University Press, Oxford, UK, 2011), Vol. 147.
- [8] M. Celio and P. F. Meier, Exact calculation of the muon polarization function, *Hyperfine Interact.* **18**, 435 (1984).
- [9] J. M. Wilkinson and S. J. Blundell, Information and Decoherence in a Muon-Fluorine Coupled System, *Phys. Rev. Lett.* **125**, 087201 (2020).
- [10] J. Lord, S. Cottrell, and W. Williams, Muon spin relaxation in strongly coupled systems, *Phys. B: Condens. Matter* **289**, 495 (2000).
- [11] A. Chiesa, F. Tacchino, M. Grossi, P. Santini, I. Tavernelli, D. Gerace, and S. Carretta, Quantum hardware simulating four-dimensional inelastic neutron scattering, *Nat. Phys.* **15**, 455 (2019).
- [12] D. Sels, H. Dashti, S. Mora, O. Demler, and E. Demler, Quantum approximate Bayesian computation for NMR model inference, [arXiv:1910.14221](https://arxiv.org/abs/1910.14221) (2019).
- [13] I. Gurevich, E. Mel'eshko, I. Muratova, B. Nikol'sky, V. Roganov, V. Selivanov, and B. Sokolov, Dipole interactions and diffusion of μ^+ meson in copper, *Phys. Lett. A* **40**, 143 (1972).
- [14] R. L. Garwin, L. M. Lederman, and M. Weinrich, Observations of the failure of conservation of parity and charge conjugation in meson decays: The magnetic moment of the free muon, *Phys. Rev.* **105**, 1415 (1957).
- [15] S. J. Blundell, Spin-polarized muons in condensed matter physics, *Contemp. Phys.* **40**, 175 (1999).
- [16] T. Lancaster, B. M. Huddart, R. C. Williams, F. Xiao, K. J. A. Franke, P. J. Baker, F. L. Pratt, S. J. Blundell, J. A. Schlueter, M. B. Mills, A. C. Maahs, and K. E. Preuss, Probing magnetic order and disorder in the one-dimensional molecular spin chains $\text{CuF}_2(\text{pyz})$ and $[\text{Ln}(\text{hfac})_3(\text{boaDTDA})]_n$ ($\text{Ln} = \text{Sm}, \text{La}$) using implanted muons, *J. Phys. Condens. Matter* **31**, 394002 (2019).
- [17] V. G. Storchak and N. V. Prokof'ev, Quantum diffusion of muons and muonium atoms in solids, *Rev. Mod. Phys.* **70**, 929 (1998).
- [18] J. Sugiyama, K. Mukai, Y. Ikedo, H. Nozaki, M. Månsson, and I. Watanabe, Li Diffusion in Li_xCoO_2 Probed by Muon-Spin Spectroscopy, *Phys. Rev. Lett.* **103**, 147601 (2009).
- [19] J. Sugiyama, H. Nozaki, I. Umegaki, Y. Higuchi, S. P. Cottrell, S. Hori, R. Kanno, and M. Mansson, in *Proceedings of the 14th International Conference on Muon Spin Rotation, Relaxation and Resonance* ($\mu\text{SR}2017$), <https://journals.jps.jp/doi/pdf/10.7566/JPSCP.21.011016>.
- [20] I. Umegaki, S. Kawauchi, H. Nozaki, H. Sawada, H. Nakano, M. Harada, S. P. Cottrell, F. C. Coomer, M. Telling, and J. Sugiyama, in *Proceedings of the 14th International Conference on Muon Spin Rotation, Relaxation and Resonance* ($\mu\text{SR}2017$), <https://journals.jps.jp/doi/pdf/10.7566/JPSCP.21.011017>.
- [21] A. D. Pant, K. Nagamine, I. Shiraki, E. Torikai, K. Shimomura, F. L. Pratt, H. Ariga, K. Ishida, and J. S. Schultz, Muonium response to oxygen content in biological aqueous solutions for cancer research, *J. Phys. Conf. Ser.* **551**, 012043 (2014).
- [22] M. Celio, Corrections to the strong collision model, *Hyperfine Interact.* **31**, 153 (1986).
- [23] E. Holzschuh and P. F. Meier, Zero-field spin relaxation of positive muons, *Phys. Rev. B* **29**, 1129 (1984).
- [24] R. Kubo and T. Toyabe, in *Magnetic Resonance and Relaxation*, edited by R. Blinc (North-Holland, Amsterdam, 1966).
- [25] P. D. de Réotier, A. Yaouanc, and S. Meshkov, Effect of nuclear spin dynamics on the μsr zero field depolarisation function, *Phys. Lett. A* **162**, 206 (1992).
- [26] P. D. de Reotier and A. Yaouanc, Quantum calculation of the muon depolarization function: Effect of spin dynamics in nuclear dipole systems, *J. Phys. Condens. Matter* **4**, 4533 (1992).
- [27] M. Celio and P. F. Meier, Spin relaxation of positive muons due to dipolar interactions, *Phys. Rev. B* **27**, 1908 (1983).
- [28] T. Lancaster, S. J. Blundell, P. J. Baker, M. L. Brooks, W. Hayes, F. L. Pratt, J. L. Manson, M. M. Conner, and J. A. Schlueter, Muon-Fluorine Entangled States in Molecular Magnets, *Phys. Rev. Lett.* **99**, 267601 (2007).
- [29] T. Lancaster, F. L. Pratt, S. J. Blundell, I. McKenzie, and H. E. Assender, Muon-fluorine entanglement in fluoropolymers, *J. Phys. Condens. Matter* **21**, 346004 (2009).
- [30] M. Celio, New Method to Calculate the Muon Polarization Function, *Phys. Rev. Lett.* **56**, 2720 (1986).
- [31] P. Dalmas de Réotier, A. Yaouanc, and J. P. Boucher, Model dependence of the longitudinal μsr depolarisation function in nuclear dipole systems, *Hyperfine Interact.* **65**, 1121 (1991).
- [32] W. Huang, V. Pacradouni, M. P. Kennett, S. Komiya, and J. E. Sonier, Precision search for magnetic order in the pseudogap regime of $\text{La}_{2-x}\text{Sr}_x\text{CuO}_4$ by muon spin relaxation, *Phys. Rev. B* **85**, 104527 (2012).
- [33] S. Blundell, A. Steele, T. Lancaster, J. Wright, and F. Pratt, A Bayesian approach to magnetic moment determination using μSR , *Phys. Procedia* **30**, 113 (2012), 12th International Conference on Muon Spin Rotation, Relaxation and Resonance ($\mu\text{SR}2011$).
- [34] P. Bonfà and R. De Renzi, Toward the computational prediction of muon sites and interaction parameters, *J. Phys. Soc. Jpn* **85**, 091014 (2016).

- [35] F. Bernardini, P. Bonfà, S. Massidda, and R. De Renzi, Ab initio strategy for muon site assignment in wide band gap fluorides, *Phys. Rev. B* **87**, 115148 (2013).
- [36] I. J. Onuorah, P. Bonfà, R. De Renzi, L. Monacelli, F. Mauri, M. Calandra, and I. Errea, Quantum effects in muon spin spectroscopy within the stochastic self-consistent harmonic approximation, *Phys. Rev. Mater.* **3**, 073804 (2019).
- [37] J. S. Möller, P. Bonfà, D. Ceresoli, F. Bernardini, S. J. Blundell, T. Lancaster, R. D. Renzi, N. Marzari, I. Watanabe, S. Sulaiman, and M. I. Mohamed-Ibrahim, Playing quantum hide-and-seek with the muon: Localizing muon stopping sites, *Phys. Scr.* **88**, 068510 (2013).
- [38] J. S. Möller, D. Ceresoli, T. Lancaster, N. Marzari, and S. J. Blundell, Quantum states of muons in fluorides, *Phys. Rev. B* **87**, 121108 (2013).
- [39] G. M. Luke, J. H. Brewer, S. R. Kreitzman, D. R. Noakes, M. Celio, R. Kadono, and E. J. Ansaldo, Muon diffusion and level crossing in copper, *Hyperfine Interact.* **64**, 721 (1991).
- [40] J. H. Brewer, S. R. Kreitzman, D. R. Noakes, E. J. Ansaldo, D. R. Harshman, and R. Keitel, Observation of muon-fluorine “hydrogen bonding” in ionic crystals, *Phys. Rev. B* **33**, 7813 (1986).
- [41] K. Nishiyama, S. Nishiyama, and W. Higemoto, Asymmetric $f \mu f$ interaction of the muon in polyfluorocarbons, *Phys. B: Condens. Matter* **326**, 41 (2003).
- [42] D. Noakes, E. Ansaldo, S. Kreitzman, and G. Luke, The ($f \mu f$)-ion in solid fluorides, *J. Phys. Chem. Solids* **54**, 785 (1993).
- [43] R. Kadono, J. Imazato, T. Matsuzaki, K. Nishiyama, K. Nagamine, T. Yamazaki, D. Richter, and J.-M. Welter, Quantum diffusion of positive muons in copper, *Phys. Rev. B* **39**, 23 (1989).
- [44] G. M. Luke, J. H. Brewer, S. R. Kreitzman, D. R. Noakes, M. Celio, R. Kadono, and E. J. Ansaldo, Muon diffusion and spin dynamics in copper, *Phys. Rev. B* **43**, 3284 (1991).
- [45] R. P. Feynman, Simulating physics with computers, *Int. J. Theor. Phys.* **21**, 467 (1982).
- [46] D. Deutsch, in *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* (The Royal Society, 1985), Vol. 400, p. 97.
- [47] P. W. Shor, in *Foundations of Computer Science, 1994 Proceedings., 35th Annual Symposium on* (Ieee, Santa Fe, New Mexico, 1994), p. 124.
- [48] L. K. Grover, in *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing* (ACM, Philadelphia, Pennsylvania, 1996), p. 212.
- [49] S. Lloyd, Universal quantum simulators, *Science* **273**, 1073 (1996).
- [50] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O’Brien, Quantum computers, *Nature* **464**, 45 EP (2010), [review Article](#).
- [51] A. Shnirman, G. Schön, and Z. Hermon, Quantum Manipulations of Small Josephson Junctions, *Phys. Rev. Lett.* **79**, 2371 (1997).
- [52] Y. Nakamura, Y. A. Pashkin, and J. S. Tsai, Coherent control of macroscopic quantum states in a single-cooper-pair box, *Nature* **398**, 786 (1999).
- [53] E. Knill, R. Laflamme, and G. J. Milburn, A scheme for efficient quantum computation with linear optics, *Nature* **409**, 46 (2001).
- [54] J. I. Cirac and P. Zoller, Quantum Computations with Cold Trapped Ions, *Phys. Rev. Lett.* **74**, 4091 (1995).
- [55] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland, Quantum dynamics of single trapped ions, *Rev. Mod. Phys.* **75**, 281 (2003).
- [56] J. Preskill, Quantum computing in the nisq era and beyond, [arXiv:1801.00862](#) (2018).
- [57] F. Arute *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
- [58] M. A. Nielsen and I. Chuang, Quantum computation and quantum information (2002).
- [59] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, *Phys. Rev. A* **86**, 032324 (2012).
- [60] J. O’Gorman and E. T. Campbell, Quantum computation with realistic magic-state factories, *Phys. Rev. A* **95**, 032338 (2017).
- [61] E. T. Campbell, B. M. Terhal, and C. Vuillot, Roads towards fault-tolerant universal quantum computation., *Nature* **549**, 172 (2017).
- [62] B. M. Terhal, Quantum error correction for quantum memories, *Rev. Mod. Phys.* **87**, 307 (2015).
- [63] S. J. Devitt, W. J. Munro, and K. Nemoto, Quantum error correction for beginners, *Rep. Prog. Phys.* **76**, 076001 (2013).
- [64] R. Raussendorf, Key ideas in quantum error correction, *Philos. Trans. R. Soc. London A: Math. Phys. Eng. Sci.* **370**, 4541 (2012), <http://rsta.royalsocietypublishing.org/content/370/1975/4541.full.pdf>.
- [65] D. A. Lidar and T. A. Brun, *Quantum Error Correction* (Cambridge University Press, Cambridge, UK, 2013).
- [66] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, Quantum computational chemistry, *Rev. Mod. Phys.* **92**, 015003 (2020).
- [67] Y. Li and S. C. Benjamin, Efficient Variational Quantum Simulator Incorporating Active Error Minimization, *Phys. Rev. X* **7**, 021050 (2017).
- [68] S. Endo, S. C. Benjamin, and Y. Li, Practical quantum error mitigation for near-future applications, [arXiv:1712.09271](#) (2017).
- [69] K. Temme, S. Bravyi, and J. M. Gambetta, Error Mitigation for Short-Depth Quantum Circuits, *Phys. Rev. Lett.* **119**, 180509 (2017).
- [70] T. Giurgica-Tiron, Y. Hindy, R. LaRose, A. Mari, and W. J. Zeng, Digital zero noise extrapolation for quantum error mitigation, [arXiv:2005.10921](#) (2020).
- [71] Z. Cai, Multi-exponential error extrapolation and combining error mitigation techniques for nisq applications, [arXiv:2007.01265](#) (2020).
- [72] A. Kandala, K. Temme, A. D. Corcoles, A. Mezzacapo, J. M. Chow, and J. M. Gambetta, Extending the computational reach of a noisy superconducting quantum processor, [arXiv:1805.04492](#) (2018).
- [73] N. P. D. Sawaya, T. Menke, T. H. Kyaw, S. Johri, A. Aspuru-Guzik, and G. G. Guerreschi, Resource-efficient digital quantum simulation of d-level systems for photonic, vibrational, and spin-s hamiltonians, *Npj Quantum Inf.* **6**, 49 (2020).

- [74] B. Koczor, Ph.D. Thesis (2019).
- [75] D. Bacon, I. L. Chuang, and A. W. Harrow, The quantum schur transform: I. efficient qudit circuits, [arXiv:quant-ph/0601001](https://arxiv.org/abs/quant-ph/0601001) (2005).
- [76] D. Bacon, I. L. Chuang, and A. W. Harrow, Efficient Quantum Circuits for Schur and Clebsch-Gordan Transforms, *Phys. Rev. Lett.* **97**, 170502 (2006).
- [77] W. M. Kirby and F. W. Strauch, A practical quantum algorithm for the schur transform, *Quantum Info. Comput.* **18**, 721 (2018).
- [78] H. Krovi, An efficient high dimensional quantum schur transform, *Quantum* **3**, 122 (2019).
- [79] A. Bärttschi and S. Eidenbenz, in *Fundamentals of Computation Theory*, edited by L. A. Gasieniec, J. Jansson, and C. Levcopoulos (Springer International Publishing, Cham, 2019), p. 126.
- [80] G. H. Low and I. L. Chuang, Hamiltonian simulation by qubitization, *Quantum* **3**, 163 (2019).
- [81] G. H. Low and I. L. Chuang, Optimal Hamiltonian Simulation by Quantum Signal Processing, *Phys. Rev. Lett.* **118**, 010501 (2017).
- [82] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, Simulating Hamiltonian Dynamics with a Truncated Taylor Series, *Phys. Rev. Lett.* **114**, 090502 (2015).
- [83] E. Campbell, A random compiler for fast hamiltonian simulation, [arXiv:1811.08017](https://arxiv.org/abs/1811.08017) (2018).
- [84] A. M. Childs, D. Maslov, Y. Nam, N. J. Ross, and Y. Su, Toward the first quantum simulation with quantum speedup, [arXiv:1711.10980](https://arxiv.org/abs/1711.10980) (2017).
- [85] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, A theory of trotter error, [arXiv:1912.08854](https://arxiv.org/abs/1912.08854) (2019).
- [86] A. M. Childs, A. Ostrander, and Y. Su, Faster quantum simulation by randomization, [arXiv:1805.08385](https://arxiv.org/abs/1805.08385) (2018).
- [87] Y. Ouyang, D. R. White, and E. Campbell, Compilation by stochastic hamiltonian sparsification, [arXiv:1910.06255](https://arxiv.org/abs/1910.06255) (2019).
- [88] M. Heyl, P. Hauke, and P. Zoller, Quantum localization bounds trotter errors in digital quantum simulation, *Sci. Adv.* **5**, eaau8342 (2019).
- [89] D. Wecker, B. Bauer, B. K. Clark, M. B. Hastings, and M. Troyer, Gate-count estimates for performing quantum chemistry on small quantum computers, *Phys. Rev. A* **90**, 022305 (2014).
- [90] R. Babbush, J. McClean, D. Wecker, A. Aspuru-Guzik, and N. Wiebe, Chemical basis of trotter-suzuki errors in quantum chemistry simulation, *Phys. Rev. A* **91**, 022311 (2015).
- [91] G. Wendin, Quantum information processing with superconducting circuits: A review, *Rep. Prog. Phys.* **80**, 106001 (2017).
- [92] V. Schäfer, C. Ballance, K. Thirumalai, L. Stephenson, T. Ballance, A. Steane, and D. Lucas, Fast quantum logic gates with trapped-ion qubits, *Nature* **555**, 75 (2018).
- [93] I. D. Kivlichan, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, W. Sun, Z. Jiang, N. Rubin, A. Fowler, A. Aspuru-Guzik, R. Babbush, and H. Neven, Improved fault-tolerant quantum simulation of condensed-phase correlated electrons via trotterization, [arXiv:1902.10673](https://arxiv.org/abs/1902.10673) (2019).
- [94] D. Wang, O. Higgott, and S. Brierley, A generalised variational quantum eigensolver, [arXiv:1802.00171](https://arxiv.org/abs/1802.00171) (2018).
- [95] E. Knill, G. Ortiz, and R. D. Somma, Optimal quantum measurements of expectation values of observables, *Phys. Rev. A* **75**, 012328 (2007).
- [96] The experimental data analyzed in this work is kindly provided by S. Blundell and J. Wilkinson. The data is taken using the pulsed muon beam at the ISIS Facility, Rutherford Appleton Laboratory, UK. Experiments are performed in zero magnetic field, at a temperature of 50 K. Around 181 million muon decay events are included in the dataset.
- [97] T. Jones, A. Brown, I. Bush, and S. C. Benjamin, Quest and high performance simulation of quantum computers, *Sci. Rep.* **9**, 10736 (2019).
- [98] Google, Cirq, <https://github.com/google/cirq> (2020).
- [99] J. R. McClean *et al.*, Openfermion: The electronic structure package for quantum computers, [arXiv:1710.07629](https://arxiv.org/abs/1710.07629) (2017).
- [100] J. P. Gaebler, T. R. Tan, Y. Lin, Y. Wan, R. Bowler, A. C. Keith, S. Glancy, K. Coakley, E. Knill, and D. Leibfried *et al.*, High-Fidelity Universal Gate set for be 9+ ion Qubits, *Phys. Rev. Lett.* **117**, 060505 (2016).
- [101] C. Ballance, T. Harty, N. Linke, and D. Lucas, High-fidelity two-qubit quantum logic gates using trapped calcium-43 ions, [arXiv:1406.5473](https://arxiv.org/abs/1406.5473) (2014).
- [102] M. Huo and Y. Li, Temporally correlated error tomography and mitigation in quantum computer, [arXiv:1811.02734](https://arxiv.org/abs/1811.02734) (2018).
- [103] A. Strikis, D. Qin, Y. Chen, S. C. Benjamin, and Y. Li, Learning-based quantum error mitigation, [arXiv:2005.07601](https://arxiv.org/abs/2005.07601) (2020).
- [104] A. Y. Kitaev, Quantum computations: Algorithms and error correction, *Russian Math. Surveys* **52**, 1191 (1997).
- [105] N. J. Ross and P. Selinger, Optimal ancilla-free clifford+*t* approximation of *z*-rotations, *Quantum Info. Comput.* **16**, 901 (2016).
- [106] S. Pallister, A Jordan-Wigner gadget that reduces *t* count by more than 6x for quantum chemistry applications, [arXiv:2004.05117](https://arxiv.org/abs/2004.05117) (2020).
- [107] D. Litinski, A game of surface codes: Large-scale quantum computing with lattice surgery, *Quantum* **3**, 128 (2019).
- [108] V. von Burg, G. H. Low, T. Häner, D. S. Steiger, M. Reiher, M. Roetteler, and M. Troyer, Quantum computing enhanced computational catalysis, [arXiv:2007.14460](https://arxiv.org/abs/2007.14460) (2020).
- [109] J. Lee, D. Berry, C. Gidney, W. J. Huggins, J. R. McClean, N. Wiebe, and R. Babbush, Even more efficient quantum computations of chemistry through tensor hypercontraction, [arXiv:2011.03494](https://arxiv.org/abs/2011.03494) (2020).
- [110] E. van den Berg and K. Temme, Circuit optimization of hamiltonian simulation by simultaneous diagonalization of pauli clusters, *Quantum* **4**, 322 (2020).
- [111] M. C. Tran, A. Y. Guo, Y. Su, J. R. Garrison, Z. Eldredge, M. Foss-Feig, A. M. Childs, and A. V. Gorshkov, Locality and Digital Quantum Simulation of Power-Law Interactions, *Phys. Rev. X* **9**, 031006 (2019).
- [112] C. Gidney, Halving the cost of quantum addition, *Quantum* **2**, 74 (2018).

- [113] Z. Cai, Resource Estimation for Quantum Variational Simulations of the Hubbard Model, *Phys. Rev. Appl.* **14**, 014059 (2020).
- [114] R. Babbush, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, A. Paler, A. Fowler, and H. Neven, Encoding Electronic Spectra in Quantum Circuits with Linear t Complexity, *Phys. Rev. X* **8**, 041015 (2018).
- [115] D. W. Berry, C. Gidney, M. Motta, J. R. McClean, and R. Babbush, Qubitization of arbitrary basis quantum chemistry by low rank factorization, [arXiv:1902.02134](https://arxiv.org/abs/1902.02134) (2019).
- [116] R. Babbush, J. McClean, C. Gidney, S. Boixo, and H. Neven, Focus beyond quadratic speedups for error-corrected quantum advantage, [arXiv:2011.04149](https://arxiv.org/abs/2011.04149) (2020).
- [117] E. Knill and R. Laflamme, Power of one bit of Quantum Information, *Phys. Rev. Lett.* **81**, 5672 (1998).
- [118] K. Fujii, H. Kobayashi, T. Morimae, H. Nishimura, S. Tamate, and S. Tani, Impossibility of Classically Simulating One-Clean-Qubit Model with Multiplicative Error, *Phys. Rev. Lett.* **120**, 200502 (2018).
- [119] A. Datta and G. Vidal, Role of entanglement and correlations in mixed-state quantum computation, *Phys. Rev. A* **75**, 042310 (2007).
- [120] T. Morimae, K. Fujii, and J. F. Fitzsimons, Hardness of Classically Simulating the One-Clean-Qubit Model, *Phys. Rev. Lett.* **112**, 130502 (2014).
- [121] T. Morimae, Hardness of classically sampling the one-clean-qubit model with constant total variation distance error, *Phys. Rev. A* **96**, 040302 (2017).