Co-Design quantum simulation of nanoscale NMR

Manuel G. Algaba [1,*,†] Mario Ponce-Martinez [1,2,*,‡] Carlos Munuera-Javaloy [3] Vicente Pina-Canelles [1]
Manish J. Thapa [1] Bruno G. Taketani [1] Martin Leib [1] Inés de Vega [1,2] Jorge Casanova [3,4] and Hermanni Heimonen [5]

[1]*IQM Quantum Computers, Nymphenburgerstr. 86, 80636 Munich, Germany*

[2]*Department of Physics and Arnold Sommerfeld Center for Theoretical Physics, Ludwig-Maximilians-Universität München, Theresienstrasse 37, 80333 Munich, Germany*

[3]*Department of Physical Chemistry, University of the Basque Country UPV/EHU, Apartado 644, 48080 Bilbao, Spain*

[4]*IKERBASQUE, Basque Foundation for Science, Plaza Euskadi 5, 48009 Bilbao, Spain*

[5]*IQM Quantum Computers, Keilaranta 19, FI-02150 Espoo, Finland*

Quantum computers have the potential to efficiently simulate the dynamics of nanoscale NMR systems. In this work, we demonstrate that a noisy intermediate-scale quantum computer can be used to simulate and predict nanoscale NMR resonances. In order to minimize the required gate fidelities, we propose a superconducting application-specific Co-Design quantum processor that reduces the number of SWAP gates by over 90% for chips with more than 20 qubits. The processor consists of transmon qubits capacitively coupled via tunable couplers to a central co-planar waveguide resonator with a quantum circuit refrigerator (QCR) for fast resonator reset. The QCR implements the nonunitary quantum operations required to simulate nuclear hyperpolarization scenarios.

I. INTRODUCTION

Computer simulations are the backbone of scientific research and technological development. Quantum computers promise in the long term to enable simulations of systems that are intractable to even the largest supercomputers [1,2]. Currently, scientists have access to so-called noisy intermediate-scale quantum (NISQ) computers [3], that present limited qubit counts without error correction. While applications of error-corrected quantum computers are well established, use cases where NISQ devices might achieve quantum advantage are still elusive [4]. In the search for these early applications, the problem must fit the hardware, and the hardware must enable implementation with minimal overheads.

Application-specific integrated chips (ASICs) are highly specialized processors optimized for specific problems when execution speed, power efficiency, or miniaturization is of utmost importance [5]. A prominent example where computational speed and energy efficiency are optimised through the use of ASICs is training of artificial neural networks using tensor processing units [6,7]. Building a general-purpose quantum computer capable of rivaling the most powerful classical computers has proven to be a difficult task, so it is likely that the first devices reaching useful quantum advantage will use quantum ASICs, also called Co-Design quantum computers.

A good example of a problem with suitable structure for simulation by quantum computers is nanoscale nuclear magnetic resonance (NMR) [8]. The problem can be described by a number of mutually interacting spins, which natively map to the qubits of a quantum computer, thereby circumventing the overheads in mapping the problem to qubits, such as in the case of fermions [9].

In general, fast and reliable quantum simulations of interacting spin systems would improve the interpretability of solid-state NMR and electron spin resonance (ESR) spectra, where advanced numerical techniques present very limited performance [10]. This shows the potential of quantum computers with a moderate number of qubits to shed light on the dynamics of these important systems. A Co-Design quantum computer that minimizes algorithm implementation overheads could be the first method to access these simulations. Note that, other NMR problems, such as zero-field NMR [11] and Hamiltonian learning [12], have already attracted research on how quantum computers can be used to tackle them and methods based on Bayesian computation [13] and generative models [14] have been developed for computing NMR spectra as well.

NMR techniques have a profound impact in research areas such as material science, chemistry, biology, and medicine [15]. Recently they have approached the nanoscale through solid-state quantum sensors such as the nitrogen vacancy (NV) center in diamond [16]. This is a particularly powerful

---

*Both authors contributed equally to this work.

†Corresponding author: manuel.algaba@meetiqm.com
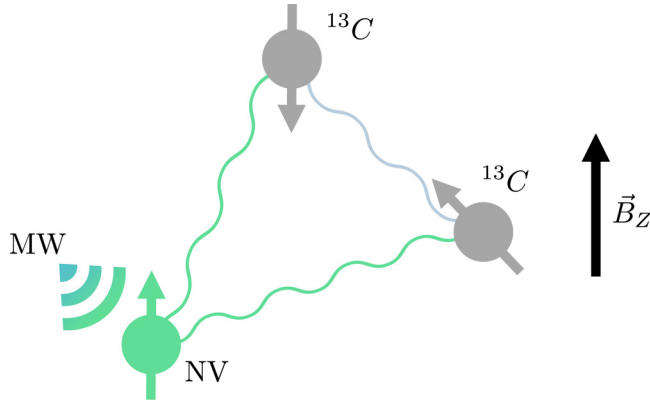
‡Corresponding author: mario.ponce@meetiqm.com

FIG. 1. NV center with a microwave drive interacting with two mutually interacting ^{13}C nuclei in a magnetic field $\vec{B}_Z$, corresponding to the Hamiltonian in Eq. (1) for $M = 1$ and $N = 2$.

quantum device, as it enables detection and control of nearby nuclear spins with nanoscale resolution [17]. Applications of the device are, e.g., the precise determination of the structure and dynamics of nuclear ensembles such as proteins [18], finding inter-label distances (via, e.g., Bayesian analysis of the NV response) in electronically labeled biomolecules [19], and the exploration of bespoke microwave (MW) sequences that efficiently transfer NV center polarization to the nuclear environment. Hyperpolarization (i.e., polarization beyond that of a thermal state in a magnetic field) of nuclear spins in diamond presents the potential to develop new and safer contrast agents for magnetic resonance imaging. This problem, which we aim to address through simulation by a quantum computer, could lead to improved detection of different malformations in tissues—such as heart or brain—without the need to deliver ionizing radiation, in contrast to other techniques [20].

This manuscript describes a Co-Design process for a quantum chip able to efficiently simulate nanoscale NMR scenarios. It is structured in three main parts, each of which is a crucial step in the Co-Design process: (1) identifying the problem (Sec. II), which here is simulating a nanoscale NMR system for hyperpolarizing nuclear spins; (2) choosing an algorithm for the nanoscale NMR problem and showing that a star-topology chip implements it with minimal overhead (Sec. III); and (3) Co-Designing the corresponding quantum chip using a central resonator bus (Sec. IV). The sections are followed by results and discussions (Sec. V) and an outlook (Sec. VI).

II. NANOSCALE NMR: HYPERPOLARIZATION

Let us consider a system consisting of M nitrogen-vacancy (NV) centers and N carbon-13 isotopes in the presence of a driving field and an external magnetic field $\vec{B}_Z$. NV centers and nuclei are all effectively described as spin-1/2 systems. The representation of such a system for $M = 1$, $N = 2$ is shown in Fig. 1. For simplicity, we consider the NV centers aligned with the external magnetic field, leading to the fol-

lowing Hamiltonian:

$$
\begin{aligned}
H = \sum_{j=1}^{M} \delta_j \sigma_j^z &- \sum_{k=1}^{N} \vec{\omega}_k^c \cdot \vec{I}_k + \sum_{j=1}^{M} \sum_{k=1}^{N} \frac{\sigma_j^z}{2} \vec{A}_{jk} \cdot \vec{I}_k \\
&+ \sum_{k>k'}^{N} g_{k'k} \left[I_{k'}^z I_k^z - \frac{1}{4}(I_{k'}^+ I_k^- + I_{k'}^- I_k^+) \right] \\
&+ \sum_{j>j'}^{M} h_{j'j} \left[\sigma_{j'}^z \sigma_j^z - 2(\sigma_{j'}^+ \sigma_j^- + \sigma_{j'}^- \sigma_j^+) \right] + H_{\mathrm{dr}}. \quad (1)
\end{aligned}
$$

In Eq. (1), we find the spin operators in the joint Hilbert space $\mathbb{C}^{2^{(M+N)}}$ of NV centers and nuclei:

$$
\sigma_j^\mu = \underbrace{\mathbb{1} \otimes \cdots \otimes \mathbb{1} \otimes \overbrace{\sigma_\mu}^{j^{\mathrm{th}}\mathrm{pos.}} \otimes \mathbb{1} \otimes \cdots \otimes \mathbb{1}}_{M \text{ factors}} \otimes \underbrace{\mathbb{1} \otimes \ldots \otimes \mathbb{1}}_{N \text{ factors}},
$$

$$
I_k^\mu = \underbrace{\mathbb{1} \otimes \ldots \otimes \mathbb{1}}_{M \text{ factors}} \otimes \underbrace{\mathbb{1} \otimes \cdots \otimes \mathbb{1} \otimes \overbrace{\frac{1}{2}\sigma_\mu}^{(M+k)^{\mathrm{th}}\mathrm{pos.}} \otimes \mathbb{1} \otimes \cdots \otimes \mathbb{1}}_{N \text{ factors}},
$$

where $(\sigma_\mu)_{2\times 2}$, $\mu \in \{x, y, z\}$ is the corresponding 2×2 Pauli matrix on the jth NV center and the kth nucleus respectively, and $\mathbb{1}$ is the 2×2 identity matrix. Accordingly, $\sigma_j^\pm = \frac{\sigma_j^x \pm i\sigma_j^y}{2}$ ($I_k^\pm = I_k^x \pm iI_k^y$) are the jth NV center (kth nucleus) ladder operators. The term δ_j is the detuning of the jth NV center with respect to the microwave drive H_{dr}. The hyperfine coupling vector $\vec{A}_{jk}$ represents the coupling between the jth NV center and the kth nucleus, while $\vec{\omega}_k^c = \gamma_c \vec{B}_Z - \frac{1}{2}\sum_{j=1}^{M} \vec{A}_{jk}$ is the modified Larmor frequency of the kth nucleus with the ^{13}C gyromagnetic ratio $\gamma_c \approx (2\pi) \times 10.7$ MHz/T, $g_{k'k}$ is the coupling between the kth and k'th nuclei, and $h_{j'j}$ is the coupling between the jth and j'th NV centers.

Note that, Eq. (1) is expressed in a rotating frame with respect to the free NV Hamiltonian, while H_{dr} represents an external driving tuned near resonance with a certain NV energy transition. The derivation of Eq. (1) can be found in Appendix A.

In order to hyperpolarize a diamond sample at room temperature, the NV centers are first optically polarized employing laser light, and then their state is transferred to the surrounding nuclei with the aid of a tailored microwave radiation scheme. The initial state of the nuclei in a room-temperature sample is well described by a fully mixed state due to the small energy splitting of the nuclear spins. By reinitializing the NV centers and repeating this procedure, the polarization transferred into the sample can be amplified. In this paper, we will consider the quantum simulation of the polarization transfer mechanism and study two different driving schemes acting on the NV centers in a room-temperature diamond.

The first driving scheme is a continuous driving whose Hamiltonian in the rotating frame mentioned earlier is $H_{\mathrm{dr}} = \frac{\Omega}{2}\sigma^\phi$, where $\sigma^\phi = e^{-i\phi}|1\rangle\langle 0| + e^{i\phi}|0\rangle\langle 1| = e^{-i\phi}\sigma^- + e^{i\phi}\sigma^+$, ϕ a phase, Ω the Rabi frequency and the kets $|1\rangle$ and $|0\rangle$ are the eigenvectors of the operator σ_z with eigenvalues ± 1 respectively. The set $\{|0\rangle, |1\rangle\}$ is called the computational basis of the state space of a two level system, and will be

our standard choice for a basis, $|0\rangle \equiv (1, 0)^t$ and $|1\rangle \equiv (0, 1)^t$. NV-nucleus polarization transfer is achieved when the Rabi frequency matches the modified nuclear Larmor frequency (i.e., when $\Omega = |\vec{\omega}_c|$), leading to the Hartmann-Hahn double-resonance condition [21]. For a single NV center and nucleus, the Hamiltonian in Eq. (1) reduces, in an interaction picture, to $H_I = \frac{A^\perp}{4}(|+\rangle\langle-|I^+ + |-\rangle\langle+|I^-)$, where $\pm\rangle| = 0\rangle| \pm 1\rangle|$, which shows a polarization transfer mechanism with the effective transfer rate $\frac{A^\perp}{4}$ (a detailed derivation can be found in Appendix B).

The second type of driving we consider is a pulsed-driving scheme, $H_{\rm dr} = \frac{\Omega(t)}{2}\sigma^\phi$, where $\Omega(t)$ is a train of π-pulses, such as the Carr-Purcell-Meiboom-Gill sequence [22,23] or the XY8 sequence [24,25]. We consider pulses with a negligible width compared to the time spacing τ between the π-pulses. If τ is selected such that $\tau = \frac{n\pi}{|\vec{\omega}^c|}$ (n being an arbitrary integer number) and the pulses are evenly spaced one finds that, in an interaction picture, for a single nucleus and NV center, the Hamiltonian reduces to $H_I = \alpha A^\perp \sigma_z I_x$, where α is a factor that depends on the integer n (see Appendix B). A phase imprinted on the pulse sequence through a time delay turns the interaction into $H_I = \alpha A^\perp \sigma_z I_y$. By combining both sequences with the appropriate rotations over the NV center, the polarization transfer interaction $H_I = -\frac{\alpha A^\perp}{4}(\sigma^+ I^- + \sigma^- I^+)$ is achieved (see Appendix B and Ref. [26] for more details).

Regarding common error sources, NV centers located at different positions in the diamond lattice experience stress conditions that lead to local energy deviations from the zero-field splitting. The corresponding term in Eq. (1) is the detuning δ_j. Another common type of imperfection appears due to unavoidable fluctuations of the Rabi frequency of the driving. This fluctuation can be modelled as an Ornstein-Uhlenbeck (OU) process [27], which has been shown to be an accurate description for NV centers [28]. It is a Gaussian process of the following form [29]:

$$X(t + \Delta t) = X(t)\,e^{-\Delta t/\tau} + \left[\frac{c\tau}{2}(1 - e^{-2\Delta t/\tau})\right]^{1/2} N(t), \quad (2)$$

where Δt is the time step, τ the correlation time, c the diffusion constant of the process, and $N(t)$ a temporally uncorrelated normally distributed random variable. It is a dimensionless term, which yields an effective Rabi frequency of $(1 + X)\Omega$. Neither of the system error types lead to considerable overheads in a simulation on a quantum computer. Finally, ^{13}C nuclear spin decay is not a relevant error source on the time scale of the protocol, since it is of the order of seconds [30], while the hyperpolarization process operates in the order of microseconds.

III. CO-DESIGN ALGORITHM

In this section, we provide an in-depth description of our Co-Design algorithm, starting with the choice of a simulation technique, followed by a short listing of hardware assumptions related to the allowed qubit operations (gates and resets), as well as the noise and errors present in the physical NMR system and in the quantum computer. Subsequently, the algorithm components are introduced. We end the section with

a discussion on layout and gate-level optimization. The high-level structure of the simulation protocol is shown in Fig. 2(a).

A. Simulation technique

The best established digital quantum simulation technique is based on decomposing the time-evolution operator into single-qubit and two-qubit gates through the Lie-Trotter-Suzuki formula [31], known as Trotterization. To simulate our problem on a quantum computer, we base our strategy on regular Trotterization [2] but we also explore the randomized Trotterization method qDRIFT [32] in Appendix C. Other, more NISQ-specific, simulation techniques such as the variational quantum simulator [33], the quantum assisted simulator [34], numerical quantum circuit synthesis [35], and a plethora of other quantum algorithms [4] can also be used as simulation methods.

One advantage of Trotterization over some of these NISQ methods is that it closely follows the real time evolution for each time step. This is particularly important for pulsed-driving schemes, where the free evolution in between different pulses always starts with a different initial state. Variational and quantum assisted methods would then require that each interpulse evolution is solved independently, making them impractical for the problem.

A second advantage of Trotterization is that its complexity and precision are straightforward to analyze. The Trotterization procedure can also be expanded to higher orders, and symmetrized expansions converge more rapidly and reduce the error with respect to the continuum time limit [36].

B. Hardware assumptions

1. Native gates

The hardware for the quantum simulation plays a major role in choosing the optimal quantum algorithm and its specific implementation. In our case, we consider a quantum computer based on superconducting qubits with the following native single-qubit gate set:

$$R_{xy}(\phi, \theta) = e^{-i(\cos\phi X + \sin\phi Y)\frac{\theta}{2}} \text{ and} \quad (3)$$

$$R_z(\theta) = e^{-iZ\frac{\theta}{2}}, \quad (4)$$

where X, Y, and Z are Pauli operators on the superconducting transmon qubits. The $R_{xy}(\phi, \theta)$ can physically be implemented through a microwave drive [37]. The gate $R_z(\theta)$ on the other hand does not need to be implemented directly, but can be performed virtually by tuning the phase of the subsequent gates applied on the qubit [38]. This reduces the number of single-qubit gates (SQGs) that need to be implemented.

The native two-qubit gate (TQG) that arises from the superconducting system Hamiltonian shown in Sec. IV and Appendix G, is a continuously parameterized controlled-Z (CZ) interaction [39], which can be transformed through local virtual R_z rotations into the form of a ZZ interaction:

$$U_{\rm ZZ}(\phi) = \begin{pmatrix} e^{-i\phi} & 0 & 0 & 0 \\ 0 & e^{i\phi} & 0 & 0 \\ 0 & 0 & e^{i\phi} & 0 \\ 0 & 0 & 0 & e^{-i\phi} \end{pmatrix}. \quad (5)$$
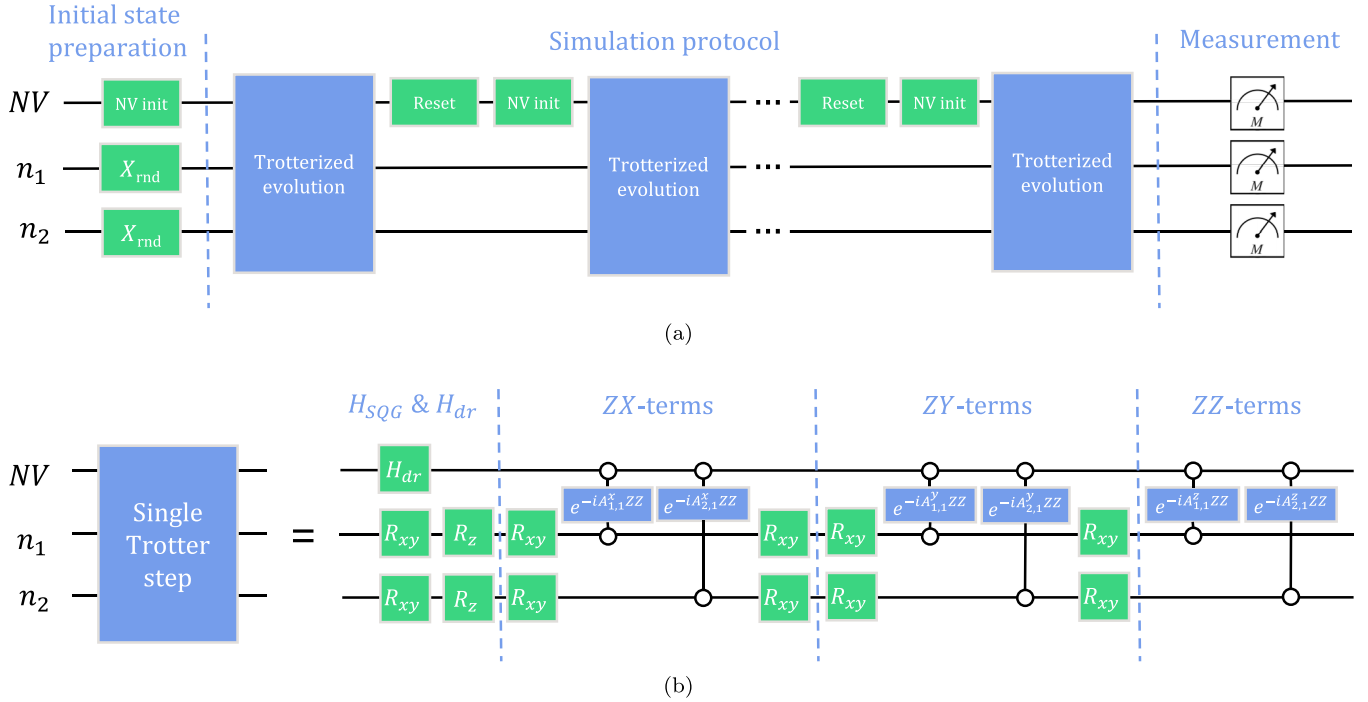
FIG. 2. (a) Sketch of the overall operation of the simulation algorithm for one NV center and two nuclei, with continuous driving; (b) corresponding gate sequence of one Trotter step on a star-topology chip for noninteracting nuclei. H_{SQG} refers to the single-qubit-gate component of the Hamiltonian, $A_{j,1}^{x,y,z}$ parameters are the various coupling strengths of the simulated system, and the X_{rnd} gates refer to X gates applied with a 50% probability to prepare an effective fully-mixed state. The initial state preparation can also be performed using the alternative random-phase approximation-inspired method. NV init is an initial-state preparation using single-qubit gates to the state required by the driving scheme. Details of the circuit components can be found in Appendix D along with a figure representing the pulsed driving case.

Even though the ZZ-interaction and the controlled-Z interactions appear different, their physical implementation is identical since they are related through virtual R_z rotations which come at no additional cost. Section IV goes into more depth on the two-qubit-gate implementation on our Co-Design quantum chip.

2. Qubit reset

In the hyperpolarization process, the state of the NV needs to be re-initialized after each cycle. It is therefore necessary to be able to reset the state of the qubit representing the NV center in the quantum computer. A qubit reset operation can be defined by two Kraus operators:

$$K_1^{reset} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad K_2^{reset} = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}. \tag{6}$$

On superconducting hardware this can be realized through connecting a quantum circuit refrigerator (QCR) to each circuit element that needs to be reset [40–43]. Different reset schemes are discussed in Sec. IV B.

3. Noise and errors

In this paper, we show that the simulation can tolerate the noise of the quantum processing unit (QPU), and that the simulation does not require large overheads to implement imperfections present in the nanoscale-NMR system, as

discussed in Sec. II. We will refer by *system imperfections* to effects in the nanoscale NMR system only, while the QPU is affected by *noise*, referring to the effect of the environment on the qubits, and *errors*, referring to inaccuracies of gates.

In our simulation of the algorithm, we use the most common noise models for superconducting transmon qubits [37], namely, an amplitude damping channel modelled by the Kraus operators:

$$K_1^{amp}(t) = |0\rangle\langle 0| + \sqrt{1-p(t)}|1\rangle\langle 1| = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-p(t)} \end{pmatrix},$$

$$K_2^{amp}(t) = \sqrt{p(t)}|0\rangle\langle 1| = \begin{pmatrix} 0 & \sqrt{p(t)} \\ 0 & 0 \end{pmatrix}, \tag{7}$$

with $p(t) = 1 - \exp(-t/T_1)$ and $T_1 = 60\,\mu s$, and a pure dephasing channel represented by the Kraus operators:

$$K_1^{deph}(t) = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-p(t)} \end{pmatrix}, \quad K_2^{deph}(t) = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{p(t)} \end{pmatrix}, \tag{8}$$

with $p(t) = 1 - \exp(-\Gamma(t))$ and $\Gamma(t)$ given by the expression $\Gamma(t) = \frac{t^2}{2}\int_0^\infty d\omega I(\omega) \coth(\frac{\beta\omega}{2}) \text{sinc}^2(\frac{\omega t}{2})$ where β is the inverse temperature of the environment. We chose the spectral function $I(\omega)$ to be of the type $1/f$ [37], and $T_2 = 60\,\mu s$. Additionally, each gate operation is assumed to be calibrated up to a two-qubit-gate (TQG) error $\varepsilon_{TQG} \in [10^{-4}, 10^{-2}]$, with the induced effective noise modelled by a depolarizing channel

defined for single-qubit gates by the Kraus operators:

$$K_1^{\text{depol}} = \sqrt{1-p}\, I,$$
$$K_2^{\text{depol}} = \sqrt{p/3}\, X,$$
$$K_3^{\text{depol}} = \sqrt{p/3}\, Y, \qquad (9)$$
$$K_4^{\text{depol}} = \sqrt{p/3}\, Z,$$

and for two-qubit gates by an analogous expression with the tensor products of two Pauli matrices and the coefficients $\sqrt{1-p}$ for the identity and $\sqrt{p/15}$ for the other operators. Single-qubit-gate (SQG) errors ε_{SQG} are assumed to be one order of magnitude lower than TQG errors.

C. Algorithm components

Our simulation of the nanoscale NMR problem follows the general structure shown in Fig. 2(a). It starts by initializing the states of all qubits, according to whether they represent a nucleus or a NV center, then evolving them using Trotter steps, followed by reset and re-initialization of the qubits representing NV centers. The cycle of time evolution and re-initialization is then repeated as many times as the protocol calls for. Finally the qubits are measured, and the polarization of the NV centers and nuclei are extracted as the expectation values of the qubit representing each element. Figure 2(a) shows the circuits for the case of continuous driving, while the details of pulsed driving schemes are shown in Fig. 10(a) in Appendix B. In the following, we go through these steps in more detail for the case of a single NV center.

1. Initial state preparation

To enable the polarization transfer, it is necessary to prepare the NV center in a specific initial state that depends on the driving scheme. For the continuous-driving scheme, it is the $|+\rangle$ or $|-\rangle$ state, and for the pulsed-driving scheme, it is one of the two computational basis states, $|0\rangle$ or $|1\rangle$.

For a diamond at room temperature, the initial state of the nuclear spins is well described by a fully mixed state $\rho_{\text{mixed}} = \frac{\mathbb{1}^{\otimes N}}{2^N}$, where $\mathbb{1}^{\otimes N}$ is the $2^N \times 2^N$ identity matrix. The state can be approximated by running the algorithm several times, each time with a different initial state obtained by applying X gates randomly on the qubits representing nuclei. A faster alternative to this sampling is the random-phase-approximation-inspired method, described in Ref. [44], and introduced into quantum computing in Ref. [45]. In this method, the qubits are all prepared in an equal superposition by applying Hadamard gates, and then the phases are randomized through the application of random phase gates. The method effectively reduces the prefactor in the scaling of the sampling error [45].

2. Time evolution

We choose to implement the time evolution generated by the Hamiltonian in Eq. (1) through Trotterization. For that, the Hamiltonian is rewritten in terms of qubit Pauli operators and arranged into noncommuting terms for an optimal Trotter splitting. The resulting circuit, which performs one Trotter step of the evolution in the continuous driving case, is de-

picted in Fig. 2(b). It consists of a set of initial single-qubit gates, including the ones corresponding to the driving and the detuning of the NV center, followed by three two-qubit gates per nucleus. There are three types of interaction terms, of the form XZ, YZ and ZZ, when no internuclear interactions are considered. With interactions there are a total of five interaction terms. Our native gate set only includes one type of two-qubit interaction as explained in Sec. III B 1. Therefore some SQGs need to be applied in order to convert the interaction terms into the right form, as discussed in Appendix D. Under specific circumstances, some TQGs can be removed by rotating the Hamiltonian into a more suitable basis as shown in Appendix E.

3. Cycles and reset

The dynamics of the system is known to produce an exchange of polarization between the NV center and the nuclei. This exchange is oscillatory, and therefore choosing a proper stopping time is important in order to achieve an effective polarization transfer from the NV center to the nuclei. In practice, a suboptimal transfer time can suffice, and the protocol is then repeated several times by resetting the NV center to its initial state and letting the system evolve under the drive again. Due to the re-initializations the full evolution of the system is nonunitary and a net gain of polarization of the system is enabled.

This structure is represented in the quantum circuit in Fig. 2(a) by the repeated Trotter evolution, followed by reset operations on the qubit representing the NV center, and a single-qubit gate to prepare the initial state of the driving protocol.

D. Layout optimization

When implementing a quantum algorithm on a superconducting QPU, the planar qubit connectivity forces us to solve the qubit-routing problem by introducing additional SWAP gates to connect distant qubits. In this section, we study the advantages of an optimized chip topology, a star topology, over a square-grid array of qubits in terms of reducing the number of SWAP gates that must be inserted to run the algorithm in Fig. 2 on the device. Different topologies will imply different counts of SWAPs added on top of the gates arising from the algorithm itself, as shown in Fig. 3. On a NISQ device, this implies different computational precision for the same gate error magnitudes. We choose the SWAP count as our metric to compare different topologies, as commonly gates have fidelities limited by calibration. The errors could be due to crosstalk, leakage, or filtering causing disturbances to the control signals. Under this scenario we want to minimize the gate count. On the other hand, for a highly tuned up device whose gates are limited by qubit coherence times, it would be optimal to minimize the circuit depth instead of the TQG count.

Assuming the gate errors are independent, the total error will be bounded by:

$$\varepsilon_{\text{gates}} = 1 - (1 - \varepsilon_{\text{TQG}})^{N_{\text{TQG}}} (1 - \varepsilon_{\text{SQG}})^{N_{\text{SQG}}}, \qquad (10)$$

where N_{TQG} is the number of two-qubit gates, N_{SQG} the number of single-qubit gates, and ε_{SQG} is the SQG error.
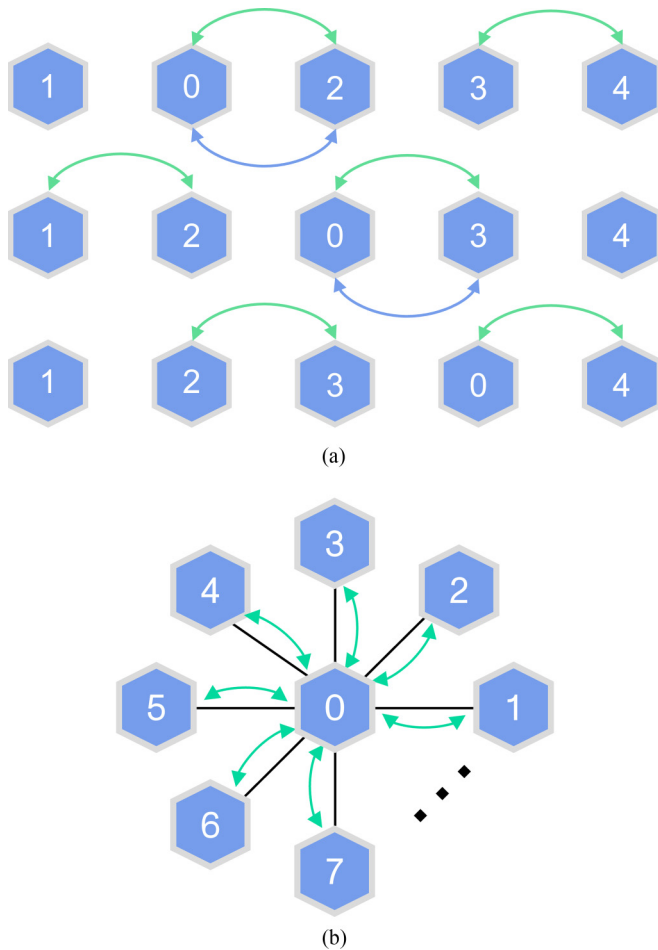
(a)

(b)

FIG. 3. (a) Three steps of the SWAP patterns in a five-qubit linear chain displayed from top to bottom. Green (blue) arrows represent the SWAP pattern for the case with (without) internuclear interactions. The green pattern is known as the "odd-even" SWAP pattern. The numbers are expressed according to the blue pattern, where label 0 represents the position of the NV center. (b) Star chip topology with the SWAP pattern for the interaction with internuclear interactions.

Consequently, reducing the gate count, especially N_{TQG}, has an exponential effect on the precision of the computations, underlining the effect of minimizing the SWAP gate overhead. As SWAP gates are not native to the hardware, but must be compiled out of three CZ gates, their effective error rate is also much higher than those of native gates.

1. Square grid

A common choice in superconducting quantum chips is the square grid of qubits. It has high connectivity and is suitable for performing the surface code error correction when scaled to large enough qubit counts with fast measurement and feedback [46]. The qubit routing problem on a square grid can be tackled using various numerical approaches [47–50]. However, these methods are inefficient. In our case, a tailored SWAP routing method, shown in Fig. 3(a), has been chosen and developed in Appendix F that can be shown to be well suited from two perspectives. First, a comparison against the cited numerical approaches (shown in Appendix F) reveals
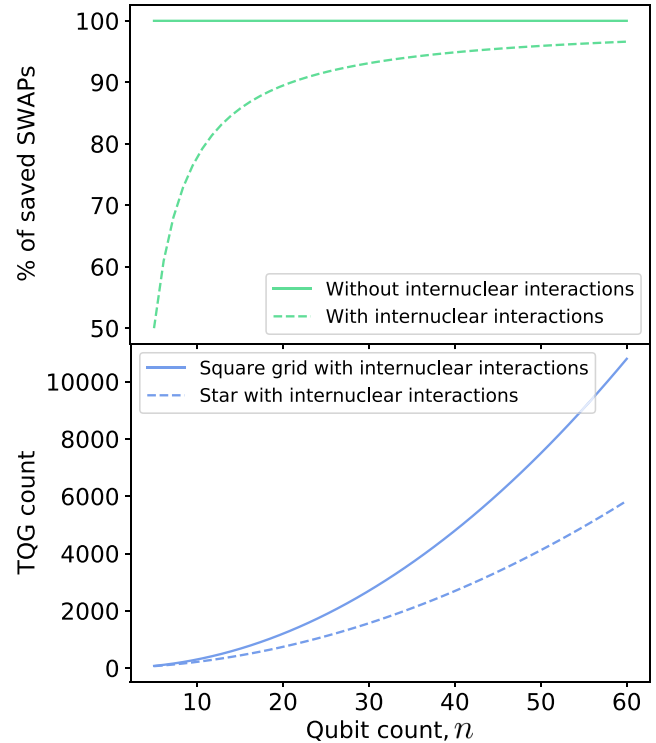

FIG. 4. The (top) panel shows the percentage of SWAP gates saved by using a star topology instead of a square grid for *n* qubits for the cases with and without internuclear interactions. The (bottom) panel shows the total TQG count against the qubit count in the interacting case for the square grid and the star architecture.

that our routing method is better in terms of number of gates. Second, it is completely deterministic and does not rely on expensive numerical optimization methods. It can also be shown not to be far from optimal: on a square grid each qubit has at most four nearest neighbors, implying that any SWAP operation provides at most three new neighbors. For an all-to-all (ATA) interacting Hamiltonian there are $n^2/2$ interactions, to leading order, for a simulation performed on n qubits (corresponding to N nuclei and one NV center). This implies a lower bound of at least $n^2/6$ SWAPs for any SWAP pattern on the square grid topology. Our SWAP pattern with $n^2/2$ SWAPs, discussed in Appendix F, is thus not far from optimal.

2. Star architecture

A star topology allows to implement the simulation of the simplified case without internuclear interactions directly, without any SWAP gates. With internuclear interactions considered, we still find a reduction in SWAP gates as compared to the square grid topology, as shown in Fig. 3(b). This reduction comes from the SWAP routing we implement, that consists of making the qubit 0 in Fig. 3(b) interact with all the external qubits and then swap its state with that of qubit 1 and repeat this process until all interactions have been performed. This allows us to use only $n - 1$ SWAP gates. The percentage of SWAP gates that can be saved can be observed in Fig. 4.

However, this improvement in the number of gates comes with a price to pay in the depth of the algorithm. We can only

TABLE I. Overheads introduced by the decomposition of $U_{ZZ}(\phi)$ gates into different examples of native TQGs in superconducting devices. The single-qubit-gate (SQG) count includes only R_{xy} rotations, as the R_z rotations can be implemented virtually.

	$U_{ZZ}(\phi)$	$U_{ZZ}(-\pi/4)$	$CZ(\pi)$	CNOT
TQGs	1	2	2	2
SQGs	0	5	3	1

do one TQG at a time in the star chip and we have $\frac{3}{2}n(n-1)$ TQGs from simulating the physical interactions and $3(n-2)$ TQGs from the SWAPs. This yields a depth for the TQGs of $\frac{3}{2}n^2 + \frac{3}{2}n - 6$ in a star chip, while for a square grid it is $6n$. Such depth increase comes from the reduction in parallelization, since all gates now act via the central qubit. On the other hand, less parallelization reduces the types of possible crosstalk errors. Adding connections between external qubits reduces the depth of the circuit, since the main cause of circuit depth is the fact that the interaction of two external qubits needs to be done exclusively by the central qubit. Further studies are required to see if the addition of more external layers to this topology (such as in a spider web) can lead to better compromises between depth and gate count, especially for simulating systems with clusters of strongly interacting nuclei.

E. Gate-level optimization

The two-qubit interactions that appear in the algorithm are the XZ, YZ and ZZ interactions, as shown in Sec. III C 2 and Fig. 2(b). When compiling the algorithm into the native gates of the device, all these interactions must be implemented in terms of some available gate set. We study in Table I the overhead introduced by decomposing these interactions into different examples of native TQGs of superconducting devices; namely, the parametrizable and fixed-phase U_{ZZ} gate, the fixed-phase controlled-Z gate CZ, and the CNOT gate. The CNOT gate is usually performed by making use of the cross-resonance gate [37,51], which introduces an U_{XZ} interaction, making it equivalent to the U_{ZZ} for the purpose of this algorithm. We assume that the SQGs that can be implemented are the R_{xy} and the R_z gates. These numbers can be further reduced if the first and last SQGs introduced by this compilation are combined with the adjacent SQGs in the algorithm.

The conclusion is that fixed-angle gates will double the number of TQGs that need to be physically performed. In Ref. [52], the improvements coming from the reduction of the gate count are compared to the new errors introduced by the interpolation of the calibrated phases. For two instances of a Quantum Approximate Optimization Algorithm (QAOA) [53], it is shown that the performance is better when using parametrized TQGs.

The gate sequences for some of the gate decompositions are shown in Fig. 5.

IV. CO-DESIGN HARDWARE

A star-architecture chip has fundamental scaling issues using a transmon as the central qubit as the number of neigh-
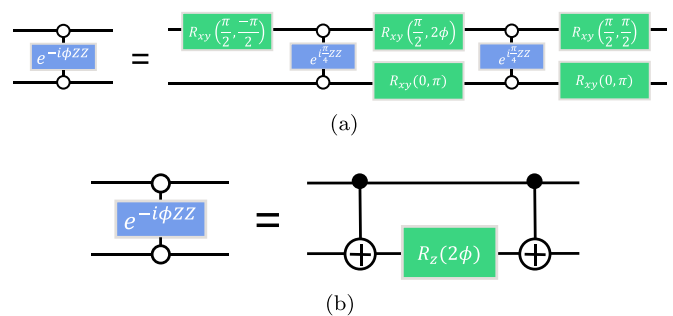

FIG. 5. (a) Gate decomposition of $e^{-i\phi ZZ}$ in terms of the fixed-phase $U_{ZZ}(\frac{\pi}{4})$ gate, (b) the *CNOT*.

bors grows. Every neighbor added to the center qubit would decrease its charging energy E_c. To keep the qubit frequency constant and anharmonicity in the transmon regime, the ratio of the qubit's Josephson energy to its charging energy, E_j/E_c, must remain unaffected. Therefore we cannot afford to change its charging energy. This leads to a trade-off between the number of coupled qubits and their coupling strength to the central element.

The spirit of Co-Design calls for replacing the central transmon with another object that enables this scaling in size. A resonator has no Josephson energy E_j, so the E_j/E_c ratio is not altered by adding more capacitive couplings to the resonator. Only small corrections to its frequency are introduced by adding coupled qubits. As a distributed element, a co-planar waveguide resonator also has physically more space for couplings than a central transmon qubit. By elongating the resonator and choosing the mode with the target frequency, the number of qubits coupled to it can further be increased. These properties make a resonator a favourable component in the center of the chip.

In the device in Fig. 6(a), the qubits are capacitively coupled to the resonator via tunable couplers [39,54,55] in the proximity of a voltage maximum of a standing wave in the resonator. As the resonator is elongated, we must use higher harmonic excitations of the resonator to keep the frequency around the operational frequency of the qubits. Tunable couplers avoid the frequency crowding issues related to direct coupling [56,57], and the linear resonator has higher connectivity in the center than ring resonator structures with quasi-all-to-all connectivities [58].

A linear resonator cannot in general be used as a qubit, since a microwave drive on it will not only populate the $\{|0\rangle, |1\rangle\}$ subspace, but also higher excited states. However, the effective interactions mediated via the tuneable coupler in Fig. 6(a) are of the type $a^\dagger a Z$ and $(a+a^\dagger)X + (a-a^\dagger)Y$ where a and $a^\dagger$ [37] are the resonator creation and annihilation operators. These types of interactions conserve excitation number, so when at most one excitation is in the qubit-resonator system, the resonator cannot be populated beyond its first excited state through interaction with a qubit mediated a tuneable coupler. CZ and iSWAP gates between the resonator and a qubit can be performed using the two interactions, and the theory is developed more fully in Sec. IV A. Then, a resonator together with an external qubit can be used as an effective central qubit in the following way.
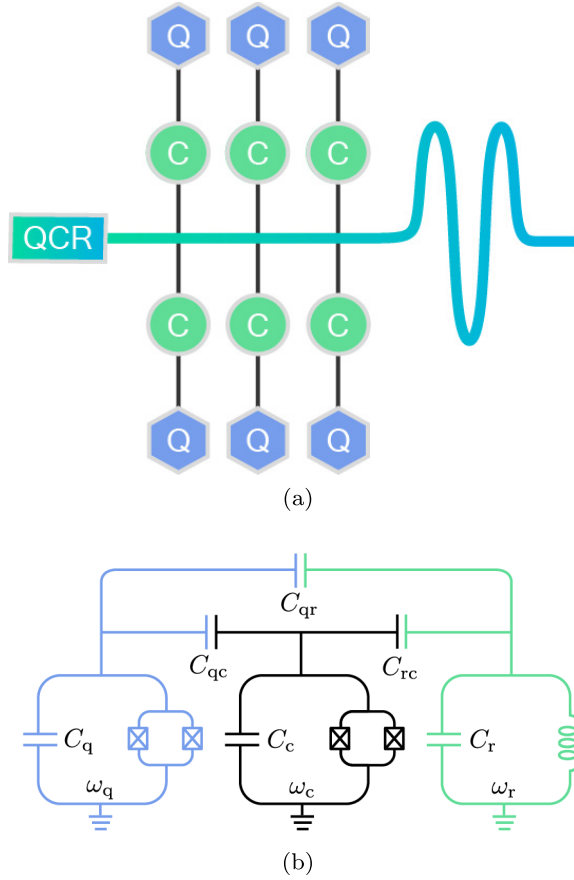
(a)

(b)

FIG. 6. (a) Central $\lambda/4$ resonator with 6 qubits coupled via tunable couplers. The resonator is also coupled to a quantum circuit refrigerator enabling fast reset. The device acts effectively as a six qubit star-architecture chip. (b) Electrical diagram of transmon qubit (left) coupled to a resonator mode (right) via a tunable coupler (center). The qubit has frequency ω_q, coupler ω_c, and resonator ω_r. The qubit and resonator have a direct capacitance C_{qr} and capacitances C_{qc} and C_{rc} respectively to the coupler.

1. Prepare all qubits and the resonator in their ground states

2. Select one qubit to form the effective central qubit together with the resonator

3. Prepare an arbitrary state in the selected qubit

4. Perform an iSWAP operation from the selected qubit to the resonator initially in the ground state

5. Perform CZ gates between the resonator and any other qubits

6. Perform an iSWAP operation back from the resonator to the selected qubit for measurement

The theoretically most straightforward protocol would be to perform a SWAP gate from the qubit to the resonator. The iSWAP, on the other hand, is a native gate that can directly be implemented on the hardware in Fig. 6(b). The iSWAP gate between the resonator and the qubit is represented by the unitary operator:

$$U_{\text{iSWAP}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & -i & 0 \\ 0 & -i & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (11)$$

TABLE II. Parameters of star-architecture chip.

Parameter	Symbol	Value
Resonator frequency	ω_r	$2\pi \times 4.3$ GHz
Qubit anharmonicity	α_q	$-2\pi \times 0.187$ GHz
Coupler anharmonicity	α_c	$-2\pi \times 0.110$ GHz
Resonator-coupler coupling	g_{rc}	$2\pi \times 98.5$ MHz
Qubit-coupler coupling	g_{qc}	$2\pi \times 101.8$ MHz
Resonator-qubit coupling	g_{rq}	$2\pi \times 8.9$ MHz
Resonator relaxation	T_1^r	$60 \mu s$
Qubit relaxation	T_1^q	$60 \mu s$
Coupler relaxation	T_1^c	$30 \mu s$
Resonator dephasing	T_2^r	$60 \mu s$
Qubit dephasing	T_2^q	$60 \mu s$
Coupler dephasing	T_2^c	$30 \mu s$

Since the CZ gates performing the computation following the iSWAP are diagonal in the computational basis, the phase introduced by the iSWAP is uninvolved in the gate. This enables substituting the SWAP gate by an iSWAP gate in the protocol to further minimize the gate count.

A. Gate theory and simulations

Here we demonstrate that in our star architecture CZ and iSWAP-type gates between any of the qubits and the $\{|0\rangle, |1\rangle\}$ subspace of a chosen resonator mode can be implemented. The operational principles of these gates are very similar to those between two qubits coupled with a tunable coupler [39,54,55,59]. The main limitation of our architecture (where one transmon is replaced by a resonator) is that iSWAP operations can only be performed in the zero- and single-excitation subspace of the two-qubit computational basis.

1. Conditional-Z gate

The CZ operation between the resonator and the qubit is described by the unitary operator:

$$\text{CZ}(\phi) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & e^{-i\phi} \end{pmatrix}. \quad (12)$$

This gate is equivalent to the $U_{ZZ}(\phi)$ gate in Eq. (5) up to two R_z rotations. To operate a CZ gate, we initialize the resonator-coupler-qubit set up shown in Fig. 6(b) at the idling configuration with zero effecting coupling between the qubit and resonator. Note that the coupler is also a transmon that shows a higher sensitivity to the magnetic flux than regular qubits. We next apply a flux pulse that lowers the coupler frequency, turning on the effective coupling between the resonator and the qubit. Depending on the flux pulse shape, the state collects conditional phase ϕ and possibly experiences population oscillations between computational and noncomputational states, as a function of the time spent at the gate-operation frequency. We optimize the pulse amplitude and duration such that after the flux pulse the CZ gate fidelity is maximized. Details of the gate theory can be found in Appendix G and the considered device parameters in Table II.
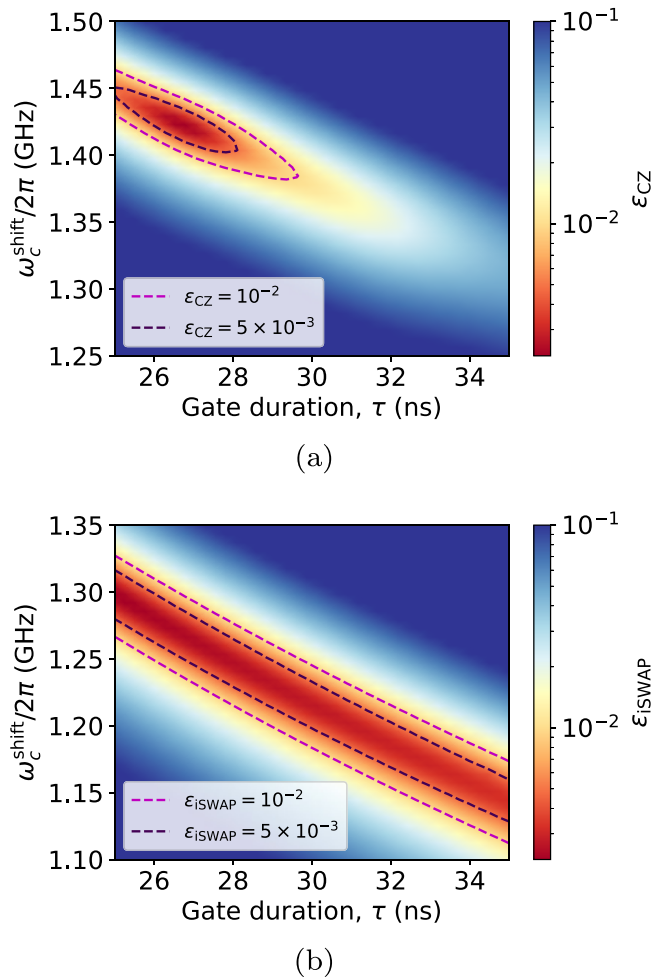
(a)

(b)

FIG. 7. (a) CZ gate error landscape averaged over random initial states. Contours with a low error are highlighted with a dashed line. (b) iSWAP gate error landscape obtained by averaging over a number of random initial states in the zero- and one-excitation manifolds. Both plots are produced using system parameters shown in Table II.

In Fig. 7(a), we operate our CZ gate by tuning the coupler frequency using a flattop Gaussian shaped flux pulse. The width of our Gaussian filter was fixed at 3 ns. Applying such a flux pulse to coupler results in a coupler frequency shift by ω_c^{shift} from the idling configuration. Then by appropriately tuning ω_c^{shift} and the gate time τ, one locates the optimal pulse configuration that minimizes the CZ(π) gate error $\varepsilon_{\text{CZ}} = 1 - (\text{tr}\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}})^2$, where σ is the target density matrix obtained after propagating some initial state $|\Psi\rangle\langle\Psi|$ with the ideal unitary of Eq. (12) and ρ the final density matrix obtained after propagating $|\Psi\rangle\langle\Psi|$ with the Lindbladian corresponding to our system defined in Eq. (G1). For our device parameters, the maximal decoherence limited CZ gate error averaged over a number of random initial states is 1.6×10^{-3}. Note that the system parameters in Table II were chosen such that they allow for the possibility to find a good idling configuration, where the residual CZ interaction vanishes before the gate operation. In our simulations, we have included environmental noise, such as amplitude damping and pure dephasing and

treated them using a Lindblad master equation solver in QUTIP [60,61].

2. iSWAP gate

Just as the CZ gate, the iSWAP gate can be natively realized in superconducting quantum computing architecture [37]. With our device, we can perform high-fidelity iSWAP gates between zero- and single-excitation computational states. The two-photon state $|1\rangle_r \otimes |1\rangle$, where $|1\rangle_r$ denotes the first excited state of the resonator, must be excluded because it resonantly interacts with the state $|2\rangle_r \otimes |0\rangle$ inducing a population exchange between the states. Hence the resulting operation in this subspace does not match the action of the targeted iSWAP operation.

The capacitive coupling between the elements of the electrical circuit shown in Fig. 6(b) gives rise to an effective XY interaction between the qubit and resonator under the rotating wave approximation. Such an interaction conserves excitation number. With only the qubit or resonator (or neither) initially populated, we stay within the single excitation subspace of the joint system, thereby minimizing leakage of quantum population into the higher excited states of the resonator. The XY interaction can be turned on by first tuning the qubit in resonance with the resonator, and then applying a flux-pulse to the coupler to turn on the coupling, similar to the CZ gate operation.

Figure 7(b) shows iSWAP gate error landscape for the same device parameters (given in Table II). The optimal average iSWAP gate error $\varepsilon_{\text{iSWAP}}$ obtained for our device is 1.7×10^{-3}. This result is obtained by averaging over a number of random initial states within the zero- and one-excitation manifolds.

The results of our two-qubit-gate simulations demonstrate that our star architecture supports operating gates with similar fidelities as regular transmon qubits coupled together. The increased local connectivity of the device reduces the need for SWAP gates to simulate the nanoscale NMR problem (and others with a similar structure) and consequently in the end improves simulation fidelities.

B. Reset

The hyperpolarization protocol described in Sec. II needs regular re-initializations of the state of the NV center. The Co-Design hardware for simulating the protocol must therefore support this operation within qubit lifetimes. This is a hardware challenge, but one with solutions in sight. In particular, the quantum circuit refrigerator (QCR) has been used to perform the reset in tens of nanoseconds [40–43], which is a similar timescale to gate operations. The advantage of using a QCR for the reset is the possibility to reset the central resonator directly, without the need transfer the resonator population back to the central qubit using an iSWAP gate. Alternatively, a fast reset is possible through applying a flux drive to a qubit to SWAP its state with its measurement line [62]. This scheme has the advantage of not requiring any additional hardware not already present on the chip, but comes with a small cost in the circuit depth, as the state of the resonator must be transported using an iSWAP gate into the
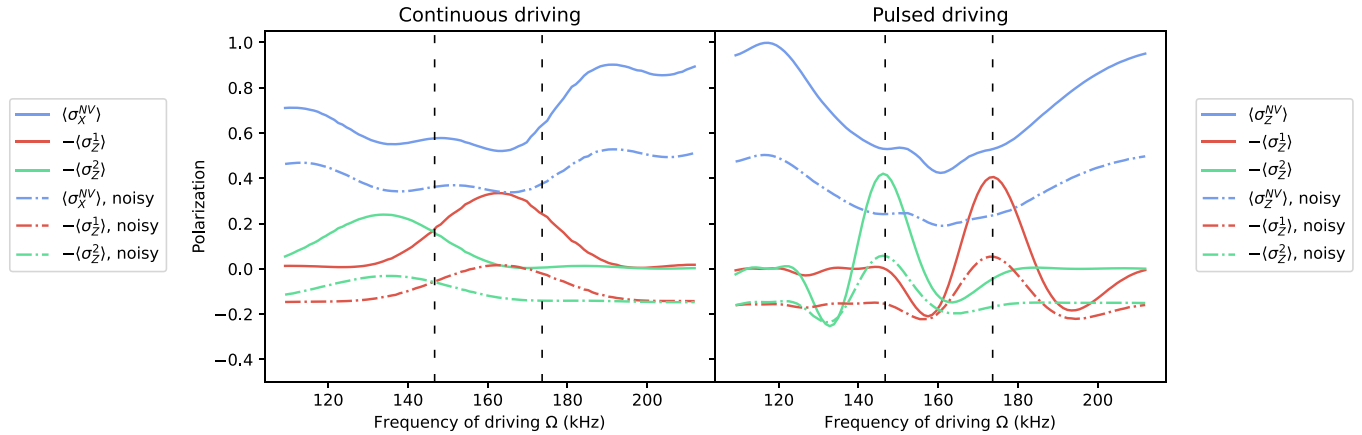
FIG. 8. Polarization transfer from one NV center to two interacting nuclei for a simulation time $t_f = 30\ \mu$s and a single cycle with $s = 32$ Trotter steps. The relevant observables to represent polarization are those included in the legends. (Left) with continuous driving and (right) with pulsed driving. Both plots depict with solid lines an exact simulation of the nanoscale NMR system including its usual imperfections, namely a detuning of $\delta_1 = 120$ kHz and a fluctuating microwave drive amplitude which follows an OU process with correlation time $\tau = 500\ \mu$s and diffusion constant $c = 4 \times 10^{-7}\ \mu$s^{-1} [see Eq. (2)]. The dotted lines show a simulation additionally including QPU noise and errors (as defined in Sec. III B 3). The dashed vertical black lines indicate the expected resonance frequencies of the nuclei, i.e., where the peaks should be centered in the absence of the detuning δ_1. The noisy QPU has $\varepsilon_{\text{SQG}} = 10^{-4}$, $\varepsilon_{\text{TQG}} = 2 \times 10^{-3}$ and amplitude damping and pure dephasing with $T_1 = 60\ \mu$s and $T_2 = 60\ \mu$s with gate durations of single- and two-qubit gates $\tau_{\text{SQG}} = 60$ ns and $\tau_{\text{TQG}} = 27$ ns, respectively. The parameters are chosen to demonstrate the performance of a state-of-the-art superconducting QPU. The characteristic effect of the detuning is to shift the curves to the left in frequency domain, as can be seen in (left). (Right) This effect is compensated by the pulsed driving, which refocuses detuning errors. The amplitude damping affecting the QPU shifts down the expectation value of all observables, while the (depolarizing) gate errors decrease the polarization transfer efficiency by reducing the visibility of the peaks.

designated central qubit and be re-initialized there. The reset timescale is also somewhat longer than when using a QCR.

V. RESULTS AND DISCUSSION

In this section, we discuss the two main results of the paper: namely the predicted performance of our proposed quantum algorithm on a regular noisy QPU, as well as the performance increase obtained with our proposed Co-Design QPU. To this aim, we will focus on the polarizations of the NV center and nuclear spins, that are relevant quantities of the problem and straightforward to measure in a quantum computer.

In Fig. 8, we compare the frequency response of the polarization transfer process on two different simulated devices: a QPU with realistic noise parameters, and an ideal noiseless QPU. We consider one NV center, two interacting nuclei and different driving frequencies for both continuous and pulsed driving schemes. In the simulation we ignore errors in the preparation of the fully mixed state of the qubits representing the nuclei. The blue curves show the remaining polarization in the NV center after one cycle of initialization and time evolution, while the red and the green curves correspond to the nuclear polarizations at the end of the cycle. For each nucleus there appears a resonance frequency in the system, for which the polarization transfer is optimal for said nucleus, depicted in the figure by the peaks of the curves.

Both simulations include the effects of the most common imperfections in nanoscale NMR systems, i.e., energy detunings and Rabi frequency fluctuations discussed in Sec. II. The simulation of the quantum algorithm additionally includes noise and gate errors present in the QPU. It is notable that

the noise affects the height and shape of the peaks more than their location.

The system imperfections include a detuning of 120 kHz of the NV center from the zero-field splitting that shifts the peaks in Fig. 8 (left) to frequencies lower than their predicted Larmor frequencies (dotted vertical black lines). Figure 8 (right) shows how the pulsed-driving scheme XY8 [24,25] acts as a robust dynamical decoupling sequence, eliminating such frequency shifts both in the ideal and noisy simulations.

Regarding the QPU noise and errors, the amplitude damping channel causes an overall shift down of all polarizations at all driving frequencies. Dephasing noise and gate errors (as modelled by depolarizing noise) cause the curves in Fig. 8 to flatten and lose contrast. While we have discussed how the product of gate errors is minimized by reducing the SWAP overhead through Co-Design hardware, the loss of contrast can also be addressed through error mitigation techniques such as zero-noise extrapolation [63–65]. Dephasing can also be reduced through dynamical decoupling techniques [37], thus extending the system coherence and increasing the effective T_2 time. The simulations presented in Fig. 8 include the decoherence times and gate fidelities that can be achieved with the hardware in Sec. IV. This implies an overestimation of the actual errors in the simulation, since the gate fidelities already include some decoherence.

To quantify the advantage of our Co-Design processor, Fig. 9 shows how the reduction in TQGs improves our ability to extract relevant information from the simulation. The figure compares the star-architecture chip to qubits connected on a square grid simulating a six qubit system with one NV center and five noninteracting nuclei. On the two chips we use SWAP patterns according to the schemes discussed in Sec. III D.
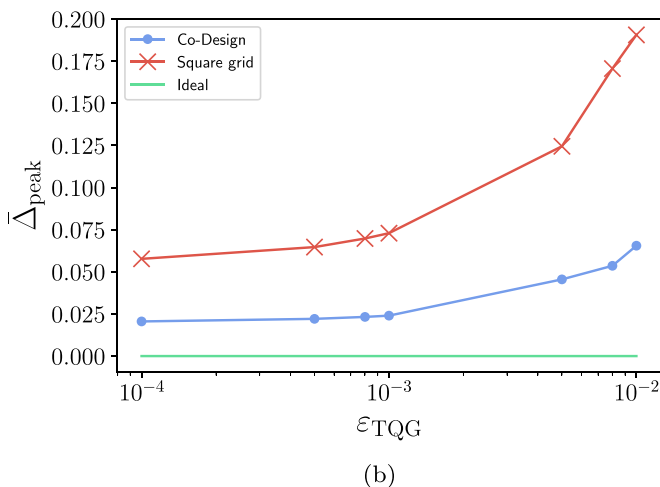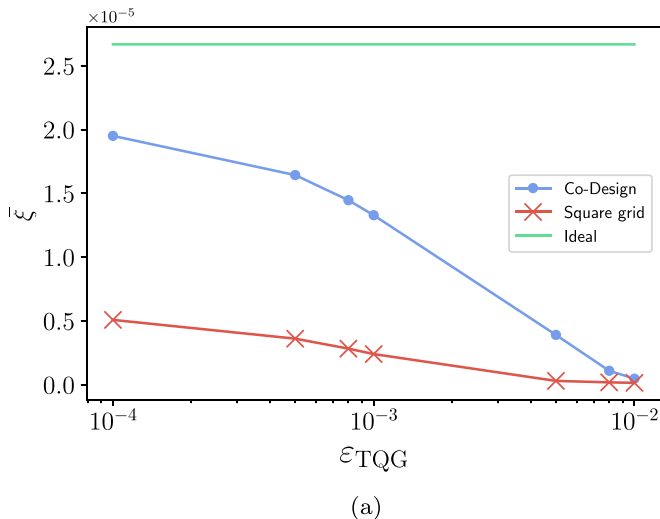
(a)

(b)

FIG. 9. Performance gain from Co-Design: a comparison between a Co-Design star-architecture against a square grid, taking as reference an ideal simulation without QPU noise. The comparison highlights the negative effect that the SWAP gates on the square grid have on extracting relevant information from the simulation. We consider two quantities: in (a), the ratio $\bar{\xi}$ between the height and the width of the polarization peaks, and in (b), the estimation error of the polarization peak center $\bar{\Delta}_{\text{peak}}$, where the bars denote an average over five nuclei for each noise level. The simulations were performed with the same parameters as in Fig. 8, except for the number of qubits, which has been increased from 2 to 5.

First, Fig. 9(a) shows the average height-to-width ratio $\bar{\xi}$ of the nuclear polarization peaks obtained with star and square grid topologies with respect to an ideal error-free simulation. It serves as an indicator of how much the QPU noise degrades the simulation for each case. The ratio $\bar{\xi}$ is computed by fitting a Gaussian function on each peak, and computing:

$$\bar{\xi} = \left\langle \frac{h}{\sigma} \right\rangle, \qquad (13)$$

where h is the height and σ the variance of the fitted Gaussian function, averaged over the five nuclei.

The curves for both topologies must coincide at $\bar{\xi} = 0$ for a maximal-error device, and at $\bar{\xi} = \bar{\xi}_{\text{ideal}}$ for an error-free quan-

tum computer, since for a maximal-error device the output is pure noise and for an error-free quantum computer the number of SWAPs is irrelevant to the precision. For NISQ devices in between these limits, a performance difference between the architectures is observed. For systems with more nuclei and NV centers, the differences between topologies start to appear at lower errors, since the number of total operations grows. This shows how the QPU topology is of great importance for the computational precision of NISQ devices, while for fault-tolerant quantum computers the precision is unaffected by the topology.

Second, Fig. 9(b) shows the average relative error in the central frequency of the NMR peaks:

$$\bar{\Delta}_{\text{peak}} = \left\langle \left| \frac{\omega_{\text{noisy}} - \omega_{\text{ideal}}}{\omega_{\text{ideal}}} \right| \right\rangle, \qquad (14)$$

where ω_{noisy} and ω_{ideal} are the peak-center frequencies extracted from the Gaussian fittings for the noisy and ideal cases, respectively. The peak centers correspond to driving frequencies that efficiently transfer polarization to different parts of the diamond lattice.

With the quantum simulation we can individually identify the nuclear resonance peaks by directly measuring the polarization of each qubit. This could enable exploration of how the polarization diffuses in the lattice with single-nucleus precision. In contrast, in a standard nanoscale NMR experiment, one typically only has only access to the excitation loss of the NV(and thus only to the average transmitted polarization). This demonstrates the advantage of simulating the system on a quantum computer, as a it provides access to the relevant microscopic details of the dynamics that are otherwise inaccessible.

The figures demonstrate that the Co-Design chip is able to detect the resonance frequencies and predict the peak heights better at all considered noise levels. The power of Co-Design is particularly evident in Fig. 9(b), where the square grid is shown to require two orders of magnitude lower noise levels to reach the same accuracy as the Co-Design chip.

VI. CONCLUSIONS AND OUTLOOK

We have presented a quantum algorithm to simulate a nanoscale NMR problem, namely a hyperpolarization protocol. We have simulated the proposed quantum algorithm with typical noise processes of a NISQ superconducting quantum computer with state-of-the-art parameters. We find that, despite considering a noisy QPU, our protocol still allows to identify the positions of the nuclear resonances (corresponding to the maximal polarizations) in the frequency domain, as well as the behavior in the vicinity of such resonant frequencies, thus enabling the exploration of optimized protocols and driving parameters to hyperpolarize the nuclear ensemble.

Moreover, we have shown that a specific Co-Design architecture adapted to the problem provides an advantage over general-purpose designs in the NISQ era, thanks to the reduction in two-qubit-gate count. Consequently, the adapted design reduces the necessary gate fidelities to solve practical problems in nanoscale NMR. This application-specific QPU consists of a central resonator, representing an NV center, coupled to a number of qubits representing the nuclei. The

design can be scaled to more NV centers and a potentially large number of qubits around them. This is an example of a shortcut to quantum advantage. Adapting more NISQ-friendly algorithm alternatives, such as those listed in [4], adapted to the problem and to the Co-Design hardware can provide further shortcuts.

Our work opens interesting directions for further investigation, since a quantum processor able to efficiently simulate nanoscale-NMR scenarios with a large number of nuclear spins would have a great impact on NMR-based applications. Fast and reliable quantum simulations of interacting spin systems would improve the interpretability of zero- and low-field NMR where spin-spin interactions become dominant [11], and nanoscale-NMR systems where a quantum sensor is strongly coupled via dipole-dipole interactions to nuclear or electron spin clusters. A possible application of the latter is the estimation of inter-label distances (via, e.g., Bayesian analysis of the NV center response) in electronically labeled biomolecules [19]. In this case, the numerical analysis of systems beyond two-electron spin labels in realistic conditions, including protein motion and decoherence channels, is already numerically challenging.

APPENDIX A: DERIVATION OF THE SYSTEM HAMILTONIAN

The Hamiltonian in Eq. (1) can be derived from first principles. Let us first assume a model including only two ^{13}C nuclei and one NV center (Fig. 1) with dipole-dipole interactions. For simplicity we also consider the NVs to be aligned with the external magnetic field. In that case, the Hamiltonian of the system reads:

$$H = DS_z^2 - \gamma_e B_z S_z - \gamma_c B_z (I_1^z + I_2^z) + \sum_{k=1}^{2} \frac{\hbar \mu_0 \gamma_e \gamma_c}{2|\vec{r}_k|^3} \left[\vec{S} \cdot \vec{I}_k - \frac{3(\vec{S} \cdot \vec{r}_k)(\vec{I}_k \cdot \vec{r}_k)}{|\vec{r}_k|^2} \right] + \frac{\hbar \mu_0 \gamma_c^2}{2|\vec{r}_{1,2}|^3} \left[\vec{I}_1 \cdot \vec{I}_2 - \frac{3(\vec{I}_1 \cdot \vec{r}_{1,2})(\vec{I}_2 \cdot \vec{r}_{1,2})}{|\vec{r}_{1,2}|^2} \right],$$

(A1)

where S_j is the jth spin component of the NV center, I_k^j the jth spin component of nucleus k, D is the zero-field splitting of the NV center, γ_e and γ_c are the gyromagnetic factors of the NV center and the nuclei respectively, B_z is the external magnetic field, which is aligned with the symmetry axis of the NV center $\vec{r}_k$ is the relative position vector between the NV center and nucleus k and $\vec{r}_{1,2}$ is the relative position vector between both nuclei.

We go into an interaction picture with respect to $H_0 = DS_z^2 - \gamma_e B_z S_z$. The NV-nuclei interaction term reads:

$$H_{\text{NV}-\text{N}}^I = \sum_{k=1}^{2} \frac{\hbar \mu_0 \gamma_e \gamma_c}{2|\vec{r}_k|^3} \left\{ \left[S_z I_k^z - \frac{3(S_z r_k^z)(\vec{I}_k \cdot \vec{r}_k)}{|\vec{r}_k|^2} \right] + U_0^\dagger \left[S_x I_k^x + S_y I_k^y - \frac{3(S_x r_k^x + S_y r_k^y)(\vec{I}_k \cdot \vec{r}_k)}{|\vec{r}_k|^2} \right] U_0 \right\},$$

(A2)

where we split the expression in commuting and noncommuting operators. The noncommuting operators pick a fast-rotating phase and can be neglected through the rotating-wave approximation. By performing an interaction-picture transformation with respect to $H_0 = -\gamma_c B_z (I_1^z + I_2^z) = \omega(I_1^z + I_2^z)$, the nucleus-nucleus interaction term reads:

$$\begin{aligned}
H_{\text{N}-\text{N}}^I &= \frac{\hbar \mu_0 \gamma_c^2}{2|\vec{r}_{1,2}|^3} U_0^\dagger \left[\vec{I}_1 \cdot \vec{I}_2 - \frac{3(\vec{I}_1 \cdot \vec{r}_{1,2})(\vec{I}_2 \cdot \vec{r}_{1,2})}{|\vec{r}_{1,2}|^2} \right] U_0 \\
&= \frac{\hbar \mu_0 \gamma_c^2}{2|\vec{r}_{1,2}|^3} \left\{ I_1^z I_2^z + \frac{1}{2}(I_1^+ I_2^- + I_1^- I_2^+) \right. \\
&\quad \left. - \frac{3[I_1^+ e^{i\omega t}(r_{1,2}^x - i r_{1,2}^y) + I_1^- e^{-i\omega t}(r_{1,2}^x + i r_{1,2}^y)][I_2^+ e^{i\omega t}(r_{1,2}^x - i r_{1,2}^y) + I_2^- e^{-i\omega t}(r_{1,2}^x + i r_{1,2}^y)]}{4|\vec{r}_{1,2}|^2} \right\},
\end{aligned}$$

(A3)

with $I_k^\pm = I_k^x \pm i I_k^y$. Applying again the rotating-wave approximation and undoing the interaction picture, we finally arrive at:

$$H_I = -\gamma_c B_z (I_1^z + I_2^z) + S_z (\vec{A}_1 \cdot \vec{I}_1 + \vec{A}_2 \cdot \vec{I}_2) + g_{1,2} \left[I_1^z I_2^z - \frac{1}{4}(I_1^+ I_2^- + I_1^- I_2^+) \right],$$

(A4)

with $\vec{A}_k = \frac{\hbar \mu_0 \gamma_e \gamma_c}{2|\vec{r}_k|^3}[\hat{z} - \frac{3(\hat{z} \cdot \vec{r}_k)\vec{r}_k}{|\vec{r}_k|^2}]$, and $g_{1,2} = \frac{\hbar \mu_0 \gamma_c^2}{2|\vec{r}_{1,2}|^3}[1 - 3(\frac{r_{1,2}^z}{|\vec{r}_{1,2}|})^2]$.

We rewrite S_z in the $\{|0\rangle, |1\rangle\}$ subspace by dropping out the $|-1\rangle$ energy state as it will not participate in the dynamics. Leakage to that state would not be a problem because of the energy difference between states $|0\rangle, |1\rangle$ and $|0\rangle, |-1\rangle$. Then by

using that $|1\rangle\langle1| = \frac{\mathbb{1}-\sigma^z}{2}$ to we get:

$$H_I = -\vec{\omega}_1^c \cdot \vec{I}_1 - \vec{\omega}_2^c \cdot \vec{I}_2 + \frac{\sigma_z}{2}(\vec{A}_1 \cdot \vec{I}_1 + \vec{A}_2 \cdot \vec{I}_2) + g_{1,2}\left[I_1^z I_2^z - \frac{1}{4}(I_1^+ I_2^- + I_1^- I_2^+)\right], \tag{A5}$$

where $\vec{\omega}_k^c = -(\frac{A_k^x}{2}, \frac{A_j^y}{2}, \frac{A_j^z}{2} - \gamma_c B_z)$ is the modified nuclear Larmor term due to the presence of the NV center.

Generalizing equation (A5) to M NV centers and N nuclei, including the detuning of the NV centers and adding the microwave driving term we obtain precisely the Hamiltonian in Eq. (1).

APPENDIX B: HYPERPOLARIZATION SEQUENCES

1. Hartmann-Hahn sequence

Here we explain the dynamics induced by the continuous driving on the hyperpolarization protocol. To illustrate the mechanism, we consider a system including a single NV center and a single nucleus. The corresponding Hamiltonian, now including the driving term, reads:

$$H = DS_z^2 - \gamma_e B_z S_z - \gamma_c B_z I_z + S_z \vec{A} \cdot \vec{I} + S_x \sqrt{2}\,\Omega\cos(\omega t - \phi). \tag{B1}$$

In the interaction picture with respect to $DS_z^2 - \gamma_e B_z S_z$, we obtain:

$$H_I = -\gamma_n B_z I_z + S_z \vec{A} \cdot \vec{I} + \frac{\Omega}{2}(e^{ip_+ t}|1\rangle\langle0| + e^{ip_- t}|$$
$$- 1\rangle\langle0| + \text{H.c.})[e^{i(\omega t - \phi)} + e^{-i(\omega t - \phi)}], \tag{B2}$$

where $p_{+/-} = D \pm |\gamma_e|B_z$. Choosing the resonance condition $\omega = p_+$ and applying the rotating-wave approximation, we get:

$$H_I = -\gamma_n B_z I_z + S_z \vec{A} \cdot \vec{I} + \frac{\Omega}{2}(e^{i\phi}|1\rangle\langle0| + e^{-i\phi}|0\rangle\langle1|). \tag{B3}$$

Finally we can use the identity $|1\rangle\langle1| = \frac{\mathbb{1}-\sigma^z}{2}$ and the fact that there will be no transitions to the $|-1\rangle$ because of energy differences:

$$H_I = -\vec{\omega}^c \cdot \vec{I} - \frac{\sigma_z}{2}\vec{A} \cdot \vec{I} + \frac{\Omega}{2}\sigma^\phi, \tag{B4}$$

where $\sigma^\phi = e^{-i\phi}|1\rangle\langle0| + e^{i\phi}|0\rangle\langle1| = e^{-i\phi}\sigma^- + e^{i\phi}\sigma^+$ and $\vec{\omega}_n = -(\frac{A_x}{2}, \frac{A_y}{2}, \frac{A_z}{2} - \gamma_n B_z)$. More details about the different terms were discussed in the main text, in Sec. II. Choosing $\phi = 0$ and further moving to an interaction picture with respect to the terms $-\vec{\omega}^c \cdot \vec{I} + \frac{\Omega}{2}\sigma_x$, we obtain:

$$H_I = \frac{e^{i\frac{\Omega}{2}\sigma_x t}\sigma_z e^{-i\frac{\Omega}{2}\sigma_x t}}{2}e^{-i\vec{\omega}^c \cdot \vec{I}t}\vec{A} \cdot \vec{I}e^{i\vec{\omega}^c \cdot \vec{I}t}. \tag{B5}$$

We choose now $\Omega = |\vec{\omega}^c|$, leading to the so-called Hartmann-Hahn double-resonance condition. Applying the identity $e^{i\vec{I}\cdot\hat{l}\phi}\vec{I} \cdot \vec{b}e^{-i\vec{I}\cdot\hat{l}\phi} = \vec{I}[(\vec{b} - (\vec{b} \cdot \hat{l})\hat{l})\cos\phi - \hat{l} \times \vec{b}\sin\phi + (\vec{b} \cdot \hat{l})\hat{l}]$ and the rotating-wave approximation to remove time-dependent terms, we get the flip-flop Hamiltonian:

$$H_I = \frac{A^\perp}{4}(|+\rangle\langle-|I^+ + |-\rangle\langle+|I^-), \tag{B6}$$

with $A^\perp = |\vec{A}_x^\perp| = |\vec{A} - (\vec{A} \cdot \hat{\omega}^c)\hat{\omega}^c|$ and the nuclear coordinates changed so that $\hat{x} = \hat{A}_x^\perp$ and $\hat{z} = \hat{A}_z^\parallel$ with $\vec{A}_z^\parallel = (\vec{A} \cdot \hat{\omega}^c)\hat{\omega}^c$.

2. Pulsed sequence

Now we consider the pulsed case, represented by the driving term $H_{dr} = \frac{\Omega(t)}{2}\sigma^\phi$, where $\Omega(t)$ is a train of π-pulses. The Hamiltonian is already expressed in the interaction picture from Eq. (B4). From there, we further move into a rotating frame with respect to the driving term. The corresponding unitary transformation is $U_0 = (-i\sigma^\phi)^k$ for the time interval between pulses k and $k+1$. This leads to:

$$H_I = -\vec{\omega}^c \cdot \vec{I} + F(t)\frac{\sigma_z}{2}\vec{A} \cdot \vec{I}, \tag{B7}$$

where $F(t)$ is the so-called filter function, with value $+1$ when k is even, and -1 when k is odd, representing the sign of the operator σ_z, flipped by the action of each pulse.

It is necessary to apply two different patterns of pulses. The "symmetric case," meaning an evenly distributed sequence of pulses for which the filter function is even and can be expanded in Fourier series of cosines as:

$$F(t) = \sum_{n=1}^{\infty} f_n \cos\left(\frac{2\pi n}{T}t\right), \tag{B8}$$

with $f_n = 0$ when n is even and $f_n = -\frac{4}{\pi n}$ when n is odd, if the pulses are distributed such that the interpulse spacing is constant. We choose the resonance condition $T = \frac{2\pi n}{|\vec{\omega}^c|}$, where n is the harmonic number. This is the same resonance condition that we introduced in Sec. II, but here it is formulated with the period T that appears in the Fourier expansion, instead of with the interpulse spacing $\tau = \frac{T}{2}$ from before. Going to an interaction picture with respect to $-\vec{\omega}^c \cdot \vec{I}$ and repeating the procedure we used above in the Hartmann-Hahn case, we get:

$$H_I = \alpha A^\perp \sigma_z I_x, \tag{B9}$$

where $\alpha = f_n/4$.

With the second pattern of pulses, called the "asymmetric case," we apply an oddly-distributed sequence of pulses for which the filter function is odd and can be expanded in a Fourier series of sines. Note that this sequence of pulses is identical to the even sequence but shifted by a $\pi/2$ phase. An analogous derivation gives:

$$H_I = \beta A^\perp \sigma_z I_y, \tag{B10}$$

with $\beta = g_m/4$ and g_m coming from the Fourier expansion of sines, analogously to f_n.

Combining these two patterns one can generate an effective Hamiltonian of the form:

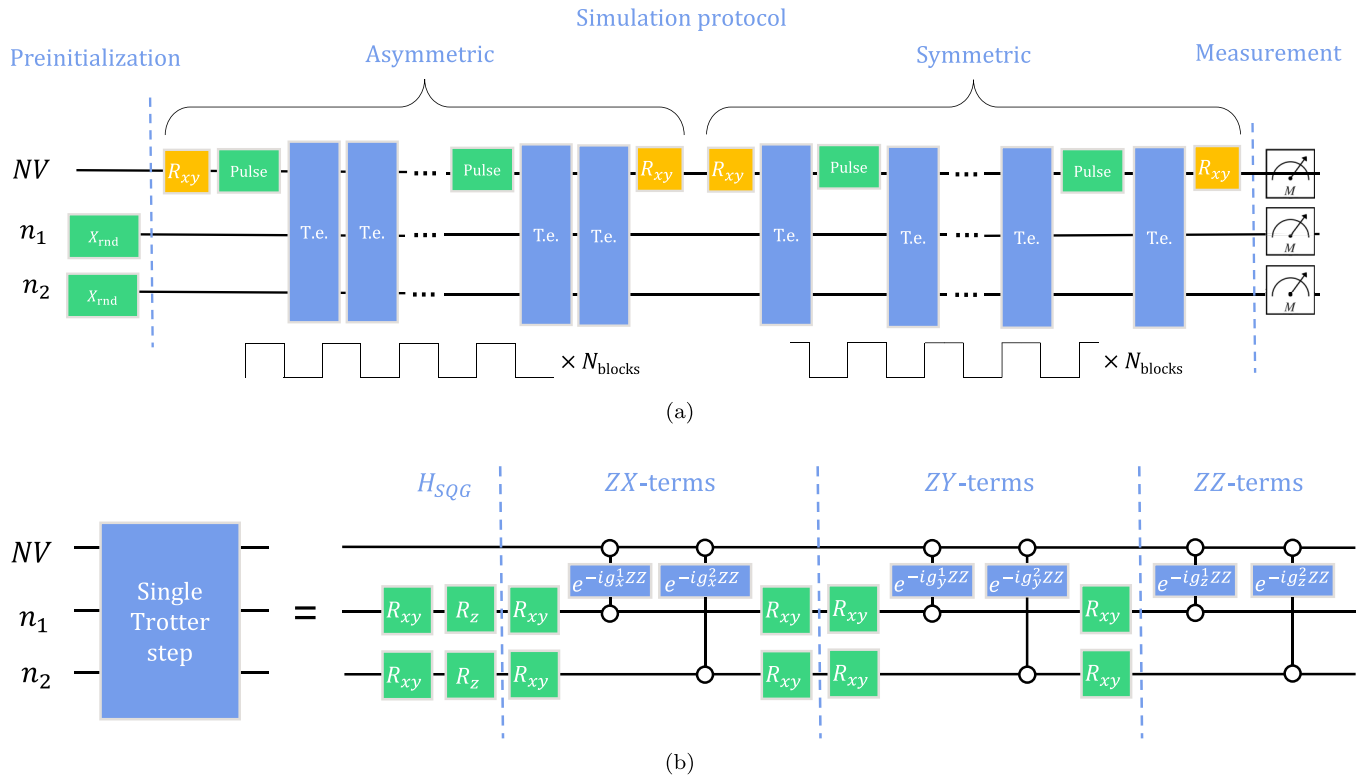$$H_I = \alpha A^\perp \sigma_z I_x + \beta A^\perp \sigma_z I_y, \tag{B11}$$

(a)

(b)

FIG. 10. (a) Sketch of the overall operation of the simulation algorithm for one NV center and two nuclei, with pulsed driving. Compare with Fig. 2. The R_{xy} gates highlighted in yellow are the rotations on the qubit representing the NV mentioned just before equation (B12). The term T.e. stands for Trotterized evolution and represents half of the free evolution of the system in between pulses, and can be devided into one or more single Trotter steps. The asymmetric and symmetric sequences of pulses are the ones discussed in Appendix B 2. The schematic drawings below the circuit in the form of square waves depict the modulation of the filter function [Eq. (B7)] under the two different pulse patterns. We choose the pulses to be either X or Y gates acting on the qubit representing the NV, following the pattern *XYXYYYXX*, which can be repeated N_{blocks} times for a stronger signal amplification. However, in our simulations a single block was enough to see clear patterns of polarization transfer, such as the ones in the right plot of Fig. 8. (b) Corresponding gate sequence of one Trotter step, as in Fig. 2, but without the H_{dr} inside, because in this case the driving is applied through the sequence of pulses.

which can be transformed with simple rotations on the qubit representing the NV into:

$$H_I = \alpha A^\perp \sigma_x I_x + \beta A^\perp \sigma_y I_y, \qquad \text{(B12)}$$

and this is equivalent to an interaction-exchange flip-flop Hamiltonian, similar to the one for the continuous-driving case (B6). A more detailed description of this whole process, including the expressions of the Fourier coefficients f_n and g_m can be found in reference [18].

In order to visualize the structure of the pulsed-driving case, we have included in Fig. 10 the circuit implementing all these terms on a quantum chip for the case of one NV center and two nuclei.

APPENDIX C: RANDOMIZED TROTTER TECHNIQUES

As explained in Sec. III, we chose Trotter expansion. Besides this, we can consider other simulation approaches such as the variational quantum simulator [33], the quantum assisted simulator [34], numerical quantum circuit synthesis [35], or a plethora of other quantum simulation algorithms aimed at NISQ devices [4].

In addition, other approaches like randomized Trotter have been recently shown to provide some advantage compared

to standard Trotter expansions [66]. We also propose to use one randomized approach, qDRIFT [32], that consists of the following: instead of splitting the whole evolution operator $e^{-it_f \sum_j h_j H_j}$ into simpler terms as done in full Trotterization, the method applies a random selection of such terms to the quantum circuit. This random selection is based on the probability distribution given by the weight of each term $h_j H_j$. For a certain evolution time, this set of gates can approximate the whole evolution operator by statistically drifting the state of the circuit towards the deterministic final state.

The error bound for this method is given as [32]:

$$\varepsilon_{\text{sim}}^{\text{qDRIFT}} \leqslant \frac{2\lambda^2 t_f^2}{N_{\text{terms}}}, \qquad \text{(C1)}$$

where $\lambda = \sum_j h_j$ and N_{terms} is the number of individual two-qubit evolution operators that are implemented. These evolution operators have the form $e^{-i\tau H_j}$, being τ a constant related to the relative weight $\frac{h_j}{\lambda}$ that the term H_j has in the Hamiltonian.

The advantage of qDRIFT compared to Trotterization is particularly apparent when dealing with Hamiltonians with a large number of terms with small coefficients, simulated for short times. While in the standard Trotter case, every term has

to be simulated for each step no matter how small its effect is, in qDRIFT this is not required. A more thorough analysis of errors in qDRIFT and gate counts can be found in [67].

This method is particularly suitable to our problem, since the range of coefficients in the Hamiltonian of a real diamond is large due to the length scales involved.

In this case, with qDRIFT the terms with smaller coefficients do not add a significant amount of gates as they would in conventional Trotterization approaches.

We note that other adapted protocols such as SparSto [68] can further enhance the simulation of this type of systems. SparSto represents a compromise between Trotterization and qDRIFT, generally guaranteeing an equal or better performance than both of them. We will not go into detail on this method since Trotterization and qDRIFT are enough to illustrate the main ideas behind this work.

APPENDIX D: HAMILTONIAN DECOMPOSITION FOR TROTTERIZED TIME EVOLUTION

In order to simulate the dynamics generated by the Hamiltonian in Eq. (1) on a quantum computer using Trotterization, we first need to express it in a suitable way. To begin with, we split the Hamiltonian into two parts:

$$H = H_{\text{SQG}} + H_{\text{TQG}}, \tag{D1}$$

which can be expressed in terms of qubit Pauli operators:

$$H_{\text{SQG}} = \sum_{k=1}^{N} \left[\frac{A_k^x}{2} \frac{X_k}{2} + \frac{A_k^y}{2} \frac{Y_k}{2} + \left(\frac{A_k^z}{2} - \gamma_c B_z \right) \frac{Z_k}{2} \right]$$
$$+ \sum_{j=1}^{M} \delta_j Z_j, \tag{D2}$$

$$H_{\text{TQG}} = \sum_{j=1}^{M} \sum_{k=1}^{N} \left[\frac{A_k^x}{2} \frac{X_k}{2} Z_j + \frac{A_k^y}{2} \frac{Y_k}{2} Z_j + \frac{A_k^z}{2} \frac{Z_k}{2} Z_j \right]$$
$$+ \sum_{k'>k=1}^{N} \frac{g_{k'k}}{4} \left[Z_{k'} Z_k - \frac{1}{2} X_{k'} X_k - \frac{1}{2} Y_{k'} Y_k \right]$$
$$+ \sum_{j>j'}^{M} h_{j'j} [Z_{j'} Z_j - X_{j'} X_j - Y_{j'} Y_j]. \tag{D3}$$

Since in the rotating frame with the drive the Hamiltonian is time independent, the time-evolution operator is simply given by:

$$U = e^{-it_f H}, \tag{D4}$$

where t_f is the time for which the simulation runs.

The time-evolution operator is split into s discrete steps through Trotter decomposition:

$$U = e^{-it_f H} = e^{-it_f(H_{\text{SQG}} + H_{\text{TQG}})} \approx \left[e^{-i\frac{t_f}{s} H_{\text{SQG}}} e^{-i\frac{t_f}{s} H_{\text{TQG}}} \right]^s$$
$$+ \mathcal{O}\left(\left(\frac{t_f}{s} \right)^2 \right). \tag{D5}$$

The evolution operator associated with single-qubit gates in each Trotter step of equation (D5) needs to be rewritten in

terms of our native gate set. It is always possible to decompose any single-qubit unitary exactly, up to a global phase, into a sequence of three single-qubit rotations such as, for example, a rotation about the y axis in between two rotations about the z axis:

$$U_1 = R_z(\beta) R_{xy}(\pi/2, \gamma) R_z(\delta), \tag{D6}$$

where the angles β, γ, and δ need to be determined from the specific entries of the unitary in question to simulate the evolution of the pth qubit:

$$U_1^p = e^{-i\frac{t_f}{s} \left(\frac{A_p^x}{2} \frac{X_p}{2} + \frac{A_p^y}{2} \frac{Y_p}{2} + \left(\frac{A_p^z}{2} - \gamma_c B_z \right) \frac{Z_p}{2} \right)}. \tag{D7}$$

From now on, we will concentrate on the case of a single NV center, which will be encoded in qubit 0. Then, the evolution operator associated to single-qubit gates for the NV center will be:

$$U_1^0 = e^{-i\frac{t_f}{s} \delta_0 Z_0}. \tag{D8}$$

Matching the entries of the matrices corresponding to the unitaries on equations (D7) and (D8) we get a system of equations for the angles β, γ, and δ for each Trotter step s.

There are 3 (5) types of interaction terms of the form $XZ, YZ, ZZ, \ldots$ in H_{TQG} without (with) internuclear interactions. Due to the native TQG being of only ZZ interaction type [see Eq. (5)], local rotations need to be introduced for simulating the rest of the TQG terms. These are $R_k^{\sigma_i \to \sigma_j}$, which have the effect of converting the Pauli operator σ_i into the Pauli operator σ_j for qubit k.

After the Trotterization introduced in equation (D5), the term H_{TQG} corresponding to TQG contains some elements which do not commute with each other, and some of them which do commute with each other. We choose to split all terms in order to express the time-evolution operator in terms of the native gates that we assumed in Sec. III B 1. Only the elements that do not commute with each other contribute to the total Trotter error, which remains of the same order:

$$e^{-i\frac{t_f}{s} H_{\text{TQG}}} \approx e^{-i\frac{t_f}{s} \left(\sum_k \frac{A_k^x}{2} \frac{X_k}{2} Z_0 \right)} e^{-i\frac{t_f}{s} \left(\sum_k \frac{A_k^y}{2} \frac{Y_k}{2} Z_0 \right)}$$
$$\times e^{-i\frac{t_f}{s} \left(\sum_k \frac{A_k^z}{2} \frac{Z_k}{2} Z_0 \right)} e^{-i\frac{t_f}{s} \left(\sum_{k'>k} \frac{g_{k'k}}{4} Z_{k'} Z_k \right)}$$
$$\times e^{i\frac{t_f}{s} \left(\sum_{k'>k} \frac{g_{k'k}}{8} X_{k'} X_k \right)} e^{i\frac{t_f}{s} \left(\sum_{k'>k} \frac{g_{k'k}}{8} Y_{k'} Y_k \right)}$$
$$+ \mathcal{O}\left(\left(\frac{t_f}{s} \right)^2 \right). \tag{D9}$$

Finally, we observe that the operators $Z_k Z_0$ (and the rest of the TQG terms) commute with each other, so the exponentials can be further split without Trotterizing:

$$e^{-i\frac{t_f}{s} \left(\sum_k \frac{A_k^z}{2} \frac{Z_k}{2} Z_0 \right)} = \Pi_k e^{-i\frac{t_f}{s} \left(\frac{A_k^z}{2} \frac{Z_k}{2} Z_0 \right)}. \tag{D10}$$

The time-evolution operator implementing the continuous sinusoidal driving σ^ϕ is:

$$e^{-i\frac{t_f}{s} \frac{\Omega}{2} \sigma^\phi} = R_{xy}\left(-\phi, \theta = \Omega \frac{t_f}{s} \right). \tag{D11}$$

The quantum algorithm for simulating the system under a pulsed-driving scheme is somewhat more involved than
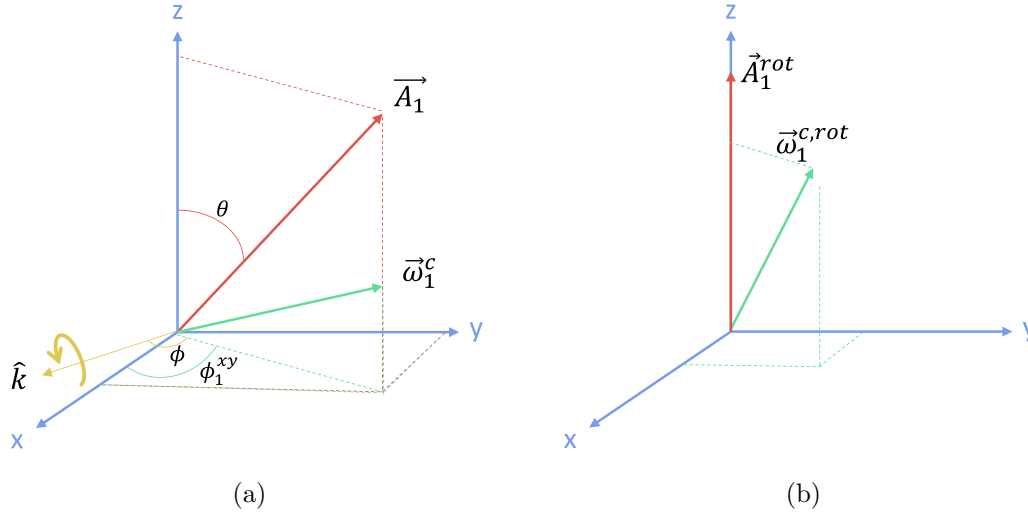
FIG. 11. (a) Coefficients vectors of the first qubit $\vec{A}_1$, $\vec{\omega}_1^c$ before the rotation, with projection over the three axis and (b) coefficient vectors of the first qubit $\vec{A}_1^{\mathrm{rot}}$, $\vec{\omega}_1^{c,\mathrm{rot}}$ after the rotation, being $\vec{A}_1^{\mathrm{rot}}$ in the Z axis.

the continuous-driving case, due to the two different time-dependent processes involved in the Trotter decomposition: the free dynamics of the spins and the sequence of pulses. The most crucial point to be aware of is the interplay between Trotter steps and interpulse spacing. The number of interpulse evolutions, i.e. number of pulses minus one, bounds from below the minimum number of Trotter steps for the simulation. Clearly, at least one Trotter step is needed for each interpulse evolution.

Taking this interplay into account, the most straightforward setup is to choose a frequency which will determine the spacing of the pulse sequence, and to identify each interpulse evolution with a single Trotter step. If the achieved precision is not high enough, more Trotter steps can be added for each interpulse evolution. Each π-pulse itself is simply implemented as an X or Y gate on the qubit representing the NV center. The OU-distributed Rabi frequency fluctuations present in nanoscale NMR systems are then simulated by over- and under-rotations of the X and Y gates.

APPENDIX E: ROTATIONAL OPTIMIZATION

In principle, we had a Hamiltonian with terms of the type ZX, ZY, and ZZ for the case of no internuclear interactions. However, we can rotate the basis so the Hamiltonian loses the ZX and ZY terms, allowing to reduce the number of TQGs. To make up for this rotation, we need to introduce different constants $\vec{A}_i^{\mathrm{rot}}$ for the problem and rotate the vector state we obtain at the end before measuring it. The rotations that we will consider are only one-qubit rotations on nuclei qubits and we are applying this just to the case with no internuclear interactions. Therefore we can consider the effect of this rotation on only one qubit representing an arbitrary nucleus. We will exemplify this procedure using nucleus 1. If we want to obtain the mean value of σ_z acting on the nucleus:

$$\langle \sigma_z \rangle = \mathrm{Tr}(\rho(t_f)\sigma_z) = \mathrm{Tr}(U(0,t_f)\rho(0)U^\dagger(0,t_f)\sigma_z), \quad (E1)$$

where $U(0,t_f)$ represents the evolution operator from $t=0$ to $t=t_f$. The density matrix $\rho(0)$ contains the state of the NV center (which is in the $|+\rangle$ state at $t=0$) and nucleus 1, i.e., $\rho(0) = |+\rangle\langle+|_{\mathrm{NV}} \otimes \frac{\mathbb{1}_1}{2}$. Our intention is to obtain an expression of this mean value in terms of the rotated evolution operators and later, we will find the appropriate rotation to be performed. Then, taking into account that the trace is invariant under a rotation $R = \mathbb{1}_{\mathrm{NV}} \otimes R_1$, we get:

$$\langle \sigma_z \rangle = \mathrm{Tr}(RU(0,t_f)\rho(0)U^\dagger(0,t_f)\sigma_z R^\dagger)$$
$$= \mathrm{Tr}(RU(0,t_f)R^\dagger R\rho(0)R^\dagger RU^\dagger(0,t_f)R^\dagger R\sigma_z R^\dagger). \quad (E2)$$

This can be expressed as:

$$\langle \sigma_z \rangle = \mathrm{Tr}(U_{\mathrm{rot}}(0,t_f)\rho_{\mathrm{rot}}(0)U_{\mathrm{rot}}^\dagger(0,t_f)R\sigma_z R^\dagger). \quad (E3)$$

The density matrix of the nucleus is the identity. Thus any rotation on nuclei qubits leaves the density matrix unaffected, leading to:

$$\langle \sigma_z \rangle = \mathrm{Tr}(U_{\mathrm{rot}}(0,t_f)\rho(0)U_{\mathrm{rot}}^\dagger(0,t_f)R\sigma_z R^\dagger). \quad (E4)$$

Then we need to rotate the system previous to the measurement. By using the invariance of the trace under cyclic permutations, we get:

$$\langle \sigma_z \rangle = \mathrm{Tr}(R^\dagger U_{\mathrm{rot}}(0,t_f)\rho(0)U_{\mathrm{rot}}^\dagger(0,t_f)R\sigma_z), \quad (E5)$$

which is equivalent to introducing a counter-rotation in the circuit before measurement.

Now let us focus on the specific rotation we have to implement. Since the constants multiplying the Pauli matrices in the Hamiltonian are $\frac{\vec{A}_1}{2}$ and $\vec{\omega}_1^c = \frac{\vec{A}_1}{2} - \gamma_c B_z \vec{e}_z$ (for nucleus 1), we can rotate the basis to obtain a representation in which the vectors have only z component for $\vec{A}_1$ and thus, XZ and YZ terms are removed. The vectors before and after the needed rotation can be seen in Fig. 11.

To compute the new vectors (and thus the new coefficients for the gates of our algorithm), we can use Rodrigues' rotation formula to rotate a vector $\vec{v}$ an angle θ around a unitary

TABLE III. Gate count for one Trotter step and for one cycle for different topologies with and without internuclear interactions.

	All-To-All	Star topology	Square grid
$N_{\mathrm{TQG}}^{\mathrm{nonint}}$	$n-1$	$n-1$	$4n-4$
$N_{\mathrm{SQG}}^{\mathrm{nonint}}$	$\frac{5}{2}n+2$	$\frac{5}{2}n+2$	$\frac{21}{2}n-\frac{47}{2}$
$N_{\mathrm{TQG}}^{\mathrm{int}}$	$\frac{3}{2}n^2-\frac{3}{2}n$	$\frac{3}{2}n^2+\frac{3}{2}n-6$	$3n^2-6n+3$
$N_{\mathrm{SQG}}^{\mathrm{int}}$	$4n^2-\frac{9}{2}n+\frac{7}{2}$	$4n^2+\frac{7}{2}n-\frac{25}{2}$	$8n^2-\frac{33}{2}n+\frac{11}{2}$

axis $\hat{k}$:

$$\vec{v}_{\mathrm{rot}} = \vec{v}\cos\theta + (\hat{k}\times\vec{v})\sin\theta + \hat{k}(\hat{k}\cdot\vec{v})(1-\cos\theta), \quad \text{(E6)}$$

being in our case, $\theta = \arccos(A_1^z/|\vec{A}_1|)$ and $\hat{k} = (\cos(\phi), \sin(\phi), 0)$, with $\phi = -\frac{\pi}{2} + \phi_{xy} = -\frac{\pi}{2} + \arctan(A_1^y/A_1^x)$.

For implementing the counter-rotation of this in the quantum circuit, we use:

$$R_1^\dagger = e^{i\frac{\theta}{2}(\cos(\phi)X - \sin(\phi)Y)}. \quad \text{(E7)}$$

APPENDIX F: SWAP ROUTING

Our qubit routing method consists of mapping the square grid to a linear chain with qubits labeled from 0 to n. Then, in the simplified case of no internuclear interactions, the optimal SWAP method for the one-to-all interaction case on a linear chain can be used. For a single NV center the protocol goes as follows.

(1) Initialize the state of the NV center in the second qubit.

(2) Perform interactions with the first and third qubits.

(3) SWAP the NV center qubit to the right.

(4) Perform interaction with right qubit.

(5) Repeat steps 3-4 until all interactions have been achieved.
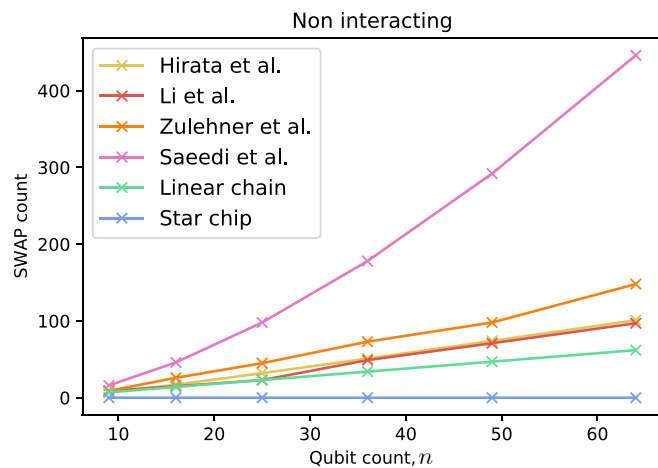
The pattern is seen in Fig. 3(a) denoted by the intense blue arrows. With internuclear interactions we need to perform a swap pattern that enables all-to-all interactions. The so-called odd-even mapping in Fig. 3(a) is an efficient one [69] represented by green arrows in Fig. 3(a). This consists of swapping first all the even qubits with their right neighbors and then swapping all the odd qubits with their right neighbors. This way, we will obtain all-to-all interactions with $\frac{1}{2}(n-1)(n-2)$ SWAP gates and a total TQG depth of $6n$. A summary of the TQG counts is shown in Table III.

To motivate the creation of a chip with a star topology and the use of an alternative linearized SWAP routing for a square grid instead of standard numerical approaches, a comparison between all the cases is provided in Fig. 12. A reduction in the number of SWAPs can be noticed for both the linear chain approach and the star-topology chip against standard numerical approaches for a square grid.

APPENDIX G: QUBIT-RESONATOR GATE THEORY

In the following discussion, we consider gate operation between the resonator and one of the qubits, and neglect any effects that arise from the interactions with spectator qubits and other resonator modes. The time dynamics in such a system are determined by the Hamiltonian:

$$H = H_0 + H_{rc} + H_{qc} + H_{rq}, \quad \text{(G1)}$$

where the uncoupled part of the total Hamiltonian $H_0 = H_r + H_c + H_q$ is:

$$\begin{aligned} H_r &= \hbar\omega_r b_r^\dagger b_r, \\ H_c &= \hbar\omega_c b_c^\dagger b_c + \frac{\hbar}{2}\alpha_c b_c^\dagger b_c^\dagger b_c b_c, \\ H_q &= \hbar\omega_q b_q^\dagger b_q + \frac{\hbar}{2}\alpha_q b_q^\dagger b_q^\dagger b_q b_q, \end{aligned} \quad \text{(G2)}$$
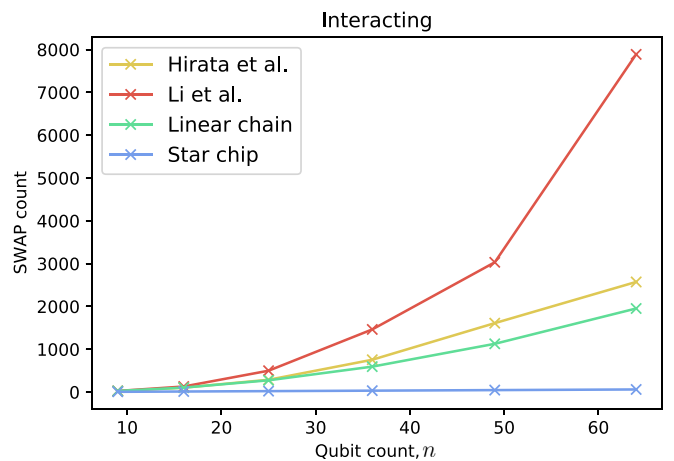
FIG. 12. (a) Comparison of the required number of SWAPs for simulating the proposed system with no internuclear interactions for each Trotter step. Numerical approaches from references are applied to a square grid. (b) Equivalent comparison with internuclear interactions. Zulehner *et al.* and Saeedi *et al.* do not improve the linear chain approach for few qubits and are intractable for larger numbers of qubits and thus are not displayed.

where b_λ and ω_λ are the annihilation operator and fundamental frequency for the mode $\lambda = \{r, c, q\}$, respectively, and α_γ is the anharmonicity of the mode $\gamma = \{q, c\}$. The interaction component of the Hamiltonian is:

$$H_{\lambda\mu} = -\hbar g_{\lambda\mu}(b_\lambda^\dagger - b_\lambda)(b_\mu^\dagger - b_\mu), \tag{G3}$$

where $\lambda\mu = \{rc, qc, rq\}$, and $g_{\lambda\mu}$ denote resonator-coupler, qubit-coupler and resonator-qubit coupling frequencies. With the Hamiltonian of Eq. (G1), we are now in a position to perform simulations of two-qubit gates by propagating a suitably chosen initial state.

Before the gate operation, we choose the idling frequencies for the qubit, resonator, and the coupler such that the CZ coupling rate ζ is minimized. This CZ coupling rate is defined as:

$$\zeta = \omega_{101} - \omega_{100} - \omega_{001} + \omega_{000}, \tag{G4}$$

where $\omega_{n_r 0 n_q}$ corresponds to the eigenenergy of Hamiltonian in Eq. (G1) with n_r excitations in resonator and n_q excitations in qubit with coupler being in the ground state. The point of minimal $|\zeta|$ is also known as the idling configuration, which we found to be at $[\omega_r, \omega_c, \omega_q]/(2\pi) = [4.30, 6.14, 4.47]$ GHz for the parameters given in Table II. The CZ gate is operated by sending a flux pulse that modifies the coupler frequency ω_c, which then in the coupled basis modifies the frequencies $\omega_{101}, \omega_{100}, \omega_{001}$ and ω_{000}. This makes ζ nonzero, so the system collects a CZ phase.

[1] R. Feynman, Simulating physics with computers, in *Feynman and Computation* (CRC Press, 2018), pp. 133–153.

[2] S. Lloyd, Universal quantum simulators, Science **273**, 1073 (1996).

[3] J. Preskill, Quantum computing in the NISQ era and beyond, Quantum **2**, 79 (2018).

[4] K. Bharti, A. Cervera-Lierta, T. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. Kottmann, T. Menke *et al.*, Noisy intermediate-scale quantum (NISQ) algorithms, Rev. Mod. Phys. **94**, 015004 (2022).

[5] M. J. S. Smith, *Application-Specific Integrated Circuits*, Vol. 7 (Addison-Wesley Reading, MA, 1997).

[6] K.-C. Hsu and H.-W. Tseng, GPTPU: Accelerating applications using edge tensor processing units, in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (Association for Computing Machinery, New York, NY, 2021), pp. 1–4, doi: 10.1145/3458817.3476177.

[7] T. Lu, T. Marin, Y. Zhuo, Y.-F. Chen, and C. Ma, Accelerating MRI reconstruction on TPUs, in *2020 IEEE High Performance Extreme Computing Conference (HPEC)* (IEEE, 2020), pp. 1–9, doi: 10.1109/HPEC43674.2020.9286192.

[8] T. Staudacher, F. Shi, S. Pezzagna, J. Meijer, J. Du, C. A. Meriles, F. Reinhard, and J. Wrachtrup, Nuclear magnetic resonance spectroscopy on a (5-Nanometer)3 sample volume, Science **339**, 561 (2013).

[9] M. Nielsen and I. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, 2000).

[10] H. Hogben, M. Krzystyniak, G. Charnock, P. Hore, and I. Kuprov, Spinach a software library for simulation of spin dynamics in large spin systems, J. Magn. Reson. **208**, 179 (2011).

[11] K. Seetharam, D. Biswas, C. Noel, A. Risinger, D. Zhu, O. Katz, S. Chattopadhyay, M. Cetina, C. Monroe, E. Demler *et al.*, Digital quantum simulation of NMR experiments, arXiv:2109.13298.

[12] 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, arXiv:2109.02163.

[13] 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).

[14] D. Sels and E. Demler, Quantum generative model for sampling many-body spectral functions, Phys. Rev. B **103**, 014301 (2021).

[15] M. H. Levitt, *Spin Dynamics: Basics of Nuclear Magnetic Resonance* (Wiley, 2013).

[16] M. W. Doherty, N. B. Manson, P. Delaney, F. Jelezko, J. Wrachtrup, and L. C. L. Hollenberg, The nitrogen-vacancy colour centre in diamond, Phys. Rep. **528**, 1 (2013).

[17] M. H. Abobeih, J. Randall, C. E. Bradley, H. P. Bartling, M. A. Bakker, M. J. Degen, M. Markham, D. J. Twitchen, and T. H. Taminiau, Atomic-scale imaging of a 27-nuclear-spin cluster using a quantum sensor, Nature (London) **576**, 411 (2019).

[18] C. Munuera-Javaloy, R. Puebla, and J. Casanova, Dynamical decoupling methods in nanoscale NMR, Europhys. Lett. **134**, 30001 (2021).

[19] C. Munuera-Javaloy, R. Puebla, B. D'Anjou, M. B. Plenio, and J. Casanova, Detection of molecular transitions with nitrogen-vacancy centers and electron-spin labels, arXiv:2110.14255.

[20] A. Ajoy, K. Liu, R. Nazaryan, X. Lv, P. Zangara, B. Safvati, G. Wang, D. Arnold, G. Li, A. Lin *et al.*, Orientation-independent room temperature optical 13c hyperpolarization in powdered diamond, Sci. Adv. **4**, eaar5492 (2018).

[21] S. Hartmann and E. Hahn, Nuclear double resonance in the rotating frame, Phys. Rev. **128**, 2042 (1962).

[22] H. Carr and E. Purcell, Effects of diffusion on free precession in nuclear magnetic resonance experiments, Phys. Rev. **94**, 630 (1954).

[23] S. Meiboom and D. Gill, Modified spin-echo method for measuring nuclear relaxation times, Rev. Sci. Instrum. **29**, 688 (1958).

[24] A. Maudsley, Modified Carr-Purcell-Meiboom-Gill sequence for NMR Fourier imaging applications, J. Magn. Reson. (1969) **69**, 488 (1986).

[25] T. Gullion, D. Baker, and M. Conradi, New, compensated Carr-Purcell sequences, J. Magn. Reson. (1969) **89**, 479 (1990).

[26] J. Casanova, Z. Y. Wang, and M. B. Plenio, Noise-Resilient Quantum Computing with a Nitrogen-Vacancy Center and Nuclear Spins, Phys. Rev. Lett. **117**, 130502 (2016).

[27] G. Uhlenbeck and L. Ornstein, On the theory of the Brownian motion, Phys. Rev. **36**, 823 (1930).

[28] J. Cai, B. Naydenov, R. Pfeiffer, L. McGuinness, K. Jahnke, F. Jelezko, M. Plenio, and A. Retzker, Robust dynamical

decoupling with concatenated continuous driving, New J. Phys. **14**, 113023 (2012).

[29] D. T. Gillespie, Exact numerical simulation of the Ornstein-Uhlenbeck process and its integral, Phys. Rev. E **54**, 2084 (1996).

[30] A. Ajoy, B. Safvati, R. Nazaryan *et al.*, Hyperpolarized relaxometry based nuclear T1 noise spectroscopy in diamond, Nat. Commun. **10**, 5160 (2019).

[31] M. Suzuki, Generalized Trotter's formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems, Commun. Math. Phys. **51**, 183 (1976).

[32] E. Campbell, Random Compiler for Fast Hamiltonian Simulation, Phys. Rev. Lett. **123**, 070503 (2019).

[33] X. Yuan, S. Endo, Q. Zhao, Y. Li, and S. Benjamin, Theory of variational quantum simulation, Quantum **3**, 191 (2019).

[34] K. Bharti and T. Haug, Quantum-assisted simulator, Phys. Rev. A **104**, 042418 (2021).

[35] E. Younis, K. Sen, K. Yelick, and C. Iancu, Qfast: Conflating search and numerical optimization for scalable quantum circuit synthesis, arXiv:2103.07093.

[36] N. Hatano and M. Suzuki, *Finding Exponential Product Formulas of Higher Orders* (Springer, Berlin, Heidelberg, 2005), pp. 37–68.

[37] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver, A quantum engineer's guide to superconducting qubits, Appl. Phys. Rev. **6**, 021318 (2019).

[38] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, Efficient Z gates for quantum computing, Phys. Rev. A **96**, 022330 (2017).

[39] F. Yan, P. Krantz, Y. Sung, M. Kjaergaard, D. L. Campbell, T. P. Orlando, S. Gustavsson, and W. D. Oliver, Tunable Coupling Scheme for Implementing High-Fidelity Two-Qubit Gates, Phys. Rev. Appl. **10**, 054062 (2018).

[40] K. Y. Tan, M. Partanen, R. E. Lake, J. Govenius, S. Masuda, and M. Möttönen, Quantum-circuit refrigerator, Nat. Commun. **8**, 15189 (2017).

[41] M. Silveri, H. Grabert, S. Masuda, K. Y. Tan, and M. Möttönen, Theory of quantum-circuit refrigeration by photon-assisted electron tunneling, Phys. Rev. B **96**, 094524 (2017).

[42] H. Hsu, M. Silveri, A. Gunyhó, J. Goetz, G. Catelani, and M. Möttönen, Tunable refrigerator for nonlinear quantum electric circuits, Phys. Rev. B **101**, 235422 (2020).

[43] V. Sevriuk, K. Y. Tan, E. Hyyppä, M. Silveri, M. Partanen, M. Jenei, S. Masuda, J. Goetz, V. Vesterinen, L. Grönberg *et al.*, Fast control of dissipation in a superconducting resonator, Appl. Phys. Lett. **115**, 082601 (2019).

[44] M. Celio, New Method to Calculate the Muon Polarization Function, Phys. Rev. Lett. **56**, 2720 (1986).

[45] S. McArdle, Learning from physics experiments with quantum computers: Applications in muon spectroscopy, PRX Quantum **2**, 020349 (2021).

[46] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, Phys. Rev. A **86**, 032324 (2012).

[47] Y. Hirata, M. Nakanishi, S. Yamashita, and Y. Nakashima, An efficient method to convert arbitrary quantum circuits to ones on a linear nearest neighbor architecture, in *2009 Third International Conference on Quantum, Nano and Micro*

Technologies, Cancun, Mexico (IEEE, 2009), pp. 26–33, doi: 10.1109/ICQNM.2009.25.

[48] G. Li, Y. Ding, and Y. Xie, Tackling the qubit mapping problem for NISQ-era quantum devices, in *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems* (ACM, New York, NY, USA, 2019), pp. 1001–1014, doi: 10.1145/3297858.3304023.

[49] M. Saeedi, R. Wille, and R. Drechsler, Synthesis of quantum circuits for linear nearest neighbor architectures, Quant. Info. Proc. **10**, 355 (2011).

[50] A. Zulehner, A. Paler, and R. Wille, An efficient methodology for mapping quantum circuits to the IBM QX Architectures, in *IEEE transactions on computer-aided design of integrated circuits and systems* (IEEE, 2019), Vol. 38, pp. 1226–1236, doi: 10.1109/TCAD.2018.2846658.

[51] E. Magesan and J. M. Gambetta, Effective Hamiltonian models of the cross-resonance gate, Phys. Rev. A **101**, 052308 (2020).

[52] N. Lacroix, C. Hellings, C. K. Andersen, A. D. Paolo, A. Remm, S. Lazar, S. Krinner, G. J. Norris, M. Gabureac, J. Heinsoo, A. Blais, C. Eichler, and A. Wallraff, Improving the performance of deep quantum optimization algorithms with continuous gate sets, PRX Quantum **1**, 020304 (2020).

[53] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, arXiv:1411.4028.

[54] M. Mariantoni, F. Deppe, A. Marx, R. Gross, F. K. Wilhelm, and E. Solano, Two-resonator circuit quantum electrodynamics: A superconducting quantum switch, Phys. Rev. B **78**, 104508 (2008).

[55] B. Foxen, C. Neill, A. Dunsworth, P. Roushan, B. Chiaro, A. Megrant, J. Kelly, Z. Chen, K. Satzinger, R. Barends *et al.*, Demonstrating a Continuous Set of Two-qubit Gates for Near-term Quantum Algorithms, Phys. Rev. Lett. **125**, 120504 (2020).

[56] C. Song, K. Xu, W. Liu, C. P. Yang, S. B. Zheng, H. Deng, Q. Xie, K. Huang, Q. Guo, L. Zhang, P. Zhang, D. Xu, D. Zheng, X. Zhu, H. Wang, Y. A. Chen, C. Y. Lu, S. Han, and J. W. Pan, 10-Qubit Entanglement and Parallel Logic Operations with a Superconducting Circuit, Phys. Rev. Lett. **119**, 180511 (2017).

[57] C. Song, K. Xu, H. Li, Y. Zhang, X. Zhang, W. Liu, Q. Guo, Z. Wang, W. Ren, J. Hao *et al.*, Generation of multicomponent atomic Schrödinger cat states of up to 20 qubits, Science **365**, 574 (2019).

[58] S. Hazra, A. Bhattacharjee, M. Chand, K. V. Salunkhe, S. Gopalakrishnan, M. P. Patankar, and R. Vijay, Ring-Resonator-Based Coupling Architecture for Enhanced Connectivity in a Superconducting Multiqubit Network, Phys. Rev. Appl. **16**, 024018 (2021).

[59] J. Chu and F. Yan, Coupler-Assisted Controlled-Phase Gate with Enhanced Adiabaticity, Phys. Rev. Appl. **16**, 054020 (2021).

[60] J. Johansson, P. Nation, and F. Nori, Qutip: An open-source Python framework for the dynamics of open quantum systems, Comput. Phys. Commun. **183**, 1760 (2012).

[61] J. Johansson, P. Nation, and F. Nori, Qutip 2: A Python framework for the dynamics of open quantum systems, Comput. Phys. Commun. **184**, 1234 (2013).

[62] Y. Zhou, Z. Zhang, Z. Yin, S. Huai, X. Gu, X. Xu, J. Allcock, F. Liu, G. Xi, Q. Yu *et al.*, Rapid and unconditional

parametric reset protocol for tunable superconducting qubits, arXiv:2103.11315.

[63] S. Endo, S. C. Benjamin, and Y. Li, Practical Quantum Error Mitigation for Near-Future Applications, Phys. Rev. X **8**, 031027 (2018).

[64] Z. Cai, Multi-exponential error extrapolation and combining error mitigation techniques for NISQ applications, npj Quantum Inf. **7**, 80 (2021).

[65] M. Krebsbach, B. Trauzettel, and A. Calzona, Optimization of Richardson extrapolation for quantum error mitigation, arXiv:2201.08080.

[66] A. M. Childs, A. Ostrander, and Y. Su, Faster quantum simulation by randomization, Quantum **3**, 182 (2019).

[67] C. F. Chen, H. Y. Huang, R. Kueng, and J. A. Tropp, Concentration for random product formulas, PRX Quantum **2**, 040305 (2021).

[68] Y. Ouyang, D. R. White, and E. T. Campbell, Compilation by stochastic Hamiltonian sparsification, Quantum **4**, 235 (2020).

[69] A. Cowtan, S. Dilkes, R. Duncan, A. Krajenbrink, W. Simmons, and S. Sivarajah, On the Qubit Routing Problem, in *14th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2019)*, Leibniz International Proceedings in Informatics (LIPIcs), Vol. 135 (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2019), pp. 51–532, doi: 10.4230/LIPIcs.TQC.2019.5.