

Prospects for NMR spectral prediction on fault-tolerant quantum computers

J. E. Elenewski^{1,*}, C. M. Camara², and A. Kalev^{3,4,†}¹*MIT Lincoln Laboratory, Lexington, Massachusetts 02421, USA*²*Department of Pediatric Oncology and the Linde Program in Cancer Chemical Biology, Dana–Farber Cancer Institute, Boston, Massachusetts, USA*³*Information Sciences Institute, University of Southern California, Arlington, Virginia 22203, USA*⁴*Department of Physics and Astronomy, and Center for Quantum Information Science & Technology, University of Southern California, Los Angeles, California 90089, USA*

(Received 20 June 2025; accepted 4 February 2026; published 24 March 2026)

Advanced atomic magnetometers have made it possible to acquire nuclear magnetic resonance spectra in zero to ultralow magnetic fields. This regime carries the benefit of compact, low-cost instrumentation with reduced spin relaxation effects and the ability to probe phenomena that are inaccessible in conventional high-field experiments. A trade-off is that the resulting spectra must be interpreted using simulations that are taxing for classical computation. Working by example for small-molecule and protein spectroscopy, we demonstrate that these simulations are a promising target for fault-tolerant quantum computation. Our holistic analysis spans from input selection to the construction of explicit circuits for qubitized quantum dynamics. By maintaining parity with experimental requirements, we demonstrate how certain cases might be especially promising for early fault-tolerant architectures. Notably, valuable natural products and small proteins could be simulated with a few hundred logical qubits and less than 10^{12} T gates, corresponding to days of runtime with foreseeable hardware based on spin or superconducting qubits.

DOI: [10.1103/km66-ltpl](https://doi.org/10.1103/km66-ltpl)

I. INTRODUCTION

Nuclear magnetic resonance spectroscopy (NMR) has a distinguished ability to resolve the atomic-scale structures and dynamics of analytes across a hierarchy of timescales [1]. This breadth has ensconced NMR as an analytical gold standard in chemistry, physics, and materials science. While most use cases are related to structure determination, more nuanced studies of reaction dynamics and macromolecular function have had outsized value [2–4]. In particular, the ability to probe near-native environments has delivered critical insights into the folding and interactions of biomolecules such as proteins and nucleic acids [2,5–11]. On a related note, magnetic resonance imaging and spatially resolved NMR spectroscopy have become standard diagnostics in clinical medicine [12–17].

The development of NMR has generally pursued higher spectral resolution over other features. This has led to instruments with large static magnetic fields, B_0 , which give a broad frequency dispersion $\Delta\omega \sim B_0$ among resonances and increase the ratio $P = \exp[-\hbar\gamma B_0/k_B T]$ of polarized spins

[1,5]. While the resulting spectra are less crowded and exhibit better signal-to-noise characteristics, the associated hardware is also costly and has high operational overhead. Nonetheless, advances in atomic magnetometry have recently opened a path toward NMR spectroscopy in zero-to-ultralow fields (ZULF; below 100 nT) where hyperpolarization can deliver spin populations that surpass high-field experiments by several orders of magnitude [18–21] (reviewed in Ref. [22]). This is of technological interest since low-field instruments are comparatively inexpensive, portable, and cryogen free. ZULF experiments can also mitigate field-associated relaxation and can target an insightful regime for spin physics [23–26]. However, only the simplest of these spectra may be interpreted without simulation.

The ultralow-field regime is characterized by long-range vector spin couplings with little separation between energy scales [27]. These features are resource-intensive for classical algorithms but are particularly natural for their quantum counterparts [28]. Attesting to this, several NMR-related tasks have been proposed for noisy intermediate-scale quantum (NISQ) computers [29–32] and the early demonstration of quantum advantage [33,34]. This has isolated highly promising use cases [31,35,36] and provided a solid foundation for future developments. This amenability is perhaps unsurprising, since the calculation of NMR spectra can be directly cast as an exercise in quantum simulation.

In this contribution, we go beyond prior work to consider NMR spectral prediction in the context of fault-tolerant quantum computation (FTQC). Our discussion is framed around experiments in the ZULF regime, though similar

*Contact author: justin.elenewski@ll.mit.edu†Contact author: amirk@isi.edu

considerations hold for strongly coupled spin systems and conventional solid-state NMR experiments at high field. We begin by identifying realistic parameters and targets for computation, including the specification of input molecules across a range of application scales and complexities. Guided by these data, we analyze circuit level implementations for state-of-the-art qubitized algorithms as well as optimized problem encodings that mitigate logical resource overhead. These results are compared with product formula methods to provide broader context. Finally, we estimate the physical requirements for select calculations when using a conventional surface code architecture and a lattice surgery scheme based on sequential Pauli product rotations [37]. We predict that meaningful spectra might be computed with overhead comparable to or less than other guidestar problems, such as factoring 2048-bit integers using Shor's algorithm [38] or simulating classically hard instances of the Fermi-Hubbard model and spin liquid Hamiltonians. This suggests that NMR spectral prediction is a robust computational objective and a potential application for early fault-tolerant quantum hardware [39].

II. SIMULATION PHYSICS

A. Experimental context

The connection between NMR and quantum computation becomes apparent when we consider a simple experiment. In this setting, nuclear spin evolution is initiated using a discrete series of radiofrequency (RF) pulses that are tuned to the Larmor frequency of a target nucleus. This causes the spin system to precess and generates a magnetization profile that is captured by pickup coils in the spectrometer. The resulting signal $\mathcal{S}(t)$ is called the free-induction decay (FID). As the name suggests, this signal is collected as the system equilibrates with its environment. Our interpretable spectrum is a sparse, frequency-domain counterpart to this profile $\mathcal{S}(\omega) = \int_0^{t_{\max}} \mathcal{S}(t) e^{i\omega t} dt$, where $t_{\max}$ is the FID duration. We can simulate this by computing a spin-spin correlator,

$$\langle S_{\text{tot}}(t) S_{\text{tot}}(0) \rangle = \text{tr} [e^{iHt} S_{\text{tot}} e^{-iHt} S_{\text{tot}} \rho], \quad (1)$$

where H is the system's Hamiltonian, ρ is its density matrix, and the operator $S_{\text{tot}} = \sum_k S_k$ captures a total spin. While a typical spectrometer will use phase-sensitive quadrature to detect $S^+ = S^x + iS^y$, this is tied to the experimental apparatus and control pulses. The problem is simpler for computation, where we only need $S_{\text{tot}} = S_{\text{tot}}^z = \sum_k S_k^z$ to capture the location of resonances [40]. This means that the simulated spectrum $\tilde{\mathcal{S}}(\omega)$ will take a simple form:

$$\tilde{\mathcal{S}}(\omega) = \int_0^{t_{\max}} \langle S_{\text{tot}}^z(t) S_{\text{tot}}^z(0) \rangle e^{i\omega t - \gamma_2 t} dt. \quad (2)$$

Note the factor containing $\gamma_2 = T_2^{-1}$, which captures relaxation mechanisms that lead to decoherence and dephasing. This can be handled using a relaxation superoperator in more nuanced treatments.

B. Nuclear spin Hamiltonian

NMR spectroscopy is predicated on the timescale disparity between electronic and nuclear dynamics in a material [1].

This effectively decouples the two subsystems, though the faster electronic component still mediates the nuclear spins through time-averaged parameters. The canonical Hamiltonian for a system of N nuclear spins follows accordingly:

$$\begin{aligned} H = & -\hbar \sum_k \gamma_k (1 - \delta_k) \mathbf{I}_k \cdot \mathbf{B}_0 - \hbar \sum_k \gamma_k \mathbf{I}_k \cdot \mathbf{B}_{\text{RF}}(t) \\ & + \hbar \sum_{k \neq l} J_{kl} \mathbf{I}_k \cdot \mathbf{I}_l \\ & + \sum_{k \neq l} b_{kl} \left[\mathbf{I}_k \cdot \mathbf{I}_l - 3 \frac{1}{|\mathbf{r}_{kl}|^2} (\mathbf{I}_k \cdot \mathbf{r}_{kl})(\mathbf{I}_l \cdot \mathbf{r}_{kl}) \right]. \end{aligned} \quad (3)$$

We adopt a convention where $\mathbf{I}_k = \mathbf{S}_k / \hbar = (S_k^x, S_k^y, S_k^z) / \hbar$ is a vector of spin angular momentum operators for the k th spin. These are defined as $S_k^\alpha = \hbar \sigma^\alpha / 2$ with σ^α a spin-1/2 Pauli operator.

The first term is a Zeeman coupling between the k th atom's nuclear spin and a static, external magnetic field $\mathbf{B}_0$. Both the nuclear gyromagnetic ratio γ_k and the diamagnetic shielding tensor δ_k appear here, reflecting a dependence on the nuclear species and its electronic environment. This confers a distinctive "chemical shift" frequency to each nucleus that vanishes at zero field. The absence of this feature causes different resonances to overlap, complicating the interpretation of ultralow-field spectra [41]. The second term constitutes a similar coupling to the magnetic component of a classical RF pulse. This is useful when simulating specific control pulses but is unnecessary when assigning resonances.

The second and third terms are scalar and dipolar spin-spin couplings, respectively. These differ in terms of prominence and prevalence. The scalar coupling (or J coupling) is an effective hyperfine shift mediated by the electronic environment of chemically bonded atoms. This perturbs the stronger Zeeman contribution at high field, splitting the resonance at a given chemical shift frequency. Conversely, the scalar coupling shifts resonances away from $\omega = 0$ in ultralow-field experiments and gives rise to the main spectral features. Scalar couplings are a sensitive probe of the spin neighborhood and molecular geometry (Fig. 1) [42]. The last term is a dipolar coupling between spatially proximate spins, where $\mathbf{r}_{kl}$ is a vector between the k th and l th nuclei. For convenience, we have absorbed the spatial dependence into the dipolar coupling strength $b_{kl} = \mu_0 \gamma_k \gamma_l \hbar / 4\pi |\mathbf{r}_{kl}|^3$. These contributions average to zero in liquids yet are prominent in the solid state or in weakly aligning media. Both scalar and dipolar couplings appear similar in spectra.

These assumptions reduce the ZULF Hamiltonian to an effective Heisenberg Hamiltonian:

$$H = \sum_{k \neq l} \tilde{J}_{kl} \vec{\sigma}_k \cdot \vec{\sigma}_l + \sum_{\substack{k \neq l \\ \alpha, \beta}} \tilde{D}_{kl}^{\alpha\beta} \sigma_k^\alpha \sigma_l^\beta - \sum_k \tilde{h}_k \cdot \vec{\sigma}_k. \quad (4)$$

This will be the target of our analysis. Here, the rescaled scalar $\tilde{J}_{kl}$ and dipolar coupling components $\tilde{D}_{kl}^{\alpha\beta}$ are specified for a fixed atomic configuration with $\alpha, \beta \in \{x, y, z\}$. We have included the field-dependent Zeeman terms for completeness even though they vanish in our calculations. The structure of this Hamiltonian makes the classical hardness of ZULF simulations apparent [26]. While certain maneuvers can simplify




FIG. 1. Components of a nuclear spin Hamiltonian. Representative chemical shifts and couplings, as illustrated using the organic molecule *N*-methylaniline. Chemically equivalent protons share the same chemical shift, which are labeled for ω_{H1} , ω_{H2} , and ω_{H3} . Electronically mediated J couplings are shown for select homonuclear (J_{HH}^n) and heteronuclear cases (J_{CH}^n , J_{NH}^n). These differ in the number of intermediate bonds n between coupled spins. A dipolar coupling J_{HH}^D is also shown. Dipolar couplings are defined between atoms independent of the J coupling.

high-field calculations—including perturbative strategies and reduced fidelity targets due to relaxation—these are inapplicable at ultralow fields [24,43]. Furthermore, the conjunction of low noise/weak relaxation and long-range interactions is especially problematic for classical tensor network simulations [44,45].

C. Input specification

We generate resource estimates for a large pool of molecular structures. These inputs are drawn from datasets that broadly sample chemical parameter space, allowing us to capture overhead for real-world problems. In order of increasing complexity they include (1) fragment and screening databases for drug discovery, (2) marketed small-molecule pharmaceuticals, (3) biological secondary metabolites or “natural products,” and (4) a series of representative biomolecules (proteins). Molecular structures are mapped to nuclear spin Hamiltonians by identifying plausible couplings between spin-1/2 nuclei. We write J_{AB}^k to denote scalar (J) couplings between nuclei of species A, B that are separated by k bonds and use J_{AB}^D to denote dipolar couplings between the same nuclei [46]. Our large datasets make it infeasible to use experimental coupling parameters for each molecule. Instead, we defer to representative scalar couplings for each chemical environment and calculate dipolar couplings based on molecular geometry [47,48]. Note that electronic structure calculations can be used to estimate spin couplings when reference data are unavailable [49,50].

We consider several schemes based on the coupled nuclei that are detected. The simplest case (denoted Proton) is based on homonuclear scalar couplings J_{HH}^k between protons (^1H) that are separated by up to four bonds. This is common for many high-field NMR spectra and has recently become feasible at ultralow fields [20,51]. We also consider a case with both heteronuclear and homonuclear couplings (Heteronuclear) across ^1H , ^{13}C , and ^{15}N nuclei. This resembles many ZULF experiments to date, which utilize at least one isotopically labeled heteronucleus. We assume that all of the respective atoms in a given molecule will be one of these magnetically active isotopes, though this is only true for a purposefully enriched sample. Thus, our estimates will be an upper bound on simulation complexity for many experiments. Note that these schemes are based on scalar couplings, under which the spin network can fragment into disconnected clusters. The required qubit count will be proportional to the number of spins in the largest cluster, capping spatial overhead and aiding parallel simulation on a large quantum computer. We assume that sufficiently small clusters can be classically outsourced.

Dipolar interactions can become observable in solid-state samples or when using alignment media. When relevant, we accommodate the former (+ Dipolar) by including dipolar terms of strength b_{kl} between $S = 1/2$ nuclei. The case of weak alignment is similar, though we rescale $b_{kl} \mapsto \kappa b_{kl}$ with $\kappa = 1.0 \times 10^{-3}$ to emulate the weaker residual dipolar couplings (+ RDCs) [52]. These are helpful when resolving molecular orientation and have been detected in ZULF experiments [25]. We only include dipolar couplings between nuclei separated by $\|\mathbf{r}_{kl}\| \leq 4.0 \text{ \AA}$ which is reasonable for RDC experiments. While this constitutes a sharp truncation for solid-state experiments, it allows us to compare overhead on an even footing.

D. Simulation parameters

Our aim is to calculate $\langle S_{\text{tot}}^z(t) S_{\text{tot}}^z(0) \rangle$ at a series of timepoints up to a maximal sampling duration t_{max} . This is done with a finite spectral resolution $\Delta\omega$ in the frequency domain that is fixed by the experimental spin-spin relaxation time $\Delta\omega \propto 1/T_2$. Relaxation also sets a minimum fidelity target $F(t)$ each evolution, with fidelity decreasing as t approaches T_2 . In the simplest treatment, a dephasing rate of $\gamma = 1/T_2$ will cause each spin to independently decohere as $F \sim \exp[-\gamma t]$. This has been used to set error thresholds of $\epsilon(t) \leq (1 - \exp[-Nt/T_2])$ for NISQ algorithms, giving a reduction in resource overhead [31]. The utility of this is weaker for qubitized algorithms, where overhead scales logarithmically with precision and a high fidelity might be needed to minimize algorithm repetitions. As a compromise, we set a target of $\epsilon = (1 - \exp[-t/T_2])$ for the target time evolution unitary but always require that $\epsilon \leq \epsilon_{\text{max}} = 5 \times 10^{-3}$. This limited fidelity is matched to the experimental T_2 , as described below, and it sets a baseline for other sources of error. Note that tighter thresholds are required for timepoints $t \ll T_2$ or for more slowly relaxing systems. Again, these only incur $O(\log 1/\epsilon)$ temporal overhead with quantum signal processing (QSP).

The maximal simulation time $t_{\max}$ and sampling schedule are also critical parameters since they define high- and low-frequency cutoffs for $\mathcal{S}(\omega)$, respectively. With regard to $t_{\max}$, experiments generally collect beyond $3T_2$ to minimize artifacts and often exceed $t_{\max} = 5T_2$ for a better signal-to-noise ratio [53]. However, fault-tolerant simulations are relaxation free so our only concern is capturing experimentally detectable dynamics. We do not aim to rival experimental data but instead assign resonances when classical simulations are costly. To this end, we assume that most coherence pathways have developed by $t_{\max} = T_2$, with the remainder falling below the detection threshold. A generic T_2 can be tricky to specify since ZULF relaxation can be slower than the high-field setting by an order of magnitude or more. We adopt a representative $T_2 = 1$ s as a compromise across molecular scales and set $t_{\max} = T_2$, though longer relaxation times may apply in practice.

III. QUANTUM ALGORITHMS

Qubitized methods, such as QSP and its descendants, are among the most resource-efficient algorithms for fault-tolerant quantum simulation [54–57]. QSP itself can approximate $\mathcal{O}(t)$ a time-evolution unitary $\mathcal{U}(t)$ up to a precision $\|\mathcal{O}(t) - \mathcal{U}(t)\|_2 \leq \epsilon$ using $O(t + \log(1/\epsilon))$ queries to the block encoding of a time-independent Hamiltonian. Spatial overhead is also mild since the encoding only requires $O(\log N)$ ancilla. We briefly review these algorithms to frame our analysis.

A. Block encoding strategy

A prerequisite for qubitized dynamics is the ability to represent our nonunitary Hamiltonian H in a quantum circuit. This is accomplished by block encoding H in a larger unitary operator U_H ,

$$U_H = \begin{pmatrix} H & * \\ * & * \end{pmatrix}, \quad (5)$$

which acts on a dilated Hilbert space. By construction, we access H using an ancillary selection state $|G\rangle \equiv |0^m\rangle$ so that

$$H = (\langle G| \otimes I_n) U_H (|G\rangle \otimes I_n). \quad (6)$$

We are only concerned with the block that encodes H and can leave the rest undefined. The term qubitization alludes to this structure, where access for H resembles a two-level system in terms of $|G\rangle$ and its orthogonal complement. Note that our embedding must be a contraction, e.g., $\|H\|_2 \leq 1$, in order to maintain unitarity. This typically requires the Hamiltonian to be rescaled $H \mapsto H/\alpha$ and can have a nontrivial impact on algorithm overhead.

Our nuclear spin Hamiltonian is characterized by Heisenberg couplings with a complicated graph topology. We must map the associated spin network onto a one-dimensional register of N system qubits. Fortunately, this reduces to simple indexing for $S = 1/2$. Each term in $H = \sum_{i=1}^M c_i \Lambda_i$ corresponds to a Pauli string with two nonidentity factors,

$$\Lambda_i = I^{\otimes p} \otimes P_k \otimes I^{\otimes q} \otimes P_l \otimes I^{\otimes (N-p-q-2)}. \quad (7)$$

These terms are unitary and thus H is manifestly a linear combination of unitary operators (LCU). This allows us to block encode using a conventional strategy.

To construct this encoding, we define a *Select oracle* that acts on the system and selection registers,

$$U_{\text{sel}} = \sum_i |i\rangle\langle i| \otimes \Lambda_i, \quad (8)$$

giving access to each term Λ_i in the LCU. A binary encoding $|i\rangle$ is used to flag Λ_i , meaning that $|G\rangle$ must be an $m = \lceil \log_2 M \rceil$ qubit state. Similarly, we define a *Prepare oracle*,

$$U_{\text{prep}} = \frac{1}{\sqrt{\|c\|_1}} \sum_i \sqrt{c_i} |i\rangle\langle G|, \quad (9)$$

which translates our selection state $|G\rangle$ into the weighted superposition required by the Select oracle. Here, $\|c\|_1 = \sum_i |c_i|$ is the 1-norm of the coefficient set $\{c_i\}$. This serves as the normalization factor $\alpha = \|c\|_1$ for our encoding. The block-encoded Hamiltonian is obtained by assembling these pieces:

$$U_H = (U_{\text{prep}}^\dagger \otimes I_n) U_{\text{sel}} (U_{\text{prep}} \otimes I_n). \quad (10)$$

B. Quantum eigenvalue transform

We generate our dynamics using a variant of QSP called the quantum eigenvalue transform (QET). The QET sequence $\mathcal{O}_{\tilde{f}}$ applies a polynomial function across the spectrum of a normal operator,

$$\mathcal{O}_{\tilde{f}} : H \mapsto f(H) = \sum_{\lambda} f(\lambda) |\lambda\rangle\langle\lambda|, \quad (11)$$

which is defined by using a series of classically optimized phase angles $\{\phi_i\}$ to specify f [54–57]. For the purposes of quantum simulation, we would like to implement $f(H) = \exp[-iHt]$. A practical QET implementation requires the block encoding U_H , a qubitized reflection $Z_{\Pi} = 2|G\rangle\langle G| \otimes I_n - I_m \otimes I_n$ on the selection subspace, and a series of qubitized rotations $Z_{\Pi, \phi} = \exp[i\phi Z_{\Pi}]$. This transform is assembled as a product,

$$\mathcal{O}_{\tilde{f}} = e^{i\phi_0 Z_{\Pi}} \prod_{k=1}^d [U_H Z_{\Pi} e^{i\phi_k Z_{\Pi}}]. \quad (12)$$

Note that the isolated reflection is often incorporated into U_H or handled through the phase convention.

While this form of the QET has favorable asymptotic scaling, the use of $U(1)$ rotations can limit expressivity. The leading obstacle is that we can only encode real or imaginary functions of definite parity. This means an evolution operator $x \mapsto \exp(-ixt) = \cos(xt) - i \sin(xt)$ must combine separate circuits for the sine and cosine as an LCU (e.g., following the symmetric convention of Ref. [58]). Other restrictions further limit the transform to $H \mapsto \exp(-iHt/2)$. If the QET sequence gives an error of ϵ in the target function, the probability of a successful shot will be $p = \|\exp(-iHt/2) \pm \epsilon\|^2$. This is close to 1/4 for small ϵ and thus algorithm repetition is inevitable. Robust oblivious amplitude amplification can increase the success probability to near unity [59,60] though this triples the algorithmic depth.

Generalized quantum signal processing (GQSP) [61] is an alternative that follows from early QSP proposals [54,55,62]. Here, the $U(1)$ signal processing rotations are swapped for $SU(2)$ operations,

$$R(\lambda, \phi, \theta) = \begin{pmatrix} e^{i(\lambda+\phi)} \cos(\theta) & e^{i\phi} \sin(\theta) \\ e^{i\lambda} \sin(\theta) & -\cos(\theta) \end{pmatrix}. \quad (13)$$

This expanded set of phases allows us to transform a block encoded operator H as

$$\left[\prod_{k=1}^d R(0, \phi_k, \theta_k) A \right] R(\lambda_0, \phi_0, \theta_0) = \begin{pmatrix} P(\mathcal{W}) & \cdot \\ Q(\mathcal{W}) & \cdot \end{pmatrix}, \quad (14)$$

where $A = |0\rangle\langle 0| \otimes \mathcal{W} + |1\rangle\langle 1| \otimes I$ is defined in terms of a walk operator $\mathcal{W} = Z_{\Pi} U_H$. The result is a pair of complex polynomials $|P(\mathcal{W})|^2 + |Q(\mathcal{W})|^2 = 1$ with fewer restrictions than conventional QSP. The GQSP formalism also specifies a simple and efficient algorithm to calculate the required phase rotations and the form of one polynomial in terms of the other. These benefits compound for Hamiltonian simulation and the eigenvalue transform. Here, the spectral representation of $\mathcal{W} = \bigoplus_{\lambda} \mathcal{W}_{\lambda}$ has a convenient form,

$$\mathcal{W} = \bigoplus_{\lambda} \begin{pmatrix} \lambda & \sqrt{1-\lambda^2} \\ -\sqrt{1-\lambda^2} & \lambda \end{pmatrix} \otimes |\lambda\rangle\langle \lambda|, \quad (15)$$

with eigenvalues of $\exp(\pm i \cos^{-1} \lambda)$ in each qubitized eigenspace. While this holds for any QET variant, the generalized scheme also gives an efficient means to implement $P(\mathcal{W}) = \exp(-iHt)$ as

$$P(\mathcal{W}_{\lambda}) = \begin{pmatrix} e^{-it \cos(\cos^{-1}(\lambda))} & 0 \\ 0 & e^{-it \cos(-\cos^{-1}(\lambda))} \end{pmatrix}. \quad (16)$$

This is not possible with conventional schemes since $P(x)$ is a complex polynomial of indefinite parity. An important consequence is that separate QET transformations are no longer required for sine and cosine components. This eliminates the associated LCU and the need for amplification to boost the success probability. We adopt this generalization of QET for our time evolution circuits [63].

C. Phase angle generation

The GQSP algorithm requires a set of phase angles that approximate $\mathcal{U}(t) = \exp(-iHt)$ to an error of ϵ . These are formally defined using a degree d Chebyshev polynomial from the Jacobi-Anger expansion [54],

$$e^{it \cos(x)} \approx \sum_{n=-d}^d i^n J_n(t) e^{inx}, \quad (17)$$

where $J_n(t)$ is the n th Bessel function of the first kind. Both the complimentary polynomial $Q(x)$ and the $SU(2)$ phase angles $(\lambda_k, \phi_k, \theta_k)$ can be refined using the algorithms of Refs. [61,64]. The result is a series of $2d + 1$ phase tuples and overhead comparable to a standard QET sequence (e.g., via Ref. [58]). In particular, the full circuit will require $2d + 1$ $SU(2)$ rotations, $2d$ qubitized reflections, and $2d$ applications of our block encoding. Moreover, a Jacobi-Anger expansion to degree $d = \lceil |\tau|/2 + \log_{10}(1/\epsilon) \rceil$ gives a truncation error

bounded by ϵ . The cited phase determination algorithms can refine up to 10^6 phases in roughly one minute using readily accessible classical hardware and stated error thresholds. This assures us that classical computations will not be a bottleneck since our small-molecule and protein estimates require less than 1.0×10^7 and 6.0×10^7 phases, respectively.

D. Circuit implementation

We use a straightforward implementation of GQSP that interleaves $SU(2)$ rotations with controlled applications of the walk operator $\mathcal{W} = Z_{\Pi} U_H$. Moreover, we alternate $\mathcal{W}$ and $\mathcal{W}^\dagger$ between repetitions as required when encoding an expansion with terms of negative degree [61]. Our arrangement is depicted in Fig. 2(a). The walk operator itself is constructed from the block encoding U_H and a multicontrolled CZ gate that implements the qubitized reflection [Fig. 2(b)]. Viewed differently, this leads to a multicontrolled $SU(2)$ rotation between repetitions of U_H or $U_H^\dagger$ similar to the qubitized rotation from conventional QET.

The block encoding U_H implements H as an LCU of Pauli strings using Select and Prepare oracles [Fig. 2(c)]. We adopt a Prepare that combines quantum read-only memory (QROM) lookup with an alias sampling strategy and which tolerates dirty ancilla outside of the selection register [65]. Alias sampling approximates LCU coefficients to a fixed bit precision using a space-time trade-off that exchanges T complexity for logical qubit overhead. This approach is particularly reasonable when inputs have a known experimental uncertainty. Our Select is more conventional, applying operators based on binary iteration over the selection register. Circuits are compiled to the universal Clifford+ T -gate set before quantifying resources, as required for quantum error correction schemes like the surface code. All explicit circuit manipulations are accomplished using the pyLIQTR software suite [66] and its Qualtran extensions [67,68]. Our NMR workflow has also been made available as part of pyLIQTR.

E. Observable estimation and sampling

The proposed simulations must begin with an initial spin state that is polarized along the z axis. For high-field spectra this is just the thermal distribution in an external field $\mathbf{B}_0$. However, ZULF is slightly more complex since it relies on nonequilibrium populations that depend on the hyperpolarization procedure [69]. High-precision calculations would aim to mimic these experimental conditions. While the required circuits are not onerous, we maintain brevity outlining a minimal approach.

In the simplest case, we would run simulations for an ensemble of S_{tot}^z eigenstates $|\psi_{0,k}\rangle$ which reflect a net positive magnetization (and conveniently coincide with computational basis states). These can be combined to yield

$$\mathcal{S}(t) = \sum_k \langle \psi_{0,k} | \mathcal{U}^{-1}(t) S_{\text{tot}}^z \mathcal{U}(t) S_{\text{tot}}^z | \psi_{0,k} \rangle, \quad (18)$$

where $\mathcal{U}(t) = \exp[-iHt]$ is the unitary time evolution operator approximated by QET. Although the number of possible states grows exponentially in N , the variance in this estimator is much smaller. Straightforward arguments show that $\mathcal{S}(t)$




FIG. 2. Circuits for GQSP. (a) A circuit implementation of GQSP in the guise of a generalized quantum eigenvalue transform $\mathcal{O}_{\{\lambda, \phi, \theta\}}(H)$. The signal processing rotations $R(\lambda, \phi, \theta)$ are operations in $SU(2)$, which we decompose according to a YZZ convention. (b) Structure of the walk operator $\mathcal{W}$ in terms of the block encoding U_H and multicontrolled CZ gate. (c) Assembly of the block encoding U_H using Select and Prepare oracles. We use this ensemble of circuit components to implement $H \mapsto \exp[-iHt]$.

can be well reproduced by sampling at most N^2 computational basis states [35]. Achieving a measurement precision of ϵ_{meas} requires $O(1/\epsilon_{\text{meas}}^2)$ samples, so we presume a worst case of $O(N^2/\epsilon_{\text{meas}}^2)$ shots. An alternative is to prepare the initial states in superposition on the quantum computer. This can be done by generating a uniform superposition and filtering with inequality tests, as in Ref. [65], to yield a target state with net positive magnetization (though, as noted earlier, more nuanced distributions will be preferred in practice). The correlation function $\mathcal{S}(t) = \langle S_{\text{tot}}^z(t) S_{\text{tot}}^z(0) \rangle$ is then extracted using phase kickback from a pair of total spin operators followed by amplitude estimation (Fig. 3). A rudimentary amplitude estimation strategy might use the Hadamard test with a $O(1/\epsilon_{\text{meas}}^2)$ overhead in the number of shots that estimate $\mathcal{S}(t)$. A robust error threshold would be on the order of $\epsilon_{\text{meas}} = 0.01$




FIG. 3. Correlation function estimator. Circuit estimating the correlation function $\langle S_{\text{tot}}^z(t) S_{\text{tot}}^z(0) \rangle$ based on phase kickback from total spin operators S_{tot}^z to the single-qubit Readout register. Time evolution is facilitated by the QET sequence $\mathcal{O}_H$, while the PUP prepare oracle is used to prepare a superposition over spin-up states. The Hadamard transform $H^{\otimes m}$ does the same over all computational basis states, and serves as a prepare oracle for the binary encoding used by the total spin operator oracles. The n -qubit state of interest is contained in the Target register, while the Hamiltonian is prepared using the $m = \lceil \log_2 M \rceil$ selection register for the LCU. An $s = \lceil \log_2 n \rceil$ register evaluates a sum of Z Pauli operations over each qubit in the target state. The total spin operator is constructed explicitly in Appendix D.

for a total of $N_{\text{shots}} = 10^4$ shots, while a more permissive value of $\epsilon_{\text{meas}} = 0.05$ gives a markedly reduced $N_{\text{shots}} = 400$ shots. Iterative phase or amplitude estimation procedures are an alternative, though they require repetition of the overall circuit [70]. A nuanced discussion can be found in Ref. [71].

Another consideration involves the set of timepoints used to reconstruct $\mathcal{S}(t)$ and thus the target spectrum $\mathcal{S}(\omega)$. A naïve approach would uniformly sample $\langle S_{\text{tot}}^z(t) S_{\text{tot}}^z(0) \rangle$ with a timestep Δt so that $N_{\text{points}} = t_{\text{max}}/\Delta t$ is much greater than N . Based on reference experiments, at least 2^{12} points might be required for a high-resolution spectrum. However, the sparse nature of $\mathcal{S}(\omega)$ is a prime candidate for compressed sensing [72], which can effectively reduce timepoint samples in NMR experiments [73–75]. This improvement ranges from a factor of two for high-precision spectra [76] to a factor of 40 for crude spectra from NISQ algorithms [31]. Again, our goal is to reproduce the location of resonances and not to mimic high-accuracy data. We will assume that $N_{\text{points}} = 400$ is suitable for most purposes.

Putting this together, we must execute $N_{\text{shots}} \times N_{\text{points}}$ time evolution circuits to predict a spectrum. This ranges between 4×10^4 and 4×10^6 shots depending on the target application. However, these points correspond to a range of simulation times. Many of these will be far shorter than the maximal simulation time, with a commensurately lower resource overhead.

IV. RESOURCE ESTIMATES

Resource estimates often use a small set of inputs and assume that they generalize to a larger problem class. This strategy is typically viable, though it can fail when there are large variations in problem structure between instances. To avoid biasing our estimates, we have generated explicit time-evolution circuits for a diverse set of roughly 1.2×10^4 small organic molecules and 211 biological macromolecules [77].

We quantify temporal overhead using the logical T -gate count (N_T) for the longest time-evolution circuit $t = t_{\text{max}}$ of the largest spin cluster used to calculate $\langle S_{\text{tot}}^z(t) S_{\text{tot}}^z(0) \rangle$. We




FIG. 4. Limiting overhead vs problem scales for small-molecule NMR simulations. Single-shot T gate count (N_T) for simulating the largest spin cluster of a molecule up to $t_{\max} = 1$ s with nuclear spin networks of size N . Datasets correspond to Hamiltonians with J couplings between magnetically active protons (Proton; ^1H) and (b) a broader set of nuclei (Heteronuclear; ^1H , ^{13}C , ^{15}N). Panels in each row correspond to these cases or the addition of residual (+ RDCs) or full (+ Dipolar) dipolar couplings. A maximal error of $\epsilon = 5 \times 10^{-3}$ is budgeted when approximating the evolution unitary with GQSP, assuming a spin-spin relaxation time of $T_2 = 1$ s. Contours are kernel density estimates for 1.2×10^4 different molecules, using 15 isovalues thresholded to enclose all points. The horizontal line is Gidney and Ekerå’s 2021 estimate of 1.08×10^{10} T -gates required to factor a 2048-bit integer in 8 h using Shor’s algorithm [38]. Vertical lines are thresholds for classical hardness based on systems with $N = 20$ and $N = 32$ spins.

also use the number of nuclear spins N as a first-order metric for problem complexity. Spatial overhead is quantified through the number of logical qubits N_L in the aforementioned circuit. This includes the N qubits that represent system as well as ancilla for the block encoding, dynamics, and observable extraction.

Qubit count is generically a poor proxy for overall computational complexity. For instance, an identity circuit can be applied to an arbitrary number of qubits with high spatial overhead but low temporal complexity. However, our nuclear spin networks have universal features that increase complexity in concert with N . These estimates are contextualized by specifying hardness scales at $N = 20$ and $N = 32$ spins, which are justified further in Appendix C. In particular, naïve classical algorithms will find explicit application of the time-evolution operator cumbersome at $N = 20$ and could be aided by a fast quantum computer [78]. Similarly, entanglement growth can make accurate matrix product state (MPS) methods with long-range interactions costly at $N = 32$. It is important to note that our goal is not quantum advantage but instead practical utility. That is, a fast quantum computer that gives meaningful speedup will be of value even if the problem is classically tractable.

A. Small-molecule spectra

Magnetically silent nuclei can disrupt scalar coupling networks to give disconnected spin clusters. The largest cluster defines a limiting simulation complexity, assuming that features of the ensemble are roughly scale invariant. Motivated by this, we extracted the largest clusters under each coupling Hamiltonian for all small-molecule instances. The single-shot logical overhead for these is presented in Fig. 4. Here, single-shot denotes a single algorithm execution for a single timepoint. Remarkably, the overhead for simulating these up to $t = t_{\max}$ requires less than 10^{10} T -gates in the J -coupled Proton or Heteronuclear cases with scalar couplings. This is also true when residual dipolar coupling (+RDCs) terms are added to give larger and more connected spin clusters. The cited bound is notable since it demarcates the previously estimated logical overhead for factoring 2048-bit integers in 8 h with Shor’s algorithm and a rotated surface code [38]. Quantum algorithms are provably advantageous and potentially interesting in this context, making the factoring result a “standard candle” for utility and computational scale. Hamiltonians with full dipolar couplings can surpass this overhead (+Dipolar), but only by a maximal factor of 10^2 . While these have almost the same number of terms as their +RDC




FIG. 5. Nuclear spin cluster statistics. Yellow bars denote the fraction of our small-molecule dataset containing a given number of disconnected nuclear spin clusters. The coupling regime is labeled in bold. Blue, green, and brown bars are fractions of the dataset with $N \geq 16$, $N \geq 20$, and $N \geq 32$ spins in the largest cluster, respectively.

counterparts, the larger coupling strengths markedly increase the normalization α and thus the length of the QET sequence. All small-molecule instances require less than 300 logical qubits, which is below the 6190 logical qubits needed for factoring.

While these estimates are suggestive, it is important to delineate problems that will truly benefit from quantum computation. As a general rule, large, highly connected clusters can be challenging for classical algorithms, particularly when they contain numerous couplings of comparable magnitude. However, there is a regime where these remain tractable using Liouville-von Neumann simulators [48], tensor networks, or other methods for highly correlated spin Hamiltonians [26,79]. We capture this feature for our nuclear spin clusters in Fig. 5.

Our simplest spin networks are defined by scalar couplings between protons (Proton) that are separated by up to four bonds. This leads to small, isolated spin clusters that are largely amenable to classical computation. Dipolar couplings modify this by linking spatially proximate protons to give higher connectivity. This is apparent in the + RDCs panels of Fig. 5, which are characterized by fewer and larger clusters in a greater proportion beyond classical hardness thresholds. Heteronuclear J couplings have a similar effect since they deliver additional coherence pathways between bonded atoms. Dipolar couplings compound this connectivity to give a single, molecule-spanning spin cluster. In this limit, over 90% of our instances should be challenging for at least one classical algorithm.

Cluster size can be proportional to classical hardness, but this it is not guaranteed. The exponential scaling of




FIG. 6. Properties of nuclear spin Hamiltonians based on heteronuclear J couplings. Distributions are presented for the number of molecules with given a nuclear spin count N and (a) number coupling terms N_c in the Hamiltonian or (b) mean distance between coupled qubits in a one-dimensional register. The single-shot T -gate count for simulations up to $t = t_{\max}$ is also presented in panel (c). Data correspond to molecules with $N \geq 20$ spins in the Heteronuclear coupling regime, which includes ^1H , ^{13}C , and ^{15}N nuclei. Resource estimates are from Fig. 4.

Hilbert space dimension with spin count will limit the naïve classical algorithms, though sparse simulators and tensor network methods can surpass full-rank simulations for many problems. To further constrain quantum utility, we consider additional metrics including number N_c of spin-spin coupling terms in a Hamiltonian and the distance q_{ij} between the coupled qubits in a one-dimensional register. These capture the sparsity of matrix Hamiltonians and the potential to generate long-range entanglement. The latter can be especially prohibitive for 1D tensor networks like MPS, where the numerical overhead scales as $O(e^S)$ in terms of the bipartite entanglement entropy S between constituent tensors. While a judicious state ordering can mitigate this in some cases [80–83], it is unlikely for spin networks with dipolar couplings.

In Fig. 6, we present a distribution of these quantities for molecules with heteronuclear and homonuclear scalar couplings. These spin networks constitute a single-molecule spanning cluster. The typical instance has many couplings per spin [Fig. 6(a)] and an ensemble-averaged coupling degree of $\langle\langle N_c/N \rangle\rangle = 9.0$. This is accompanied by a mean distance of $\langle\langle d_{ij} \rangle\rangle = 12.6$ between qubits representing coupled spins in




FIG. 7. Aggregate single-shot overhead for small-molecule NMR simulations. Estimates of the aggregate T -gate count required to perform a single-shot dynamics simulation for all quantum-relevant clusters ($N \geq 20$ spins) over $N_{\text{step}} = 400$ logarithmically spaced timepoints up to $t = t_{\text{max}}$. Parameters and precision targets match Fig. 4. Datasets contain 156 molecules (Heteronuclear), 1806 molecules (+ RDCs), and 1806 molecules (+ Dipolar), respectively. These are smaller than the full dataset since some molecules lack clusters beyond the $N \geq 20$ threshold.

our one-dimensional register [Fig. 6(b)], and an mean maximum distance of $\langle \langle \max_{ij} d_{ij} \rangle \rangle = 30.5$ qubits. While MPS evolution algorithms can treat long-range couplings up to $N \sim 50$ spins with some efficiency, these states are often ordered so the coupling strength decays algebraically with distance. That is not the case for our nuclear spin Hamiltonians, where the numerous long-range couplings are often comparable in strength to their local counterparts (or up to 10^3 -fold larger for full dipolar couplings in the solid state). These parameters lie near the mode of the distribution, as 43% of molecules have spin networks that exceed the mean node degree and 39% exceed the mean distance. Based on this, it is likely that a substantial fraction of molecules would see a benefit from quantum computation when some dipolar couplings are present. It is significant that these parameters generally increase with the spin count N and are positively correlated with a greater T -gate count [Fig. 6(c)]. Thus, classical and quantum hardness are expected to scale in concert for this problem.

So far, we have only considered the overhead for a single quantum dynamics simulation. However, a major difference between number theoretic algorithms and the estimation of a correlation function is the need to consider a multitude of timepoints. We must also simulate all classically hard spin clusters in a given molecule. The latter will not increase the maximal number of required qubits since the estimates of Fig. 4 already correspond to the largest spin cluster. In fact, nearly all instances from the Heteronuclear + RDCs case comprise a single cluster that spans across the entire molecule [84]. The single-shot T -gate count for the leading cluster is also bounded since these estimates correspond to the longest simulation time t_{max} (though this does not have to bound the maximal T -gate count for individual clusters).

A useful metric is the aggregate overhead for a single shot of all calculations (clusters and timepoints) used in estimating a nuclear spin correlation function. That is, we estimate T -gate counts for running N_{step} simulations of duration $t_k \in [0, t_{\text{max}}]$ and aggregate these over a set of M relevant clusters: $N_{T,\text{agg.}} = \sum_{k=1}^{N_{\text{step}}} \sum_{m=1}^M N_T(m, t_k)$. Since we can use a compressed sensing protocol, it is likely that the time samples will be weighted toward early points. We approximate this by using a logarithmic discretization of the time domain so that

$t_k \in [1/(2f_{\text{max}}), t_{\text{max}}]$. The frequency f_{max} is an upper bound on the fastest coupling in the spin network and the factor of two is added to satisfy Nyquist requirements. Since small clusters can be treated classically, we only include those with $N \geq 20$ in our estimates. This is beyond the scale of exact simulation, though we presume that some approximate method (e.g., tensor networks or otherwise) will have classical utility near or below this scale. Again, a sufficiently fast quantum computer might have utility even when a cluster is classically tractable. These assumptions lead to the estimates of Fig. 7. On the single-shot level, we see that the aggregate T -gate count exceeds our longest simulation of the largest cluster in each molecule by roughly a factor of 10^2 . The number of required shots N_{shots} will depend on the target measurement precision ϵ_{meas} and protocol as discussed earlier.

B. Biological macromolecules

Our estimates suggest that quantum computation might have sub-Shor utility for simulating the nuclear spin networks of certain small organic molecules. Macromolecular spin networks are even more taxing for classical computation and they have greater spectroscopic utility [2–4,6–11]. To assess this case, we have examined a series of proteins within the Proton and Heteronuclear coupling regimes with optional residual dipolar couplings (+ RDC). This macromolecular dataset includes drug-discovery relevant peptides as well as larger protein structures from a previously designed benchmark [85]. Here, we find a tight correlation between problem qubit N - and T -gate counts (Fig. 8). This reflects the compact spatial conformations of folded proteins and the fact that these biopolymers assembled from a fixed pool of amino acid blocks. Interestingly, all molecules in this dataset still lie below the logical qubit count for 2048-bit integer factoring. The largest proteins that we consider are ion channels—the P7 channel and the human voltage-dependent anion channel (VDAC)—and a maltodextrin binding protein. All three are challenging for NMR experiments and have no feasible path toward classical computation in the ZULF regime.

We contextualize this overhead by comparing to key quantum dynamics problems, including a J_1 – J_2 Heisenberg




FIG. 8. Overhead estimates for simulating biomolecular (proton) spin networks. Single-shot T -gate (N_T) and logical qubit (N_L) counts for simulating 211 different protein spin networks up to $t_{\max} = 1$ s. Hamiltonians are defined for the largest spin cluster in each regime, which is generally a single cluster than spans the entire protein. Parameters match those of Fig. 4 and estimates are presented as a kernel density at three contour levels. Reference estimates are also shown for the single-orbital Fermi-Hubbard model on a square lattice (green diamonds) and a $J_1 - J_2$ Heisenberg antiferromagnet on a triangular lattice (yellow hexagons). Lattice Hamiltonians and parameters are described in Appendix C.

antiferromagnet on a triangular lattice and the single-orbital Fermi-Hubbard model on a square lattice (see Appendix C for details). These are effective models for spin liquids and strongly correlated electronic materials and large instances of both remain intractable to classical computation for large instances. More specifically, 16×16 site lattices lie near the envelope of current classical computations, though unambiguously valuable information should become apparent at 128×128 lattices and beyond. Our estimates suggest that simulating protein NMR spectra has temporal complexity comparable to reproducing correlation functions for these condensed matter systems. Moreover, this can be done with fewer logical qubits.

Our analysis has assumed that molecules are at most weakly aligned. Realistic solid-state NMR predictions would




FIG. 9. Categorical distribution of single-shot overhead for small-molecule NMR simulations. Estimates of the aggregated T -gate count needed to perform a single-shot dynamics simulation for all quantum-relevant spin clusters ($N \geq 20$ spins) at $N_{\text{step}} = 400$ logarithmically spaced timepoints up to $t = t_{\max}$. Data, parameters, and precision targets match the heteronuclear regime of Fig. 4.

require longer distance cutoffs for dipolar interactions. Moreover, individual dipolar couplings would be roughly 1000-fold larger than the RDC counterparts, substantially increasing the normalization factor α and thus the effective simulation cost. This would give a markedly elevated T -gate count, though the spatial resource utilization would be comparable.

C. Utility and Physical Resource Overhead

ZULF NMR is an emerging technique with potential impacts throughout chemistry and materials science. Our input molecules are drawn from real-world datasets, allowing us to speculate on the value of quantum computation in supporting this spectroscopy. In Fig. 9, we present a categorical classification of our small-molecule estimates with scalar heteronuclear couplings. From a practical standpoint, only larger pharmaceuticals or natural products are very likely to remain inaccessible with foreseeable improvements in classical algorithms. This suggests that the strongest candidates for quantum utility would be spectroscopically challenging small-molecules or macromolecules. The former might uniquely benefit from ZULF spectroscopy since high-field experiments can struggle to resolve an unambiguous structure.

Since the proposed calculations would facilitate experiments, their utility is inherently tied to their runtime. We have estimated this for spectroscopically challenging reference molecules and therapeutic-scale peptides in Fig 10. This analysis assumes a planar layout of spin or superconducting qubits with nearest-neighbor couplings and a quantum error correction layer based on lattice surgery and the rotated surface code (detailed in Appendix E). While simulations with the minimal number of problem qubits would be prohibitively slow, a quantum computer that could factor 2048-bit integers (6190 logical qubits) or simulate a 128×128 site instance of the

Molecule	N	N_L	(d_2, d_1)	N_{fact}	N_{phys}	Minimal Qubits	Gidney–Ekerå	FH 128
Proton + RDCs								
IPC	18	74	(13, 7)	4	1.8×10^5	4294.0 h	39.4 h	16.4 h
Parthenolide	20	76	(13, 7)	4	1.8×10^5	4743.3 h	43.9 h	18.2 h
Strychnine	22	78	(13, 7)	4	1.9×10^5	3300.7 h	30.8 h	12.8 h
Heteronuclear								
IPC	27	106	(15, 9)	4	2.9×10^5	32286.0 h	467.9 h	193.3 h
Parthenolide	35	114	(15, 9)	4	2.9×10^5	44630.3 h	666.1 h	273.8 h
Strychnine	45	124	(15, 9)	4	3.0×10^5	64655.9 h	994.7 h	409.2 h
Proteins (Proton)								
α -Conotoxin	99	168	(13, 7)	4	2.5×10^5	7205.9 h	91.2 h	37.7 h
Gramicidin	124	193	(13, 7)	4	2.7×10^5	7620.0 h	103.0 h	42.8 h
Mastoparan	126	195	(15, 9)	4	3.7×10^5	18719.8 h	353.2 h	145.1 h

FIG. 10. Physical overhead for spectroscopically challenging molecules. Runtimes reflect the concurrent execution of algorithm repetitions on systems (1) with the minimal physical qubits for simulation, (2) capable of factoring 2048-bit integers in 8 h (Gidney–Ekerå; 6190 logical/20 M physical qubits) and (3) capable of simulating the single-orbital Fermi-Hubbard model a 128×128 lattice (FH 128; 32 805 logical/48.1 M physical qubits). Estimates are based on our explicit logical counts and a fast lattice surgery layout using the rotated surface code. We assume an error budget of $\epsilon_{\text{target}} = 0.1\%$ for the QEC layer, measurement depth $\mathcal{D}_{\text{meas}} = N_T$, cycle time of $t_{\text{cycle}} = 1 \mu\text{s}$, reaction time $t_{\text{react}} = 10 \mu\text{s}$, and an error rate of 0.01% for physical two-qubit gates. T -factory distances (d_2, d_1) and counts N_{fact} are optimized to meet this budget, giving physical qubit counts N_{phys} and runtimes following the methodology of Appendix E. The code distance $d = d_2$ of data tiles is matched to our level-two factories.

Fermi-Hubbard model (32 805 logical qubits) allow sufficient parallelization to be useful. However, a quantum computer could also be advantageous for even smaller molecules if it is sufficiently fast. An interesting takeaway from this analysis is the interplay between the Hamiltonian regime, molecular scale, and physical overhead. Notably, small molecules with heteronuclear couplings can surpass the quantum resource requirements for simulating drug-like peptides with a simple spin Hamiltonian. These systems are likely to have higher utility and value should increase with the size of a protein. We do not press this analysis further since larger molecules will likely to require a greater number of timepoints due to spectral crowding.

We have provided estimates based on the surface code due to its maturity and ubiquity. Other error correction schemes such as quantum low-density parity-check codes (qLDPCs) and yoked surface codes can offer more robust encoding rates and potentially lower resource requirements [86–88]. This has been established in revised estimates for factoring problems [89], though the formalism and architectures for these alternatives are not universally well developed.

D. Alternative quantum algorithms

Qubitized algorithms are nearly optimal in terms of resource overhead. This limits opportunities for competitors, though it is important to differentiate between analytical scaling and empirical cost. For example, the best commutator bounds for product formulas are known to overestimate temporal overhead by 10^2 - to 10^3 -fold when compared to empirical simulations of one-dimensional Heisenberg models [90–92]. This can outperform QSP while requiring fewer logical qubits, albeit in a Hamiltonian-dependent manner. A caveat is that empirical results must be extrapolated from classically tractable instances to computationally hard scales. While this makes it difficult to know which algorithm will be optimal for a given problem, is still worthwhile to estimate how our qubitized workflow might compare.

To approach this question, we have used comparatively tight commutator bounds to estimate the T -gate overhead $N_T^{\text{PF}^2}$ for product formula simulations using a second-order Trotter-Suzuki scheme (see Appendix F). We present the ratio $R_T = N_T^{\text{GQSP}}/N_T^{\text{PF}^2}$ of our GQSP estimates to these counts N_T^{GQSP} in Fig. 11. This is assessed for a series of nuclear spin Hamiltonians as a function of the anticipated simulation error ϵ . Based on these data, qubitized algorithms will outperform product formulas across error scales for all but the largest spin networks. In particular, the $O(n^k \|H\|_1 t + \log(1/\epsilon))$ scaling of GQSP for k -local Hamiltonians is increasingly favorable at small ϵ due to the polynomial overhead of second-order product formulas.

Most of our resource estimates correspond to systems that are impractical or intractable using classical computation. However, approximate tensor network algorithms can simulate some for relatively short timescales. We use this strategy to perform an explicit error analysis that we also present in Fig. 11. These empirical estimates require 10^3 -fold fewer T gates than predicted by commutator bounds, consistent with prior calculations for related spin systems [90–92]. The T -gate count is likewise consistent across similarly sized spin networks as expected from Fig. 6. A crossover favoring GQSP occurs below $\epsilon \sim 10^{-5}$ while product formula outperform by less than an order of magnitude up to $\epsilon \sim 10^{-3}$. Given these ambiguities, we rely on qubitization for our physical resource estimates while acknowledging that they are likely to be pessimistic.

Our analysis does not account for problem-optimized block encodings or the empirical performance of qubitization. These factors will further reduce the T count in favor of GQSP and it might be argued that $\epsilon \sim 10^{-3}$ would be closer to a realistic crossover. Since R_T varies roughly tenfold across the problem scales in Fig. 6, it is reasonable that our resource estimates for $\epsilon = 5 \times 10^{-3}$ lie near the upper limit in ϵ where qubitization might be useful. That being said, this error target is based on a conservatively short coherence time for ultralow fields and it




FIG. 11. Relative logical overhead for product formula simulations. The ratio $R_T = N_T^{\text{GQSP}}/N_T^{\text{PF2}}$ of T -gate overhead for GQSP and second-order Trotter-Suzuki (PF2) simulations is based on analytical bounds (solid) and empirical (dotted line) measures of product formula error. The latter are driven by TEBD simulations at different Trotter step counts r and thus time step $\delta t = t_{\text{max}}/r$. Two-site TDVP calculations provide a reference with $r = 8000$. Error ϵ is quantified between terminal states from TEBD and TDVP evolutions as described in Appendix F. All tensor network simulations adopt maximal bond dimension $\chi = 64$, no singular value cutoff, and a simulation time of 1.0×10^{-3} s which is extrapolated to $t_{\text{max}} = 1.0$ s. Nuclear spin Hamiltonians include heteronuclear J couplings and residual dipolar couplings up to 4 Å at varying spin counts (Cycloserine: 11, Caffeine: 22, IPC: 28, Strychnine: 45, Ciguatoxin: 137, Mastoparan: 126, BPTI: 428, Cytochrome c: 802).

only applies to the longest simulation times. Many systems will have a longer T_2 and will require a smaller value of ϵ to yield accurate spectra. GQSP will be advantageous for these simulations due to its $O(\log(1/\epsilon))$ scaling, though again the optimal choice of method is likely to be system dependent. Early-time correlators are another case with tight error thresholds and these are relevant to all spectral simulations. We should note that the choice of algorithm is not guided by T count alone. Product formulas can excel when spatial resources are limited since they generally do not require ancilla (e.g., for block encoding). This can be valuable in the near-term when logical qubits will be scarce.

V. CONCLUSIONS

In this contribution, we have analyzed how fault-tolerant quantum computation might augment certain low-field NMR spectroscopies. A unique aspect of this work is that it addresses a large number of inputs and provides explicit, circuit-based estimates using state-of-the-art qubitized algorithms (following Ref. [62]). Based on our estimates, several classes of molecules remain at or beyond the limits of classical computation yet feasible for quantum computation. For example, our resource estimates suggest that small-molecule NMR spectra can be simulated with fewer than 10^{10} T gates,

a computational burden comparable to factoring 2048-bit integers using Shor's algorithm (though recent estimates have pushed this to lower resource counts [89]). Macromolecular systems with up to 124 spins in their largest cluster require fewer than 10^5 logical qubits and T -gate counts comparable to those required for condensed-matter models like 128×128 Fermi-Hubbard lattices (when using a qubitized approach) or spin liquid models. This positions ZULF NMR spectral prediction as a high-impact application for early fault-tolerant quantum hardware, with significant implications for spectroscopy-driven research in chemistry, biology, and materials science.

A quantum computer can have utility for classically tractable problems if it is inexpensive and high throughput. This is particularly true for NMR prediction since it does not enjoy the same large, dedicated high-performance computing allocations as many frontier-scale chemistry and physics problems. The stated limitations are not just a matter of resource availability. Since spectral predictions should complement experiments, a task that requires weeks of compute time could have limited utility for many small molecules yet remain relevant for challenging targets like natural products and peptides. Conversely, the ability to simulate spectra for large numbers of molecules could have added value by virtue of sheer computational throughput.

Our effort follows several discussions of NMR simulation in an NISQ context [31,35,36]. Beyond low-field instruments, it has been proposed that quantum computation can find utility predicting macromolecular and solid-state materials spectra for high-field solid-state NMR. This is certainly possible based on our observations. However, the large number of strong, long-range interactions will markedly increase the block encoding normalization and linearly increase the runtime by the same factor. The cost would be even greater in a high-field regime, where the nonzero chemical shift terms ω_{k0} can be a factor of 10^6 larger than other couplings. Thus, the most plausible use case for solid-state spectra remains in the ZULF regime. Another opportunity for quantum simulation involves the design of pulse sequences for spectrometers [5]. This is an exercise in quantum control and the ability to tackle large spin systems could lead to unanticipated experimental techniques. However, both this and solid-state NMR require methods for the simulation of explicitly time-dependent Hamiltonians. These are challenging to implement for qubitization but straightforward in the case of product formulas [93–96].

As a final note, we have steered toward pessimistic estimates for resource overhead. This is particularly true when considering spectra with heteronuclear couplings. However, many experiments are performed for homonuclear proton networks, with select heteronuclei included through isotopic enrichment. This means that practical Hamiltonians would contain fewer terms with a commensurate reduction in T complexity.

ACKNOWLEDGMENTS

This material is based upon work supported in part by the Defense Advanced Research Projects Agency (DARPA) under Air Force Contract No. FA8702-15-D-0001. In addition,

this research was developed in part with funding from the DARPA under Contract No. HR00112330014. C.M.C. also acknowledges support under NIH Grant No. 5T32HL007574.

Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views or policies of the Defense Advanced Research Projects Agency, the Department of Defense, or the U.S. Government.

DATA AVAILABILITY

The data that support the findings of this article are openly available [66]; embargo periods may apply.

APPENDIX A: DATASET CURATION AND INPUT PREPARATION

Small-molecule structures were prepared using classical, molecular mechanics based methods. Initial all-atom geometries were generated from a standard library of bond angles, as defined by the molecular connectivity and resulting orbital hybridization. These preliminary geometries were relaxed using conjugate gradient minimization according to the Merck molecular mechanics force field (MMFF94) [97]. No electrostatic or van der Waals cutoffs were applied, and optimizations were conducted to a relative energy gradient of 1×10^{-8} . This strategy is sufficient for many (nonreactive) biomolecular simulations, where the high environmental temperature often permits a classical description of molecular geometry. Numerical optimizations were handled using the Open Babel toolkit [98]. Small-molecule inputs were drawn from Maybridge libraries (Building Block, RO3 Diversity, and Screening collections), pharmaceuticals from nonbiologic drugs approved by the United States Food and Drug Administration (the FDA “Orange Book”), and natural products were curated from PubChem databases [99–101].

Protein inputs are based on solution NMR structures, as drawn from the Protein Data Bank (PDB) [102] and established reference sets [85]. No further optimizations were applied since these geometries are realistic targets for the proposed quantum algorithm workflow.

APPENDIX B: CLASSICAL COMPUTATION AND QUANTUM UTILITY THRESHOLDS

Quantum utility is a moving target that is defined, in part, by the frontier of classical computation. In this section, we briefly remark on prior work and leading classical algorithms that guide our hardness thresholds.

Classical simulations of ZULF spectra have generally utilized full-rank density matrix evolutions and a judicious choice of basis to achieve a per-time-step complexity of $O(2^{2N})$ in problem size N [28]. These employ near-optimal matrix multiplication algorithms to go beyond a $O(2^{3N})$ scaling but are less efficient than sparse statevector simulators and tensor network methods. To put these methods in perspective, a system of $N = 20$ spins requires 16 384 Gb to store the full time-evolution operator or a density matrix (assuming 64-bit complex entries). While formally tractable, this requires the resources of a small cluster and makes integration up to relevant timescales cumbersome.

Tensor networks—such as MPS—offer a systematic trade-off between accuracy and computational cost [103,104]. MPS factorizations are driven by the singular value decomposition (SVD) and reduce the overhead to store a statevector from $O(2^N)$ to $O(dN\chi^2)$ using iterative truncations that are optimal in the Frobenius norm. A sim for operators). Here, $d = 2$ is the dimension of our computational basis and the bond dimension χ reflects the maximal number of classical parameters (singular values) that are retained between truncations. This ansatz associates a tensor $A_{s_{\alpha-1}s_{\alpha}}^{j_{\alpha}}$ with each physical site $\alpha \in \{1, \dots, N\}$, where states are index by $j_{\alpha} \in \{0, 1\}$ in our $d = 2$ computational basis. The tensors are contracted by summing over indices $1 \leq s_{\alpha} \leq \chi$ to give a compressed statevector:

$$|\Psi\rangle = \sum_{j_1 j_2 \dots j_N} C_{j_1 j_2 \dots j_N} |j_1\rangle \otimes |j_2\rangle \dots |j_N\rangle \\ \approx \sum_{j_1 j_2 \dots j_N, \{s\}} A_{s_1}^{j_1} A_{s_1 s_2}^{j_2} \dots A_{s_{N-1}}^{j_N} |j_1\rangle \otimes |j_2\rangle \dots |j_N\rangle, \quad (\text{B1})$$

where $C_{j_1 j_2 \dots j_N}$ is the unfactorized coefficient tensor.

The main algorithms driving MPS time-evolution are time-evolving block decimation (TEBD) and the time-dependent variational principle (TDVP). TEBD applies a series of two-site gates to the MPS at each timestep following a Trotter-Suzuki expansion as detailed in Appendix G. The cost of each gate is dominated by the SVD, which scales as $O(d^3 \chi^3)$ for adjacent tensors in the basis ordering. If we have M the number of spin coupling terms in our Hamiltonian, then the overall scaling will be $O(Md^3 \chi^3)$ scaling per step. This is much better than the full-rank evolution if a reasonably small χ is sufficiently accurate for a given application. The TDVP algorithm takes a different strategy by constraining an evolution to the tangent space of an MPS manifold. In particular, a series of projections are used to generate a local Hamiltonian for the tangent space and the equations of motion are integrated to define a time-optimal evolution. The two-site variant of TDVP has a comparable time complexity of $O(N(\chi^2 d^3 D^2 + \chi^3 d^2 D + \chi^3 d^3))$, where N is the number of MPS sites and D is a bond dimension designated for local updates [105].

Long-range couplings lead to broadly disseminated entanglement, which can scale dramatically with the effective spatial dimensionality of a system. This also increases the required dimension of bonds interstitial to coupled sites in the MPS. Notably, a half-chain entanglement entropy S is associated with a bond dimension $\chi \sim O(\exp S)$. In the worst case, this means that the simulation cost can grow exponentially between successive TEBD iterations (at fixed accuracy). These factors limit efficient MPS to systems with quasi one-dimensional entanglement structures. Our simulations are precisely the opposite: They have many short- and long-range interactions of comparable magnitude. This gives broadly disseminated entanglement and makes it hard to find a reasonable trade-off between bond dimension (numerical overhead) and acceptable accuracy. Based on these factors, we anticipate that systems with $N = 32$ will be safely cumbersome for long-timescale TEBD or TDVP simulations. While other tensor networks can tolerate more complicated entanglement patterns, these are costly and not well adapted for practical use.

APPENDIX C: CONDENSED-MATTER MODELS

To frame our estimates more broadly, we have calculated the overhead for simulating prototypical condensed-matter models on a fault-tolerant quantum computer. These systems carry intrinsic theoretical value, though the classical overhead of large, high-precision calculations lies beyond the scope of existing numerical methods.

The first model we consider is the J_1 – J_2 Heisenberg Hamiltonian ($S = 1/2$) on a two-dimensional triangular lattice,

$$H_{12} = J_1 \sum_{\langle i,j \rangle} \vec{\sigma}_i \cdot \vec{\sigma}_j + J_2 \sum_{\langle\langle i,j \rangle\rangle} \vec{\sigma}_i \cdot \vec{\sigma}_j, \quad (\text{C1})$$

where $\vec{\sigma}_i = (\sigma_i^x, \sigma_i^y, \sigma_i^z)$ is a vector of Pauli operators acting on the i th lattice site. Repeated brackets denote summation over nearest and next-nearest-neighbor pairs, respectively. This frustrated spin system is believed to exhibit spin-liquid order at intermediate values of J_2/J_1 . Nonetheless, numerical and analytical efforts remain ambiguous regarding the precise nature of this phase [106–108]. The J_1 – J_2 model can be encoded in the same manner as our NMR Hamiltonians since both are defined in terms of Pauli strings.

Our second reference is the Fermi-Hubbard model, which is believed to capture the leading behavior of strongly correlated materials. This problem has garnered significant attention as a target for quantum computation [65,71,109,110]. We consider the single-orbital case on an $N_x \times N_y$ square lattice as specified by the Hamiltonian

$$H_{\text{FH}} = - \sum_{\langle i,j \rangle, \alpha} (J_{ij} c_{i,\alpha}^\dagger c_{j,\alpha} + \text{H.c.}) + U \sum_i n_{i,\downarrow} n_{i,\uparrow}. \quad (\text{C2})$$

Here, $\{c_{i\alpha}, c_{j\beta}^\dagger\} = \delta_{ij}\delta_{\alpha\beta}$ are fermionic creation and annihilation operators for a particle of spin α or β on the

i th or j th lattice site, respectively. By extension, we write $n_{i\alpha} = c_{i\alpha}^\dagger c_{i\alpha}$ for the fermionic particle number operator. This system is parameterized by an intersite hopping strength J_{ij} and a many-body interaction U between particles of different spin on the same lattice site. We consider an isotropic limit where $J_{ij} = J$ is constant and the ratio U/J determines the underlying physics.

Unlike our other targets, the Fermi-Hubbard model is defined in terms of fermions. We can tackle this by using the Jordan-Wigner transform to map spinless fermions to qubits,

$$c_j = (-Z)^{\otimes(j-1)} \otimes \sigma_j^- \otimes I^{\otimes(N-j)}, \quad (\text{C3})$$

$$c_j^\dagger = (-Z)^{\otimes(j-1)} \otimes \sigma_j^+ \otimes I^{\otimes(N-j)}, \quad (\text{C4})$$

$$n_j = c_j^\dagger c_j = (I_j + Z_j)/2, \quad (\text{C5})$$

where $\sigma_j^\pm = (X_j \pm iY_j)/2$ are spin ladder operators. To accommodate our spinful case, we double the number of lattice sites within each row to give a $2N_x \times N_y$ lattice. We then designate even and odd sublattices as spin-down and spin-up orbitals, respectively, so that row-adjacent pairs of different spin orbitals can be mapped to the same physical site. Thus, the hopping terms from H_{FH} are confined to sublattices of the same parity while the many-body interaction occurs between different spins on the same physical site (which lie on different lattices). A mapping between the 2D square lattice and a 1D chain is accomplished by defining a zigzag path between rows, giving a constant Pauli string length for coupled sites in a column. Putting these pieces together, we obtain an amenable representation for our Hamiltonian:

$$\begin{aligned} H_{\text{FH}} = & \frac{J}{2} \sum_{j=1}^{N_y} \sum_{k=1}^{N_x-1} [X_{(j,2k-1)} \bar{Z} X_{(j,2k+1)} + Y_{(j,2k-1)} \bar{Z} Y_{(j,2k+1)} + X_{(j,2k)} \bar{Z} X_{(j,2k+2)} + Y_{(j,2k)} \bar{Z} Y_{(j,2k+2)}] \\ & + \frac{J(-1)^{N_x}}{2} \sum_{j=1}^{N_y-1} \sum_{k=1}^{2N_x} [X_{(j,k)} \bar{Z} X_{(j+1,k)} + Y_{(j,k)} \bar{Z} Y_{(j+1,k)}] + \frac{U}{4} \sum_{j=1}^{N_y} \sum_{k=1}^{N_x} (I + Z_{(j,2k-1)})(I + Z_{(j,2k)}). \end{aligned} \quad (\text{C6})$$

The parenthetical notation $(j, k) = 2N_x(j-1) + k$ denotes the index of the subscripted lattice site in the one-dimensional qubit register. Note that we do not include a chemical potential term since the filling fraction is set by the initial state in this arrangement.

Our resource estimates for the J_1 – J_2 model assume GQSP dynamics generated by Eq. (C1), with a $t_{\text{max}} = 1.0$ s and $J_1 = \text{Hz}$ and $J_2 = 0.5$ Hz. Similarly, we set $U/J = -4.0$ for Fermi-Hubbard model and adopt a maximal simulation time of $t_{\text{max}} = 2\pi \cdot 200.0$. This long simulation would set a lower bound on the energy scale for any correlation functions. Hamiltonians are block encoded as an LCU over Pauli strings using the same alias sampling strategy as NMR Hamiltonians. We specify a precision of 0.1% for the Hamiltonian coefficients and a phase angle sequence that reproduces the target unitary with a fidelity of $\epsilon = 0.001$. It should be noted that more efficient encoding strategies exist for the Fermi-Hubbard

model [65]. However, we follow the outlined approach since it is comparable in structure to our other representations.

APPENDIX D: TOTAL SPIN OPERATOR

We present a straightforward quantum circuit for the total spin operator S_{tot}^α in Fig. 12. This has low temporal overhead and the logical qubit requirements parallel our qubitized simulation algorithm.

APPENDIX E: PHYSICAL RESOURCE ESTIMATION

Our physical resource estimates are based on lattice surgery with the rotated surface code [37,111]. This is natural for planar devices with nearest-neighbor connections, though other hardware layouts can be used with modest overhead. To guide our analysis, we adopt design principles on parity with Ref. [38]. Doing so maintains a common reference point




FIG. 12. Total spin operator. Circuit implementing the the total spin operator S_{tot}^α , where $\alpha \in \{x, y, z, +, -\}$ denotes the constituents. This arrangement requires $s = \lceil \log_2 n \rceil$ ancilla, which must be put into an equal superposition using a Hadamard transformation and uncomputed after calculating the correlation function.

for overhead though the overall implementation is not highly optimized. The parameters for our estimates are summarized in Fig. 13.

1. Fault-tolerant layout

A foundational and potentially rate-limiting requirement is the production of magic states. We accomplish this using AutoCCZ factories, which can deliver $|T\rangle$ -states by generating of an intermediate $|\text{CCZ}\rangle = \text{CCZ}|+\rangle^{\otimes 3}$ and performing a catalytic $|\text{CCZ}\rangle \rightarrow 2|T\rangle$ conversion [112,113]. In brief, the entry point is an injection of T states at code distance $d_0 = \lfloor d_1/2 \rfloor$ using well-established methods [114]. These states are subsequently passed through a round of 15:1 Reed-Muller distillation in “level-1” T -subfactories [113]. The result is a

	Parameter	Symbol	Value
Level 0 (Injection)	Volume	$\mathcal{V}_{T_1}$	100 cells
Level 1 (T_1 -Factory)	Area	$\mathcal{A}_{T_1}$	$4d_1 \times 8d_1$
	Depth	$\mathcal{D}_{T_1}$	$5.75 d_1$
	Volume	$\mathcal{V}_{T_1}$	1100 cells
Level 2 (CCZ, Catalyst)	Area	$\mathcal{A}_{\text{CCZ}}$	$6d_2 \times 3d_2$
		$\mathcal{A}_{\text{cat}}$	$4d_2 \times 4d_2$
	Depth	$\mathcal{D}_{\text{CCZ}}$	$5.00 d_2$
		$\mathcal{D}_{\text{cat}}$	$1.00 d_2$
	Volume	$\mathcal{V}_{\text{CCZ}} + \mathcal{V}_{\text{cat}}$	1000 cells
Hardware	Cycle Time	t_{cycle}	1 μs
	Reaction Time	t_{react}	10 μs

FIG. 13. Parameters for physical resource estimation. Input parameters are presented for the AutoCCZ constructions from Ref. [112] while minding the improvements from Refs. [38,115].

distance d_1 refinement that we denote as $|T_1\rangle$. A number η of the $|T_1\rangle$ -states are then fed into a specialized level-two subfactory that produces a $|\text{CCZ}\rangle$ at distance $d_2 > d_1$. When desired, catalytic circuitry can mediate the $|\text{CCZ}\rangle \rightarrow 2|T_2\rangle$ conversion. Since the level-two CCZ-factory takes multiple inputs it must be accompanied by

$$\mathcal{N}_{T_1} = \left\lceil \eta \times \frac{\mathcal{D}_{T_1}}{\mathcal{D}_{\text{CCZ}}} \right\rceil \quad (\text{E1})$$

of the level-one T factories to avoid a bottleneck. Here, $\mathcal{D}_{T_1}$ and $\mathcal{D}_{\text{CCZ}}$ are cycle depths for the respective factories while η is a constant that depends to the target state ($\eta = 8$ for a terminal $|\text{CCZ}\rangle$ and $\eta = 4$ for a terminal $|T\rangle$). The overall cycle depth for a round of magic state distillation is then

$$\mathcal{D}_{\text{distill}} = \max(\mathcal{D}_{T_1}, \mathcal{D}_{\text{CCZ}} + \mathcal{D}_{\text{cat}}). \quad (\text{E2})$$

This expression also includes a $\mathcal{D}_{\text{cat}} = 1 \cdot d_2$ cycle overhead for catalysis. The AutoCCZ factory’s spatial footprint is an aggregate of these components,

$$\mathcal{A}_{\text{fact}} = \eta \mathcal{A}_{T_1} + \mathcal{A}_{\text{CCZ}} + \mathcal{A}_{\text{cat}} + \mathcal{A}_{\text{store}}. \quad (\text{E3})$$

Each contribution $\mathcal{A}_x = 2n_x d_x^2$ is specified in terms of the number of n_x of logical code tiles in each subcomponent and the number $2d_x^2$ of physical qubits in the corresponding tiles. The term $\mathcal{A}_{\text{store}}$ captures routing and storage for level-one T states (e.g., for caching redundant catalytic states to avoid failure).

Our magic state factories aim to deliver T states. Nonetheless, the $|\text{CCZ}\rangle$ is a valuable magic state since it can directly implement Toffoli gates. A refined strategy might leverage both $|\text{CCZ}\rangle$ and $|T_2\rangle$ -states, particularly in light of the Toffoli-based logic from our block encoding. Optimized layouts for QROM access have assumed this execution model [116]. We do not pursue this route since our logical estimates are based on T counts alone. Instead, we defer to a standard fast block layout which requires only $\mathcal{N}_{\text{data}} = 2N_{\text{logical}} + \sqrt{8N_{\text{logical}}} + 1$ tiles for the algorithm, ancilla, and routing [37]. This arrangement permits the consumption of at least one $|T\rangle$ -state every reaction cycle. We assume that magic state factories are placed around the periphery of this layout. This is important since our goal is to run a reaction-limited computation.

2. Error budget

Our foundational unit is the topological error per spacetime cell,

$$\mathbb{E}(d) = 0.1 \left(\frac{p_{\text{phys}}}{p_{\text{thresh}}} \right)^{(d+1)/2}, \quad (\text{E4})$$

where d is the code distance for the underlying logical qubit, $p_{\text{thresh}} = 0.01$ is the code error threshold, and p_{phys} is the physical error rate (e.g., for a two-qubit gate). We then subdivide our total nonalgorithmic error,

$$\epsilon_{\text{phys}} = \mathcal{D}_{\text{meas}} \mathcal{N}_{\text{data}} \mathbb{E}(d_2) + \frac{\mathcal{N}_T}{2} \epsilon_{T_2}, \quad (\text{E5})$$

where $\mathcal{D}_{\text{meas}}$ is the measurement depth. The first term captures topological error for data/routing qubits while the second term accounts for imperfectly prepared magic states. Following Litinski’s execution model, we translate T gates into Pauli

product rotations and commute Clifford gates to the end of the circuit [37]. The latter can be absorbed into terminal Pauli product measurements which allows us to forego their accounting in our estimates. We will assume a worst-case scenario where $\mathcal{D}_{\text{meas}}$ is equal to the T depth, though in principle this can be reduced through analysis of the target circuit. Note that Eq. (E5) contains a factor of $1/2$ since ϵ_{T2} corresponds to the catalytic production of two T -states. This error is calculated by propagating error through each stage of distillation,

$$\epsilon_{\text{inj}} = \mathcal{V}_{\text{inj}} \mathbb{E}(d_1/2) + p_{\text{phys}}, \quad (\text{E6})$$

$$\epsilon_{T1} = \mathcal{V}_{T1} \mathbb{E}(d_1) + 35\epsilon_{\text{inj}}^3, \quad (\text{E7})$$

$$\epsilon_{\text{CCZ}} = \mathcal{V}_{\text{CCZ}} \mathbb{E}(d_2) + 28\epsilon_{T1}^2, \quad (\text{E8})$$

$$\epsilon_{T2} = \mathcal{V}_{\text{cat}} \mathbb{E}(d_2) + \epsilon_{\text{CCZ}}, \quad (\text{E9})$$

where $\mathcal{V}_n$ is the spacetime volume of the corresponding factory in unit cells (e.g., one distance d logical qubit over d syndrome extraction cycles).

3. Operation and timing

Our objective is to run calculations that are limited by the reaction time t_{react} for an error correction cycle. In this case, the wall runtime is effectively $T_{\text{wall}} = \mathcal{D}_{\text{meas}} t_{\text{react}}$. We will require $\mathcal{N}_{\text{fact}} = N_T \mathcal{D}_{\text{fact}} t_{\text{cycle}} / T_{\text{wall}}$ magic state factories to maintain this rate by ensuring a consistent pool of T states. This defines the physical qubit footprint

$$N_{\text{phys}} = \mathcal{N}_{\text{fact}} \mathcal{A}_{\text{fact}} + 2\mathcal{N}_{\text{data}} d_2^2, \quad (\text{E10})$$

for our computation. This quantity is specified by the minimal code distance that performs our calculation with error $\epsilon_{\text{phys}} < \epsilon_{\text{target}}$ according to Eq. (E5).

APPENDIX F: PRODUCT FORMULA ESTIMATES

Product formula (PF) algorithms are an competitive alternative to qubitization [92] and drove the first proposals for gate-based simulation on a quantum computer [117]. The simplest of these require no ancilla for encoding, meaning that spatial overhead is limited to the target computational qubits. This is particularly valuable for near-term FTQC where logical qubits are anticipated to be scarce. Product formulas also

provide a systematic path toward increasingly accurate dynamics, albeit with an asymptotic scaling that is less favorable than qubitization. As we will show, the crossover from finite size to asymptotic behavior has significant implications when simulating NMR spectra.

In brief, product formula algorithms split the infinitesimal evolution $\mathcal{U}(\delta t) = \exp[-iH\delta t]$ generated by a Hamiltonian $H = \sum_{m=1}^M H_m$ into a product of terms,

$$\mathcal{S}(\delta t) = \prod_{v=1}^{\gamma} \prod_{\gamma=1}^M e^{-i a_v H_{\pi_v(\gamma)} \delta t}, \quad (\text{F1})$$

where v is termed the stage of the expansion, $\pi_v(\gamma)$ is a permutation that accounts for an optimal ordering of H_m , and a_v is a real coefficient. We assume that each H_m is unitary and that each factor in Eq. (F1) can be implemented as a (multi)qubit rotation. An extension to longer times $t \gg \delta t$ is generated by consecutive evolutions $\tilde{\mathcal{U}}(t) = \prod_{k=1}^r \mathcal{S}(\delta t)$, where $r = \lceil t/\delta t \rceil$ is the number of timesteps in the simulation.

The utility of a PF method will be contingent on the splitting strategy [118]. In particular, we want to minimize the computational overhead at a given level of error ϵ in an order k expansion,

$$\epsilon = \|\mathcal{U}(t) - \tilde{\mathcal{U}}(t)\| \leq O(\delta t^{k+1}), \quad (\text{F2})$$

where $\mathcal{U}(t)$ is the exact evolution at time t . It is generally difficult to compute ϵ explicitly to give tight bounds for a generic p -local Hamiltonian [90,92,119]. Our discussion will consider a specific splitting based on the k th order Trotter-Suzuki formula. This is defined by starting with the second-order expansion,

$$S_2(\delta t) = \left(\prod_{\gamma=1}^M e^{-i H_{\gamma} \delta t/2} \right) \left(\prod_{\gamma=M}^1 e^{-i H_{\gamma} \delta t/2} \right), \quad (\text{F3})$$

and recursively generating its higher-order counterparts,

$$S_{2k}(\delta t) = [S_{2k-1}(u_k \delta t)]^2 S_{2k-1}((1 - 4u_k) \delta t) S_{2k-1}(u_k \delta t)^2, \quad (\text{F4})$$

where $u_k = 1/(4 - 4^{1/(2k-1)})$ [120–122]. We focus on the second-order case ($k = 2$), which has a robust trade-off between error scaling and computational overhead. The error in time-evolution unitaries is likewise bounded:

$$\|\mathcal{U}(t) - \tilde{\mathcal{U}}(t)\| \leq \frac{t^3}{12r^2} \sum_{i=1}^M \left(\left\| \sum_{j,k=i+1}^M [H_j, [H_k, H_i]] \right\| + \frac{1}{2} \left\| \sum_{j=i+1}^M [H_j, [H_j, H_i]] \right\| \right). \quad (\text{F5})$$

The only free parameter in our analysis is the number of product formula steps r , making it the sole determinant of product formula error.

While formally useful, direct computation of the operator norm in Eq. (F5) is not always numerically convenient. Instead, we quantify error $\Delta E = |E - \tilde{E}|$ in terms of the energy eigenvalues for our target Hamiltonian H and the effective Hamiltonian $\tilde{H}$ that is realized by the product

formula. A straightforward expansion gives us the bound $\Delta E \leq \|\mathcal{U}(t) - \tilde{\mathcal{U}}(t)\|/t$ which equates to the operator norm at unit time. We adopt this arrangement.

To derive a resource estimate we must evaluate the combinatorially large number of commutator terms or, more practically, bound them with the triangle inequality. However, product formula methods will often empirically outperform their bounds, including Eq. (F5), by a substantial margin




FIG. 14. Circuit implementing a Pauli product rotation. We implement an arbitrary Pauli product rotation $\exp[-i\alpha P]$ using a ladder of CNOT gates and a single-qubit rotation in the z basis. This is depicted for a string $P = P_1 \otimes P_2 \otimes P_3$, where factors other than $P_i = Z$ are obtained by conjugating with Clifford gates (e.g., $C_i = H$ when $P_i = X$ and $C_i = HS$ when $P_i = Y$). Our nuclear spin Hamiltonians will only require one- and two-qubit rotations.

[90–92,123]. This makes it difficult to reconcile the temporal resources that are required to meet an error threshold. To address the discrepancy, we have also explicitly simulated a series of classically tractable nuclear spin Hamiltonians at increasing r as described in Appendix G. Logical resources are then quantified for both formal and empirical error bounds by constructing explicit time evolution circuits. In practice, the circuit constituents amount to a series of Pauli product rotations with one- or two-qubit support. These are implemented using the convention from Fig. 14.

APPENDIX G: TENSOR NETWORK SIMULATIONS

We have used tensor network simulations to capture the empirical scaling of product formula error ϵ with respect to timestep $\delta t = T/r$ for a series of nuclear spin Hamiltonians. This is straightforward since the second-order Trotter-Suzuki decomposition from Eq. (F3) corresponds to the TEBD of an MPS [103,124–126]. The general structure of an MPS and a high-level description of this approximation were outlined in Appendix C.

We implement these calculations by mapping our logical qubit register onto the one-dimensional tensor train from an MPS [Eq. (B1)]. Each Pauli rotation gate $\exp[-i\alpha P]$ then translates to a tensor G_{ij}^{kl} that can be contracted with the MPS in direct analogy with the PF circuit layout. This arrangement has implications for the feasibility of classical computation. In particular, the required tensor contractions are made efficient by using gauge transformations to place segments of the MPS on either side of the target qubits in left- or right-canonical form in Fig. 15(a). Tensors to the left (right) of this orthogonality center then become isometries $\sum_i (A_i^{[n]})^\dagger A_i^{[n]} = I$ and explicit summation can be limited indices between sites where the gate is applied. The numerical cost is ultimately dominated by a subsequent SVD, which factors the resulting rank-4 tensor into a pair of rank-3 tensors $\Gamma^{[n]}, \Gamma^{[n+1]}$ and a diagonal singular value matrix Λ^n . This reestablishes the MPS ansatz while incurring a $O(Md^3\chi^2)$ overhead when applying M gates to a state with physical dimension $d = 2$ and bond dimension χ . From here, the MPS can be placed in left (right) canonical form through left (right) multiplication of Λ^n into an adjacent site tensor.




FIG. 15. Application of two-qubit gates to an MPS. (a) A local gate G_{ij}^{kl} can be efficiently applied to consecutive sites of a left-canonicalized MPS. To accomplish this, we apply a gauge transformation to the chain and make site n the orthogonality center. The gate tensor G_{ij}^{kl} is then contracted with physical indices of $A^{[n]}$ and $A^{[n+1]}$ to give a rank-4 tensor. We recover the MPS ansatz by performing a singular value decomposition. If desired, the diagonal singular value matrix Λ^n can be absorbed into the left or right factors. (b) Long-range gates are implemented between $A^{[n]}, A^{[n+2]}$ by using swap gates $S_{n,n+1}$ to exchange consecutive site indices. Note that longer-range interactions will require additional swaps.

The preceding algorithm is defined for gates between consecutive qubits, though our nuclear spin Hamiltonians have a preponderance of long-range couplings. We accommodate these by using effective swap gates to reorder the MPS and make the target indices adjacent as in Fig. 15(b). Although the required swaps can be performed at nominal cost, this shuffling will lead to pronounced long-range entanglement. This increases the bond dimension χ between all intermediate tensors in the MPS and thus the cost of applying subsequent gates. Alternatively, one build a matrix product operator (MPO) that implements the long-range gate. This strategy is comparable for two-qubit gates but can be convenient for higher-weight Pauli products [127].

Rapid entanglement generation, the associated truncation errors, and deviations from unitarity are known to limit the accuracy of TEBD when long-range interactions are present [128] (though these are, in part, due to the underlying product formula decomposition). Conversely, tangent-space methods like the TDVP can circumvent these shortcomings at the cost of a fully dynamic perspective [103,129–132]. Both methods are foundational to our strategy. In particular, we leverage the TDVP evolution $|\Phi_{\text{ref}}(t)\rangle$ generated by a given spin Hamiltonian as a high-accuracy approximation to the physical dynamics. We then use TEBD to generate a series of states $|\Psi_r(t)\rangle$ corresponding to r steps of a Trotter-Suzuki algorithm and calculate energy $\Delta E = |E_r - E_{\text{ref}}|$ and state error $\epsilon_s = 1 - |\langle \Psi_r(t) | \Phi_{\text{ref}}(t) \rangle|$ with respect to the TDVP reference. The iTensor library is used to facilitate numerical tensor operations [133]. Note that local TDVP can simulate quantum

gates and may offer better performance TEBD, though we do not adopt it here [105].

The state error from our procedure has multiple contributions with respect to the exact, physical dynamics $|\Phi_{\text{exact}}(t)\rangle$. We decompose these as

$$\epsilon_{\text{state}} = \epsilon_{\text{method}} + \epsilon_{\text{trunc}} + \epsilon_{\text{PF}}. \quad (\text{G1})$$

This includes a term ϵ_{method} inherent to our implementation of the tensor network algorithms, a term ϵ_{trunc} arising from MPS truncation at the designated bond dimension χ and singular value λ_{min} cutoffs, as well as the desired product formula error ϵ_{PF} (an analogous decomposition holds for the energy error). The component ϵ_{method} is largely defined by the under-

lying numerical methods and floating point precision, though these may have some emergent cutoff dependence. We can minimize this and ϵ_{trunc} by adopting a high bond dimension cutoff χ relative to the problem physics and by using a very low cutoff for singular value truncation. In particular, an increasingly large value of χ will minimize the distance between the calculated states and $|\Phi_{\text{exact}}(t)\rangle$ under any unitarily invariant norm. While ϵ_{PF} may have a different scaling between the exact dynamics and our approximate reference problem, we expect this to be negligible for most Hamiltonians. Thus, we set λ_{min} to the limit of floating point precision, use the largest χ that is feasible for practical calculations, and treat ϵ_{state} as a tight upper bound on the product formula error ϵ_{PF} without qualifying it further.

-
- [1] R. R. Ernst, G. Bodenhausen, and A. Wokaun, *Principles of Nuclear Magnetic Resonance in One and Two Dimensions* (Wiley, Hoboken, NJ, 1990).
 - [2] Y. Hu, K. Cheng, L. He, X. Zhang, B. Jiang, L. Jiang, C. Li, G. Wang, Y. Yang, and M. Liu, NMR-based methods for protein analysis, *Anal. Chem.* **93**, 1866 (2021).
 - [3] M. Marušič, M. Toplishek, and J. Plavec, NMR of RNA—Structure and interactions, *Curr. Opin. Struct. Biol.* **79**, 102532 (2023).
 - [4] C. Fontana and G. Widmalm, Primary structure of glycans by NMR spectroscopy, *Chem. Rev.* **123**, 1040 (2023).
 - [5] G. S. Rule and T. K. Hitchens, *Fundamentals of Protein NMR Spectroscopy* (Springer, New York, 2006).
 - [6] R. Sprangers, A. Velyvis, and L. E. Kay, Solution NMR of supramolecular complexes: Providing new insights into function, *Nat. Methods* **4**, 697 (2007).
 - [7] M. Weingarth and M. Baldus, Solid-state NMR-based approaches for supramolecular structure elucidation, *Acc. Chem. Res.* **46**, 2037 (2013).
 - [8] R. Rosenzweig and L. E. Kay, Solution NMR spectroscopy provides an avenue for the study of functionally dynamic molecular machines: The example of protein disaggregation, *J. Am. Chem. Soc.* **138**, 1466 (2016).
 - [9] R. Puthenveetil and O. Vinogradova, Solution NMR: A powerful tool for structural and functional studies of membrane proteins in reconstituted environments, *J. Biol. Chem.* **294**, 15914 (2019).
 - [10] P. Bayer, A. Matena, and C. Beuck, NMR spectroscopy of supramolecular chemistry on protein surfaces, *Beilstein J. Org. Chem.* **16**, 2505 (2020).
 - [11] B. Reif, S. E. Ashbrook, L. Emsley, and M. Hong, Solid-state NMR spectroscopy, *Nat. Rev. Methods Primers* **1**, 2 (2021).
 - [12] K. M. Cecil, Proton magnetic resonance spectroscopy: Technique for the neuroradiologist, *Neuroimaging Clin. North Am.* **23**, 381 (2013).
 - [13] G. Öz, *et al.*, Clinical proton MR spectroscopy in central nervous system disorders, *Radiology* **270**, 658 (2014).
 - [14] J. M. Tognarelli, M. Dawood, M. I. F. Sharif, V. P. B. Grover, M. M. E. Crossey, I. J. Cox, S. D. Taylor–Robinson, and M. J. W. McPhail, Magnetic resonance spectroscopy: Principles and techniques: Lessons for clinicians, *J. Clin. Exp. Hepatol.* **5**, 320 (2015).
 - [15] R. W. Brown, Y.-C. N. Cheng, E. M. Haacke, M. R. Thompson, and R. Venkatesan, *Magnetic Resonance Imaging: Physical Principles and Sequence Design* (Wiley, 2014).
 - [16] J. M. Soares, R. Magalhães, P. S. Moreira, A. Sousa, E. Ganz, A. Sampaio, V. Alves, P. Marques, and N. Sousa, A Hitchhiker's guide to functional magnetic resonance imaging, *Front. Neurosci.* **10**, 515 (2016).
 - [17] C. Westbrook and J. Talbot, *MRI in Practice* (Wiley, 2018).
 - [18] J. Bernarding, G. Buntkowsky, S. Macholl, S. Hartwig, M. Burghoff, and L. Trahms, *J*-coupling nuclear magnetic resonance spectroscopy of liquids in nT fields, *J. Am. Chem. Soc.* **128**, 714 (2006).
 - [19] J. W. Blanchard, D. Budker, and A. Trabesinger, Lower than low: Perspectives on zero- to ultralow-field nuclear magnetic resonance, *J. Magn. Reson.* **323**, 106886 (2021).
 - [20] S. J. DeVience, M. Greer, S. Mandal, and M. S. Rosen, Homonuclear *J*-coupling spectroscopy and low magnetic fields using spin-lock induced crossing, *ChemPhysChem* **22**, 2128 (2021).
 - [21] T. Theis, P. Ganssle, G. Kervern, S. Knappe, J. Kitching, M. P. Ledbetter, D. Budker, and A. Pines, Parahydrogen-enhanced zero-field nuclear magnetic resonance, *Nat. Phys.* **7**, 571 (2011).
 - [22] D. A. Barskiy, J. W. Blanchard, D. Budker, J. Eills, S. Pustelny, K. F. Shebertsov, M. C. D. Tayler, and A. H. Trabesinger, Zero- to ultralow-field nuclear magnetic resonance, *Prog. Nucl. Magn. Reson. Spectrosc.* **148**, 101558 (2025).
 - [23] J. W. Blanchard, M. P. Ledbetter, T. Theis, M. C. Butler, D. Budker, and A. Pines, High-resolution zero-field NMR *J*-spectroscopy of aromatic compounds, *J. Am. Chem. Soc.* **135**, 3607 (2013).
 - [24] M. Emondts, M. P. Ledbetter, S. Pustelny, T. Theis, B. Patton, J. W. Blanchard, M. C. Butler, D. Budker, and A. Pines, Long-lived heteronuclear spin-singlet states in liquids at a zero magnetic field, *Phys. Rev. Lett.* **112**, 077601 (2014).
 - [25] J. W. Blanchard, T. F. Sjolander, J. P. King, M. P. Ledbetter, E. H. Levine, V. S. Bajaj, D. Budker, and A. Pines, Measurement of untruncated nuclear spin interactions via zero- to ultralow-field nuclear magnetic resonance, *Phys. Rev. B* **92**, 220202(R) (2015).

- [26] S. Appelt, F. W. Häsing, U. Sieling, A. Gordji-Nejad, S. Glöggler, and B. Blümich, Paths from weak to strong coupling in NMR, *Phys. Rev. A* **81**, 023420 (2010).
- [27] A. Karabanov, I. Kuprov, G. T. P. Charnock, A. van der Drift, L. J. Edwards, and W. Köckenberger, On the accuracy of the state space restriction approximation for spin dynamics simulations, *J. Chem. Phys.* **135**, 084106 (2011).
- [28] A. Wilzewski, S. Afach, J. W. Blanchard, and D. Budker, A method for measurement of spin-spin couplings with sub-mHz precision using zero- to ultralow-field nuclear magnetic resonance, *J. Magn. Reson.* **284**, 66 (2017).
- [29] M. G. Algaba, M. Ponce-Martinez, C. Munuera-Javaloy, V. Pina-Canelles, M. J. Thapa, B. G. Taketani, M. Leib, I. de Vega, J. Casanova, and H. Heimonen, Co-design quantum simulation of nanoscale NMR, *Phys. Rev. Res.* **4**, 043089 (2022).
- [30] T. E. O'Brien, L. B. Ioffe, Y. Su, D. Fushman, H. Neven, R. Babbush, and V. Smelyanskiy, Quantum computation of molecular structure using data from challenging-to-classically simulate nuclear magnetic resonance experiments, *PRX Quantum* **3**, 030345 (2022).
- [31] K. Seetharam, D. Biswas, C. Noel, A. Risinger, D. Zhu, O. Katz, S. Chattopadhyay, M. Cetina, C. Monroe, E. Demler, and D. Sels, Digital quantum simulation of NMR experiments, *Sci. Adv.* **9**, eadh2594 (2023).
- [32] A. Burov, O. Nagl, and C. Javerzac-Galy, Towards quantum utility for NMR quantum simulation on a NISQ computer, *arXiv:2404.17548*.
- [33] D. A. Abanin *et al.* (Google Quantum AI and Collaborators), Observation of constructive interference at the edge of quantum ergodicity, *Nature (London)* **646**, 825 (2025).
- [34] C. Zhang, *et al.*, Quantum computation of molecular geometry via many-body nuclear spin echoes, *arXiv:2510.19550*.
- [35] D. Sels, H. Dashti, S. Mora, O. Demler, and E. Demler, Quantum approximate Bayesian computation for NMR model inference, *Nat. Mach. Intell.* **2**, 396 (2020).
- [36] K. Seetharam, D. Sels, and E. Demler, Platform tailored code-sign of gate-based quantum simulation, *Phys. Rev. Res.* **6**, 033059 (2024).
- [37] D. Litinski, A game of surface codes: Large-scale quantum computing with lattice surgery, *Quantum* **3**, 128 (2019).
- [38] C. Gidney and M. Ekerå, How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits, *Quantum* **5**, 433 (2021).
- [39] R. Babbush, R. King, S. Boixo, W. Huggins, T. Khattar, G. H. Low, J. R. McClean, T. O'Brien, and N. C. Rubin, The grand challenge of quantum applications, *arXiv:2511.09124*.
- [40] The most straightforward digital quantum simulations will reproduce a pure state density matrix as opposed to the mixed state generated by relaxation operators. Instead, we treat relaxation through the decaying exponential factor in Eq. (2). Thus, the longitudinal and transverse magnetization profiles become equivalent for determining resonances in the NMR spectrum. It may be prudent to work with S^+ or S^- when developing pulse sequences due to the difference in objectives and potential methodology.
- [41] There will be a chemical shift at small but finite fields, though resonances can still be indistinguishable due to finite linewidths.
- [42] Stated generally, the scalar coupling will be a tensor quantity. However, only the scalar isotropic component is relevant for the liquid-phase experiments that we consider. Other protocols can access different components of the coupling tensor and extract different information.
- [43] M. C. Tayler and L. F. Gladden, Scalar relaxation of NMR transitions at ultralow magnetic field, *J. Magn. Reson.* **298**, 101 (2019).
- [44] Y. Zhou, E. M. Stoudenmire, and X. Waintal, What limits the simulation of quantum computers? *Phys. Rev. X* **10**, 041038 (2020).
- [45] T. Ayral, T. Louvet, Y. Zhou, C. Lambert, E. M. Stoudenmire, and X. Waintal, Density-matrix renormalization group algorithm for simulating quantum circuits with a finite fidelity, *PRX Quantum* **4**, 020304 (2023).
- [46] We drop the isotopic label in subscripts since all nuclei are assumed to be the most abundant spin-1/2 isotope.
- [47] W. Kemp, *NMR in Chemistry: A Multinuclear Introduction* (Macmillan Education Limited, 1986).
- [48] H. J. Hogben, M. Krzystyniak, G. T. P. Charnock, P. J. Hore, and I. Kuprov, *Spinach*—A software library for simulation of spin dynamics in large systems, *J. Magn. Reson.* **208**, 179 (2011).
- [49] D. Cremer and J. Gräfenstein, Calculation and analysis of NMR spin-spin coupling constants, *Phys. Chem. Chem. Phys.* **9**, 2791 (2007).
- [50] T. Helgaker, M. Jaszunski, and P. Swider, Calculation of NMR spin-spin coupling constants in Strychnine, *J. Org. Chem.* **81**, 11496 (2016).
- [51] S. J. DeVience and M. S. Rosen, Homonuclear J -coupling spectroscopy using J -synchronized echo detection, *J. Magn. Reson.* **341**, 107244 (2022).
- [52] P. Tzvetkova, U. Sternberg, T. Gloge, A. Navaro-Vázquez, and B. Luy, Configuration determination by residual dipolar couplings: Accessing the full conformational space by molecular dynamics with tensorial constraints, *Chem. Sci.* **10**, 8774 (2019).
- [53] The experimental budget is also guided by other parameters, $t_{\max} = n_{\text{points}}/2W$, such as the desired spectral width W and the number of points n_{points} sampled in the signal. There is likewise a trade-off between satisfying Nyquist sampling requirements and practical spectrometer availability. Thus, it difficult to map literature parameters back to quantum simulation inputs.
- [54] G. H. Low and I. L. Chuang, Optimal Hamiltonian simulation by quantum signal processing, *Phys. Rev. Lett.* **118**, 010501 (2017).
- [55] G. H. Low and I. L. Chuang, Hamiltonian simulation by qubitization, *Quantum* **3**, 163 (2019).
- [56] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC 2019)* (Association for Computing Machinery, New York, NY, 2019), pp. 193–204.
- [57] J. M. Martyn, Z. M. Rossi, A. K. Tan, and I. L. Chuang, Grand unification of quantum algorithms, *PRX Quantum* **2**, 040203 (2021).
- [58] Y. Dong, X. Meng, K. B. Whaley, and L. Lin, Efficient phase-factor evaluation in quantum signal processing, *Phys. Rev. A* **103**, 042419 (2021).

- [59] 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).
- [60] J. M. Martyn, Y. Liu, Z. E. Chin, and I. L. Chuang, Efficient fully-coherent quantum signal processing algorithms for real-time dynamics simulation, *J. Chem. Phys.* **158**, 024106 (2023).
- [61] D. Motlagh and N. Wiebe, Generalized quantum signal processing, *PRX Quantum* **5**, 020368 (2024).
- [62] R. Rines, K. Obenland, and I. Chuang, Empirical determination of the simulation capacity of a near-term quantum computer, [arXiv:1905.10724](https://arxiv.org/abs/1905.10724).
- [63] Our practical use of the GQSP formalism actually corresponds to a quantum eigenvalue transform. However, we retain the name GQSP to maintain consistency with literature. However, it is currently unclear if GQSP can be extended to the more general context of the quantum singular value transform (QSVT).
- [64] S. Yamamoto and N. Yoshioka, Robust angle finding for generalized quantum signal processing, [arXiv:2402.03016](https://arxiv.org/abs/2402.03016).
- [65] 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).
- [66] pyLIQTR: A Python library for fault-tolerant quantum algorithms, GitHub (2024), <https://github.com/isi-usc-edu/pyLIQTR>.
- [67] M. P. Harrigan, T. Khattar, C. Yuan, A. Peduri, N. Yosri, F. D. Malone, R. Babbush, and N. C. Rubin, Expressing and analyzing quantum algorithms with Qualtran, [arXiv:2409.04643](https://arxiv.org/abs/2409.04643).
- [68] QUALTRAN: A quantum algorithms translator, GitHub, 2024, <https://github.com/quantumlib/Qualtran>.
- [69] J. Eills, D. Budker, S. Cavagnero, E. Y. Chekmenev, S. J. Elliott, S. Jannin, A. Lesange, J. Matysik, T. Meersmann, T. Prisner, J. A. Reimer, H. Yang, and I. V. Koptiyug, Spin hyperpolarization in modern magnetic resonance, *Chem. Rev.* **123**, 1417 (2023).
- [70] D. Grinko, J. Gacon, C. Zoufal, and S. Woerner, Iterative quantum amplitude estimation, *npj Quantum Inf.* **7**, 52 (2021).
- [71] A. A. Agrawal, J. Job, T. L. Wilson, S. N. Saadatmand, M. J. Hodson, J. Y. Mutus, A. Caesura, P. D. Johnson, J. E. Elenewski, K. J. Morrell, and A. F. Kemper, Quantifying fault-tolerant simulation of strongly correlated systems using the Fermi-Hubbard model, [arXiv:2406.06511](https://arxiv.org/abs/2406.06511).
- [72] D. Donoho, Compressed sensing, *IEEE Trans. Inf. Theory* **52**, 1289 (2006).
- [73] M. J. Bostock, D. J. Holland, and D. Nietlispach, Improving resolution in multidimensional NMR using random quadrature detection with compressed sensing reconstruction, *J. Biomol. NMR* **68**, 67 (2017).
- [74] M. Bostock and D. Nietlispach, Compressed sensing: Reconstruction of non-uniformly sampled multidimensional NMR data, *Concepts Magn. Reson. Part A* **46A**, e21438 (2017).
- [75] S. Robson, H. Arthanari, S. G. Hyberts, and G. Wagner, Nonuniform sampling for NMR spectroscopy, *Methods Enzymol.* **614**, 263 (2019).
- [76] F. Delaglio, G. S. Walker, K. A. Farley, R. Sharma, J. C. Hoch, L. W. Arbogast, R. G. Brinson, and J. P. Marino, Non-uniform sampling for all: More NMR spectral quality, less measurement time, *Am. Pharm. Rev.* **20**, 339681 (2017).
- [77] R. A. E. Carr, M. Congreve, C. W. Murray, and D. C. Rees, Fragment-based lead discovery: Leads by design, *Drug Discov. Today* **10**, 987 (2005).
- [78] At $N = 16$, a matrix representation of the time-evolution operator would be 64 Gb (assuming 64-bit complex entries). This would be time-consuming to integrate on the timescales required for simulating spectra. At $N = 20$, the operators would require 16 384 Gb matrices, which is prohibitive for any classical calculation.
- [79] N. C. Menicucci and C. M. Caves, Local realistic model for the dynamics of bulk-ensemble NMR information processing, *Phys. Rev. Lett.* **88**, 167901 (2002).
- [80] M. M. Rams and M. Zwolak, Breaking the entanglement barrier: Tensor network simulation of quantum transport, *Phys. Rev. Lett.* **124**, 137701 (2020).
- [81] G. Wójtowicz, J. E. Elenewski, M. M. Rams, and M. Zwolak, Open system tensor networks and Kramers' crossover for quantum transport, *Phys. Rev. A* **101**, 050301(R) (2020).
- [82] G. Wójtowicz, J. E. Elenewski, M. M. Rams, and M. Zwolak, Dual current anomalies and quantum transport within extended reservoir simulations, *Phys. Rev. B* **104**, 165131 (2021).
- [83] J. E. Elenewski, G. Wójtowicz, M. M. Rams, and M. Zwolak, Performance of reservoir discretizations in quantum transport simulations, *J. Chem. Phys.* **155**, 124117 (2021).
- [84] While these clusters span the whole molecule, the overhead does not correspond to the maximum number of possible Hamiltonian terms (and thus T -gate count). This is due to our finite range for dipolar interactions, which is smaller than the maximal extent of the molecule. Stated differently, these models have a fully connected spin network, but do not saturate the edge count.
- [85] P. Klukowski, F. F. Damberger, F. H.-T. Allain, H. Iwai, H. Kadavath, T. A. Ramelot, G. T. Montelione, R. Riek, and P. Güntert, The 100-protein NMR spectra dataset: A resource for biomolecular NMR data analysis, *Sci. Data* **11**, 30 (2024).
- [86] N. P. Breuckmann and J. N. Eberhardt, Quantum low-density parity-check codes, *PRX Quantum* **2**, 040101 (2021).
- [87] S. Bravyi, A. W. Cross, J. M. Gambetta, D. Maslov, P. Rall, and T. J. Yoder, High-threshold and low-overhead fault-tolerant quantum memory, *Nature (London)* **627**, 778 (2024).
- [88] C. Gidney, M. Newman, P. Brooks, and C. Jones, Yoked surface codes, *Nat. Commun.* **16**, 4498 (2025).
- [89] C. Gidney, How to factor 2048 bit RSA integers with less than a million noisy qubits, [arXiv:2505.15917](https://arxiv.org/abs/2505.15917).
- [90] A. M. Childs, D. Maslov, Y. Nam, N. J. Ross, and Y. Su, Toward the first quantum simulation with quantum speedup, *Proc. Natl. Acad. Sci. USA* **115**, 9456 (2018).
- [91] M. C. Tran, S.-K. Chu, Y. Su, A. M. Childs, and A. V. Gorshkov, Destructive error interference in product-formula lattice simulation, *Phys. Rev. Lett.* **124**, 220502 (2020).
- [92] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, Theory of Trotter error with commutator scaling, *Phys. Rev. X* **11**, 011020 (2021).
- [93] Y.-H. Chen, A. Kalev, and I. Hen, Quantum algorithm for time-dependent Hamiltonian simulation by permutation expansion, *PRX Quantum* **2**, 030342 (2021).
- [94] D. An, D. Fang, and L. Lin, Time-dependent Hamiltonian simulation of highly oscillatory dynamics and superconvergence for Schrödinger equation, *Quantum* **6**, 690 (2022).

- [95] H. Zhao, M. Bukov, M. Heyl, and R. Moessner, Adaptive Trotterization for time-dependent Hamiltonian quantum dynamics using piecewise conservation laws, *Phys. Rev. Lett.* **133**, 010603 (2024).
- [96] J. L. Bosse, A. M. Childs, C. Derby, F. M. Gambetta, A. Montanaro, and R. A. Santos, Efficient and practical Hamiltonian simulation from time-dependent product formulas, *Nat. Commun.* **16**, 2673 (2025).
- [97] T. A. Halgren, Merck molecular force field. I. Basis, form, scope, parametrization, and performance of MMFF94, *J. Comput. Chem.* **17**, 490 (1996).
- [98] N. M. O’Boyle, M. Banck, C. A. James, C. Morley, T. Vandermeersch, and G. R. Hutchinson, Open Babel: An open chemical toolbox, *J. Cheminform.* **3**, 33 (2011).
- [99] T. Scientific, Maybridge fragment libraries (2025), <https://www.thermofisher.com/us/en/home/industrial/pharmabiopharma/drug-discovery-development/screening-compounds-libraries-hit-identification/maybridge-fragment-libraries.html>.
- [100] S. Kim, J. Chen, T. Cheng, A. Gindulyte, J. He, S. He, Q. Li, B. A. Shoemaker, P. A. Thiessen, B. Yu, L. Zaslavsky, J. Zhang, and E. E. Bolton, PubChem 2025 update, *Nucleic Acids Res.* **53**, D1516 (2025).
- [101] U.S. Food and Drug Administration, Approved Drug Products with Therapeutic Equivalence Evaluations (2025), <https://www.accessdata.fda.gov/scripts/cder/ob/index.cfm>.
- [102] RCSB Protein Databank (2025), <https://www.rcsb.org/>.
- [103] S. Paeckel, T. Köhler, A. Swoboda, S. R. Manmana, U. Schollwöck, and C. Hubig, Time-evolution methods for matrix-product states, *Ann. Phys.* **411**, 167998 (2019).
- [104] J. I. Cirac, D. Pérez-García, N. Schuch, and F. Verstraete, Matrix product states and projected entangled pair states: Concepts, symmetries, theorems, *Rev. Mod. Phys.* **93**, 045003 (2021).
- [105] A. Sander, M. Fröhlich, M. Ali, M. Eigel, J. Eisert, M. Hintermüller, C. B. Mendl, R. M. Milbradt, and R. Wille, Quantum circuit simulation with a local time-dependent variational principle, *arXiv:2508.10096*.
- [106] Z. Zhu and S. R. White, Spin liquid phase of the $S = \frac{1}{2}J_1-J_2$ Heisenberg model on the triangular lattice, *Phys. Rev. B* **92**, 041105(R) (2015).
- [107] W.-J. Hu, S.-S. Gong, W. Zhu, and D. N. Sheng, Competing spin-liquid states in the spin-1/2 Heisenberg model on the triangular lattice, *Phys. Rev. B* **92**, 140403(R) (2015).
- [108] N. E. Sherman, M. Dupont, and J. E. Moore, Spectral function of the Heisenberg model on the triangular lattice, *Phys. Rev. B* **107**, 165146 (2023).
- [109] C. Cade, L. Mineh, A. Montanaro, and S. Stanisic, Strategies for solving the Fermi-Hubbard model on near-term quantum computers, *Phys. Rev. B* **102**, 235122 (2020).
- [110] A. Kan and B. Symons, Resource-optimized fault-tolerant simulation of the Fermi-Hubbard model and high-temperature superconductor models, *npj Quantum Inf.* **11**, 138 (2025).
- [111] D. Horsman, A. G. Fowler, S. Devitt, and R. Van Meter, Surface code quantum computing by lattice surgery, *New J. Phys.* **14**, 123011 (2012).
- [112] C. Gidney and A. G. Fowler, Efficient magic state factories with a catalyzed $|CCZ\rangle$ to $2|T\rangle$ transformation, *Quantum* **3**, 135 (2019).
- [113] A. G. Fowler and C. Gidney, Low overhead quantum computation using lattice surgery, *arXiv:1808.06709*.
- [114] Y. Li, A magic state’s fidelity can be superior to the operations that created it, *New J. Phys.* **17**, 023037 (2015).
- [115] C. Gidney and A. G. Folwer, Flexible layout of surface code computations using Auto CCZ states, *arXiv:1905.08916*.
- [116] J. Lee, D. W. Berry, C. Gidney, W. J. Huggins, J. R. McClean, N. Wiebe, and R. Babbush, Even more efficient quantum computations of chemistry through tensor hypercontraction, *PRX Quantum* **2**, 030305 (2021).
- [117] S. Lloyd, Universal quantum simulators, *Science* **273**, 1073 (1996).
- [118] J. Ostmeier, Optimized Trotter decompositions for classical and quantum computing, *J. Phys. A: Math. Theor.* **56**, 285303 (2023).
- [119] 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).
- [120] M. Suzuki, Decomposition formulas of exponential operators and Lie exponentials with some applications to quantum mechanics and statistical physics, *J. Math. Phys.* **26**, 601 (1985).
- [121] M. Suzuki, Quantum statistical Monte Carlo methods and applications to spin systems, *J. Stat. Phys.* **43**, 883 (1986).
- [122] M. Suzuki, Fractal decomposition of exponential operators with applications to many-body theories and Monte Carlo simulations, *Phys. Lett. A* **146**, 319 (1990).
- [123] C.-F. Chen and F. G. S. L. Brandão, Average-case speedup for product formulas, *Commun. Math. Phys.* **405**, 32 (2024).
- [124] I. V. Oseledets, Tensor-train decomposition, *SIAM J. Sci. Comput.* **33**, 2295 (2011).
- [125] G. Vidal, Efficient classical simulation of slightly entangled quantum computations, *Phys. Rev. Lett.* **91**, 147902 (2003).
- [126] G. Vidal, Efficient simulation of one-dimensional quantum many-body systems, *Phys. Rev. Lett.* **93**, 040502 (2004).
- [127] M. P. Zaletel, R. S. K. Mong, C. Karrasch, J. E. Moore, and F. Pollmann, Time-evolving a matrix product state with long-ranged interactions, *Phys. Rev. B* **91**, 165112 (2015).
- [128] K. Hémerly, F. Pollmann, and D. J. Luitz, Matrix product state approaches to operator spreading in ergodic quantum systems, *Phys. Rev. B* **100**, 104303 (2019).
- [129] J. Haegeman, J. I. Cirac, T. J. Osborne, and I. Pižorn, Time-dependent variational principle for quantum lattices, *Phys. Rev. Lett.* **107**, 070601 (2011).
- [130] J. Haegeman, T. J. Osborne, and F. Verstraete, Post-matrix product state methods: To tangent space and beyond, *Phys. Rev. B* **88**, 075133 (2013).
- [131] J. Haegeman, C. Lubich, I. Oseledets, B. Vandereycken, and F. Verstraete, Unifying time-evolution and optimization with matrix product states, *Phys. Rev. B* **94**, 165116 (2016).
- [132] L. Vanderstraeten, J. Haegeman, and F. Verstraete, Tangent-space methods for uniform matrix product states, *SciPost Phys. Lect. Notes* **7** (2019).
- [133] M. Fishman, S. R. White, and E. M. Stoudenmire, The ITensor software library for tensor network calculations, *SciPost Phys. Codebases* **4** (2022).