

Comparing planar quantum computing platforms at the quantum speed limit

Daniel Basilewitsch^{1,*}, Clemens Dlaska¹, and Wolfgang Lechner^{1,2}
¹*Institute for Theoretical Physics, University of Innsbruck, A-6020 Innsbruck, Austria*
²*Parity Quantum Computing GmbH, A-6020 Innsbruck, Austria*

 (Received 24 May 2023; revised 30 November 2023; accepted 12 March 2024; published 5 April 2024)

An important aspect that strongly impacts the experimental feasibility of quantum circuits is the ratio of gate times and typical error time scales. Algorithms with circuit depths that significantly exceed the error time scales will result in faulty quantum states and error correction is inevitable. We present a comparison of the theoretical minimal gate time, i.e., the quantum speed limit (QSL), for realistic two- and multi-qubit gate implementations in neutral atoms and superconducting qubits. Subsequent to finding the QSLs for individual gates by means of optimal control theory we use them to quantify the circuit QSL of the quantum Fourier transform and the quantum approximate optimization algorithm. In particular, we analyze these quantum algorithms in terms of circuit run times and gate counts both in the standard gate model and the parity mapping. We find that neutral atom and superconducting qubit platforms show comparable weighted circuit QSLs with respect to the system size.

DOI: [10.1103/PhysRevResearch.6.023026](https://doi.org/10.1103/PhysRevResearch.6.023026)

I. INTRODUCTION

Quantum computers promise to solve computational problems that are deemed hard or even intractable for classical computers. Their potential applications include prime factoring of large integers [1], quantum simulation [2], quantum chemistry [3], combinatorial optimization [4], and even problems in finance [5]. Currently, quantum computing is in the so-called noisy intermediate-scale quantum (NISQ) era [6], characterized by imperfect qubit control, and qubit numbers that prohibit quantum error correction [7] for relevant problem sizes. Nevertheless, recent proof-of-principle experiments [8–11] demonstrated that a computational quantum advantage over classical computers can be reached already with NISQ hardware. However, it remains a crucial challenge to go beyond the proof-of-principle stage, i.e., to demonstrate a quantum advantage for practically relevant computational tasks on resource-limited present-day devices.

To reach a practical quantum advantage regime [2] in NISQ-era digital quantum computing it is of crucial importance to execute quantum algorithms as efficiently as possible in order to minimize the time for noise mechanisms to impair the quantum information processing. This effectively makes a minimization of the quantum algorithm run times and gate counts desirable—a task that can be addressed in various ways. One option is to find an algorithm’s optimal circuit representation, i.e., a circuit requiring a minimal circuit

depth together with a minimal gate count for a given set of available gates. Quantum circuit optimization has, e.g., been done heuristically [12,13] or by machine learning techniques [14,15] and several open-source packages are readily available [16–18].

Another option for minimizing the algorithm run time is to minimize the time for each elemental quantum gate of a given quantum circuit. While protocols for fast, high-fidelity quantum gates are nowadays routinely available and implemented on all major quantum computing platforms like neutral atoms [19,20], superconducting circuits [8,21] or trapped ions [22,23], the ideal would be to execute every quantum gate at its quantum speed limit (QSL). In general, the QSL denotes the shortest time needed to accomplish a given task [24]. It constitutes a fundamental limit in time and depends on the system under consideration, i.e., its Hamiltonian and the control knobs available to steer the dynamics.

Here, we determine the QSLs of quantum gates for two major quantum computing platforms that allow for two-dimensional (2D) qubit arrangements—neutral atoms and superconducting circuits [25]. Our study reveals how close the current experimental gate protocols are in comparison to their QSLs and thus exemplifies what can theoretically still be gained from further speeding up gate protocols. Moreover, provided that every gate could be experimentally realized at the QSL, our analysis gives an estimate of how many gates can be executed realistically before decoherence takes over and renders longer quantum circuits practically infeasible. To this end, we consider two prototypical quantum computing algorithms: (i) the quantum Fourier transform (QFT), required for Shor’s algorithm for integer factorization [1], and (ii) the quantum approximate optimization algorithm (QAOA), used to solve combinatorial optimization problems [4]. Considering standard NISQ devices for both neutral atoms and superconducting circuits with qubits arranged in a 2D grid

*daniel.basilewitsch@uibk.ac.at

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

architecture with only nearest-neighbor connectivity, we calculate the circuit run times with gates at the QSL for both algorithms. This allows for a direct comparison of both platforms in terms of the maximal problem sizes that should currently be feasible on their NISQ representatives.

A common challenge arising in 2D platforms with nearest-neighbor connectivity is the requirement to perform gates between non-neighbor qubits. In the standard gate model (SGM), such gates can be replaced by sequences of universal single- and two-qubit gates using the available local connectivity. However, this comes at the price of increasing the circuit depths and gate counts. As an alternative to the SGM, we also examine circuit representations using the so-called parity mapping (PM). In brevity, the PM for quantum computing [26] and quantum optimization [27,28] is a problem-independent hardware blueprint that only requires nearest-neighbor connectivity at the cost of increased qubit numbers. Since for QAOA circuits in the PM it is beneficial to use local three- and four-qubit gates [29,30], we also determine their QSLs on both platforms.

It should be noted that this paper focuses entirely on the determination of the QSLs for various quantum gates and how to turn these into a fair comparison of circuit run times across platforms. A more detailed discussion of gate protocols for specific quantum gates or platforms can, e.g., be found in Refs. [31–33] for neutral atoms or in Refs. [34–36] for superconducting circuits. Moreover, it should be noted that we only consider the gate times, respectively circuit run times, as well as the number of gates as indicators for the feasibility of quantum circuits. While both quantities are doubtlessly important, they are by no means the only quantities impacting a circuit’s feasibility. A holistic figure of merit assessing a circuit’s feasibility would also need to account for state preparation and measurement errors and various other error sources.

The paper is organized as follows. In Sec. II we present our main result, i.e., an overview of circuit run times when using gate times from the literature and gate times at the QSL—evaluated both for neutral atoms and superconducting circuits as well as for circuits in the SGM and the PM. In Sec. III we then present the details of our numerical model. Section IV B introduces the basic notion of QSLs and how quantum optimal control theory (OCT) can be used for determining QSLs. A detailed discussion of QSLs as well as which combinations of available control fields allow to reach the QSLs, is given in Sec. V. Section VI concludes the paper.

II. MAIN RESULT: QUANTUM ALGORITHMS AT THE QUANTUM SPEED LIMIT

In this section, we compare the run times of QFT and QAOA quantum circuits using gate set implementations available on neutral atom and superconducting circuit hardware. To do so we consider two scenarios. In the first scenario, we use literature values for gate times of state-of-the-art gate implementations of the minimal universal gate set [referred to as “standard gate set” (SGS) from now on] native to each platform, which thus represents the canonical way of converting quantum algorithms into executable quantum circuits. In the second scenario, we use an extended set of gates available at each platform with gate times at the QSL [referred to as

“QSL gate set” (QGS) from now on], which thus yields the current fundamental limits in circuit run times. We describe both gate sets in Sec. II A and use them to analyze circuit run times of QFT and QAOA quantum circuits both in the SGM in Sec. II B and the PM in Sec. II C. A brief review of the basic concepts and quantum circuits of a QFT and a single QAOA step within the SGM and the PM is given in Appendix A.

A. Standard and QSL gate sets

In the SGM, each quantum algorithm is converted into quantum circuits using gates from a universal set of quantum gates. Such a universal set typically contains all the single-qubit gates and at least one entangling two-qubit gate, e.g., the CNOT gate [38]. The row named “standard gate set” in Table I summarizes the native universal gate sets and their typical gate times for both platforms at comparable gate fidelities. For neutral atoms, we consider the controlled-Z gate, $CZ = \text{diag}\{1, 1, 1, -1\}$, as the typical entangling two-qubit gate. This is a common choice for neutral atoms [20] as it has been successfully used for implementing quantum algorithms [39,40]. For superconducting circuits, we have chosen the iSWAP-like Sycamore gate as an entangling two-qubit gate, motivated by its short gate time and successful usage in recent quantum advantage experiments with quantum processors based on tunable couplers [8,9].

In addition, we consider an extended, platform-independent set of quantum gates operated at the QSL (see “QSL gate set” in Table I). This set consists of all single-qubit gates as well as several multi-qubit gates implementable on both platforms: CZ, CNOT, SWAP, ZZZ, and ZZZZ. The availability of a wider range of gates allows for more flexibility in finding circuit representations with fewer gates and shorter circuit run times. In addition, we also assume that every gate in this set is executed at the QSL, which allows for another speed-up of circuit run times. The run times obtained for the QGS should therefore be viewed as an estimate for the QSL of the circuit itself, i.e., the circuit QSL—provided that no other, potentially faster and/or better suited gates are available [41]. Note that the details regarding the method to determine the QSLs and their results are presented in Secs. IV B and V. The row “QSL gate set” in Table I summarizes the gates of the extended set and lists the QSL times for both platforms.

Table I indicates that absolute circuit run times for neutral atoms will be longer compared to those for superconducting circuits since their elemental gate times differ by more than an order of magnitude. However, in order to ensure a fair comparison of circuit run times the absolute gate times need to be weighted by the finite coherence time and lifetime of qubits and other levels involved in the gate mechanisms.

For neutral atoms, we take both the coherence time of the qubit states, given by the dephasing time $T_2^* = 4 \text{ ms}$ [39], and the lifetime of the Rydberg state, $T_{\text{Ryd}} = 150 \mu\text{s}$ [39], as typical error time scales against which we compare circuit run times. We take the former as reference for single-qubit gates, since they do not occupy the Rydberg level [42], and the latter for two-qubit gates, where controlled transitions via Rydberg levels constitute the primary gate mechanism [43]. For superconducting circuits, we take their intrinsic T_1 time, $T_1 = 15 \mu\text{s}$

TABLE I. Overview of different gate sets and corresponding gate times on neutral atoms and superconducting circuits. While the row “standard gate set” (SGS) represents typical gates and times used on the respective platform and are taken from literature, the multi-qubit gates in the row named “QSL gate set” (QGS) correspond to an extended gate set with gate times at the QSL. These times are determined within this study. The columns labeled as “time (rel.)” indicate the gate times weighted by the corresponding dissipation/decoherence times for the respective platforms (see main text for explanation).

	Neutral atoms			Superconducting circuits		
	gate	time (ns)	time (rel.)	gate	time (ns)	time (rel.)
Standard gate set (SGS)	local	~ 1000 [20]	$\sim 2.5 \times 10^{-4}$	local	25 [37]	1.7×10^{-3}
	CZ	350 [20]	2.3×10^{-3}	Sycamore	12 [37]	8.0×10^{-4}
QSL gate set (QGS)	local	~ 1000 [20]	$\sim 2.5 \times 10^{-4}$	local	25 [37]	1.7×10^{-3}
	CNOT ^a	300	2.0×10^{-3}	CNOT	14	9.3×10^{-4}
	CZ	350	2.3×10^{-3}	CZ	10	6.7×10^{-4}
	SWAP	400	2.7×10^{-3}	SWAP	12	8.0×10^{-4}
	ZZZ	600	4.0×10^{-3}	ZZZ	24	1.6×10^{-3}
	ZZZZ	600	4.0×10^{-3}	ZZZZ	80	5.3×10^{-3}

^aThe CNOT gate is listed for completeness but not used in any circuits for neutral atoms, see Sec. V for details.

[8], as typical error time scale since it applies for both single- and two-qubit gates. However, note that these choices are rather conservative. For neutral atoms, any two-qubit gate dynamics will naturally also involve the qubit levels, which have much longer coherence times. Weighting two-qubit gates exclusively by the Rydberg lifetime thus overestimates the lifetime-induced error probability. For superconducting circuits, longer T_1 times up to $500 \mu\text{s}$ have already been reported [44].

B. Circuit times in the standard gate model

In the following, we calculate the circuit run times for QFTs and single QAOA steps in the SGM using gates from both the SGS and the QGS. As already outlined above, we assume both hardware platforms to consist of qubits arranged in 2D arrays with only nearest-neighbor connectivity (see Sec. III for details regarding the model). This requires replacing all gates between nonconnected qubits with gate sequences between physically connected qubits.

1. Standard gate set circuits

Using gates from the SGS, the circuit’s gate sequences consist of single-qubit gates plus CZ gates in case of neutral atoms and single-qubit gates plus Sycamore gates in case of superconducting circuits (cf. Table I). Since the circuit for the QFT does by construction not require any gate between non-neighboring qubits [45], we only replace the controlled-phase gate and SWAP gate by gate sequences from the SGS. The situation changes for the QAOA circuits, as it assumes an all-to-all connected architecture and finding a circuit representation with minimal depth while only requiring nearest-neighbor connectivity is likely NP-hard [46]. As a remedy, we use the open-source *pytket* compiler [16] to translate the QAOA’s bare quantum circuits to executable quantum circuits that are in agreement with the hardware’s nearest-neighbor connectivity and SGSs. We furthermore use the compiler’s quantum circuit optimization features to opti-

mize the circuits in order to minimize gate counts and circuit depths.

The squares in Figs. 1(a) and 1(c) show the resulting run times for quantum circuits corresponding to QFTs and QAOA steps, respectively, when executed on neutral atoms (orange) and superconducting circuits (purple). Within each line, the size of the problem instance increases from the lower left to the upper right. Note that for a fair comparison between platforms, all gate times have been weighted by each platform’s intrinsic error time scales as described previously. Circuits with run times significantly exceeding the platform’s intrinsic, level-dependent error time scales, highlighted by the gray area in Figs. 1(a) and 1(c), will most likely yield unreliable results. We observe that despite the different time scales for gates, the weighted circuit run times are almost identical for both platforms and for both the QFT and QAOA, see squares in Figs. 1(a) and 1(c), respectively. However, only problem instances with relatively small qubit numbers seem to be currently doable on both platforms.

While the circuit run time is one deciding factor in whether it is feasible on current NISQ hardware, this measure neglects the fact that every gate comes with an intrinsic error probability. Assuming gate errors on the order of 0.1–1%, which are realistic both for neutral atoms [20,40] and superconducting circuits [37], it is clear that also the gate count limits the feasibility of quantum circuits and lower gate counts are thus preferable. To this end, the squares in Figs. 1(b) and 1(d) show the number of two-qubit gates required for QFT and single-step QAOA circuits, respectively. Note that circuit representations in terms of the SGS require the same number of two-qubit gates both on neutral atom and superconducting qubit hardware. Hence, both cases are represented by a single line in Figs. 1(b) and 1(d).

2. QSL gate set circuits

The circles in Figs. 1(a) and 1(c) show the circuit run times for the same quantum algorithms and complexity levels as used for the squares but when the QGS is employed to generate circuit representations. The corresponding two-qubit

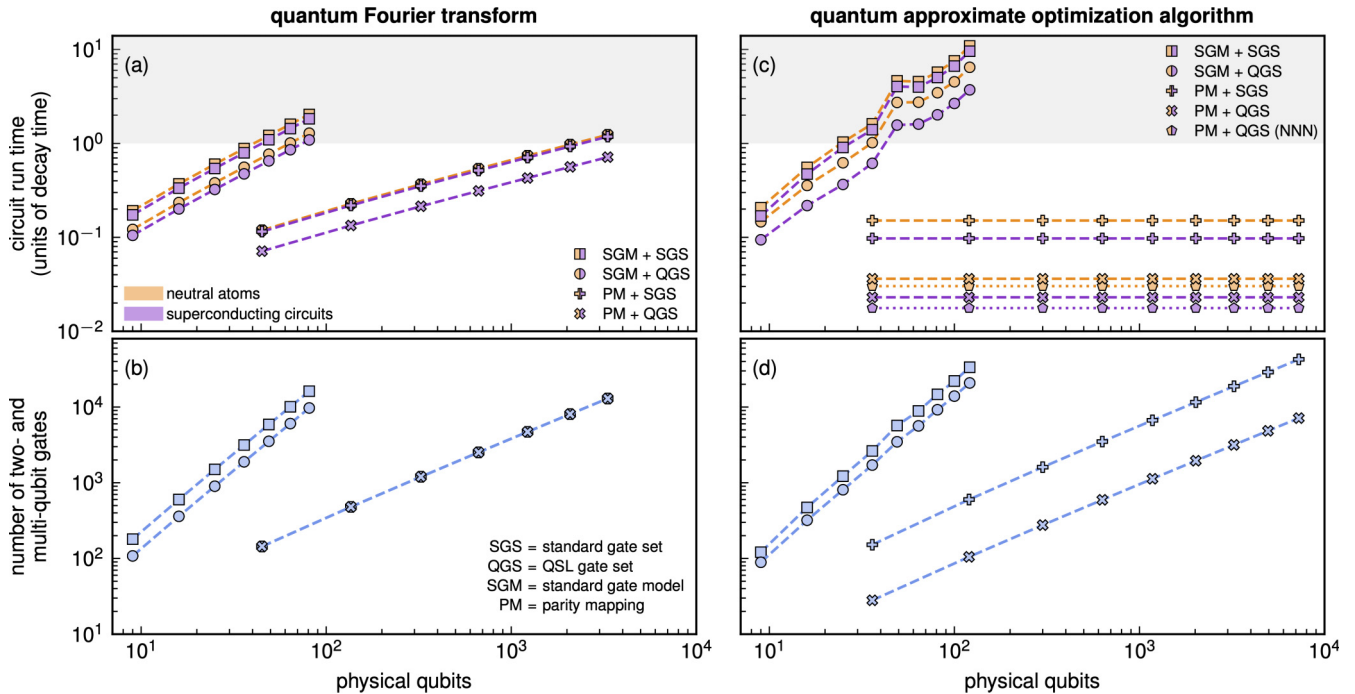


FIG. 1. Comparison of the circuit run times and gate counts for the quantum Fourier transform (QFT) and quantum approximate optimization algorithm (QAOA) between neutral atoms and superconducting circuits. The left (right) column corresponds to the QFT (QAOA) and the orange (purple) marker correspond to neutral atoms (superconducting circuits). The circuit run times for both algorithms and various problem instances of different size, $N = 9, 16, \dots, 81$ qubits for the QFT and $N = 9, 16, \dots, 121$ qubits for the QAOA, are given in panels (a) and (c), respectively. The numbers of two- and multi-qubit gates in the corresponding circuits are given in panels (b) and (d), respectively. The data for the squares [circles] is obtained using the standard gate model (SGM) using the gates and times from the standard gate set (SGS) [QSL gate set (QGS)], cf. Table I. In contrast, the plus signs [crosses] represent the results when the same algorithms are realized in the parity mapping (PM) using the SGS [QGS]. The pentagons correspond to an altered 2D architectures allowing for next-to-nearest-neighbor (NNN) coupling between qubits. The run time for each circuit is given in units of typical platform-specific error times. The gray area in panels (a) and (c) indicates where the circuit run times exceeds the error times.

gate counts are illustrated by the circles in Figs. 1(b) and 1(d). We observe an overall reduction in circuit run times and gate counts for both platforms, both algorithms and any considered problem instance. This improvement is a combined effect of using maximally fast quantum gates at the QSL and taking advantage of the increased flexibility provided by the extended set of gates. Especially the availability of SWAP gates needs to be stressed as it directly reduces the gate count and thus leads to a reduction in circuit run time even without an additional speedup in gate times. Due to the fact that all two-qubit gates in the QGS are by construction only between nearest-neighbors and have identical gate errors—due to our optimization—only the total number of gates matters. In order to quantify the improvements achievable through the QGS, Fig. 2(a) lists the average reduction in circuit run times and gate counts for each algorithm and platform.

Our analysis demonstrates to which extent circuit run times and gate counts can—from a theoretical perspective—still be improved if standard gates are replaced by an extended gate set with gate times at the QSL. However, even with the improved gate set and its reduction in run time and gate count, most of the quantum circuits, i.e., circles in Fig. 1, remain likely infeasible using current NISQ hardware. Quantum error correction codes could in principle address the issue of circuit run times and gate counts exceeding the limits set by finite

lifetimes and gate errors. Nevertheless, since both lifetimes [44,47] and qubit numbers [48] are constantly increasing, more complex quantum circuits will likely reach the feasible regime in the near future.

C. Circuits times in the parity mapping

Besides the representation of quantum algorithms using the SGM, we now consider the representation of the same algorithms within the PM. While the PM was originally designed to tackle combinatorial optimization problems via quantum annealing [27], it can also be utilized for digital quantum optimization algorithms such as QAOA [30] as well as to achieve universal quantum computing [26]. At its core, the PM circumvents the need for long-range interactions between qubits, which in turn renders gates between nonadjacent qubits obsolete. However, this comes at the expense of requiring more physical qubits and many-body constraints on 2×2 plaquettes of qubits as specified in detail in Appendix A 2. Similar to our analysis for the SGM (see Sec. II B) we use either gates from the SGS or the QGS in the following.

1. Standard gate set circuits

The plus signs in Fig. 1(a) [1(b)] show the circuit run times [two-qubit gate counts] of QFTs in the PM for the same

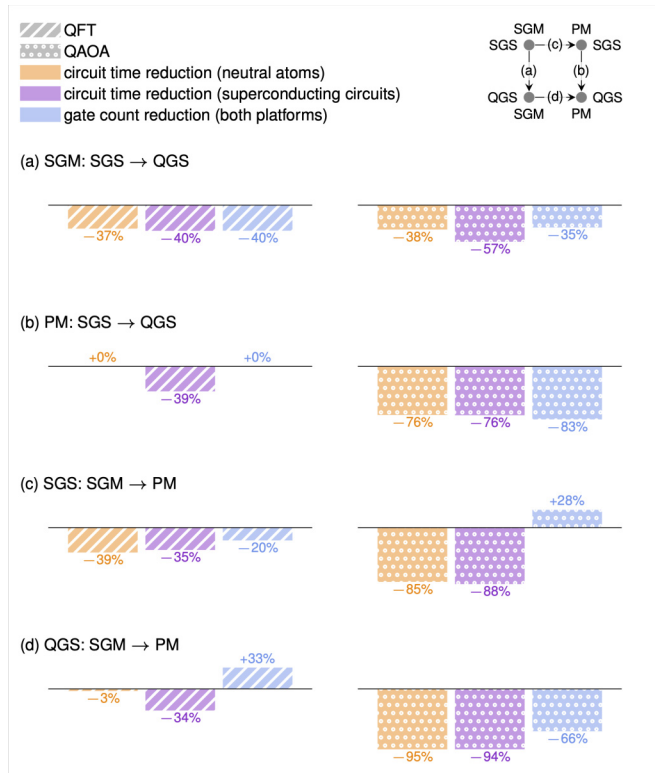


FIG. 2. Reduction in weighted circuit run times and (two- and multi-qubit) gate counts for the scenarios and data shown in Fig. 1. In panels (a) and (b) the circuits in the standard gate model (SGM) and parity mapping (PM) are compared when using the “QSL gate set” (QGS) instead of the “standard gate set” (SGS). In panels (c) and (d) the circuits in the SGS and QGS are compared when using the PM instead of the SGM. All presented numbers are the average over problem instances of different sizes, i.e., different numbers of qubits.

problem instances as used for the SGM [squares]. In both cases and for both platforms, the gates and times of the SGS have been used. We observe a reduction in circuit run times and gate counts for both platforms with the comparison being made with respect to the values for the SGM [see also Fig. 2(c)]. The same comparison can be done for the QAOA steps with their circuit run times and gate counts given by the plus signs in Figs. 1(c) and 1(d), respectively. Regarding single-step QAOA resource requirements, we observe a significant reduction in circuit run times [see also Fig. 2(c)] due to the constant circuit depth in the PM [29,30]. However, note that the circuit depth differs for neutral atoms and superconducting circuits due to the different coupling mechanisms between qubits with neutral atoms having the deeper circuits, cf. Appendix A 2. The reduction in circuit run time is accompanied by an increase in two-qubit gate counts compared to the SGM, which we attribute to the decomposition of the constraint gates, cf. Eq. (A6), into the natively available gates within each platform.

2. QSL gate set circuits

The crosses in Fig. 1 show the results for circuit representations in the PM when employing gates and times from

the QGS. For the QFT on neutral atoms, we do not observe any further reduction in circuit time or gate counts compared to its representation utilizing the SGS. This is because its circuit representations [49] for neutral atoms contain only single-qubit and CZ gates—gates for which the gate times are identical within the SGS and the QGS. In contrast, for superconducting circuits, we still observe an improvement in circuit time since the CZ gate becomes directly available in the QGS and must no longer be replaced by Sycamore and single-qubit gates. However, the representations in the SGS and QGS only differ by single-qubit gates and hence their two-qubit gate count is identical and also identical to that of the neutral atoms. This is reflected in Fig. 1(c) by a single line of superimposed plus signs and crosses.

For the single-step QAOA circuits in the PM, we observe a significant reduction both in circuit run time and gate counts on both platforms when the QGS is used instead of the SGS. This is due to the availability of the multi-qubit gates $ZZZ(\gamma)$ and $ZZZZ(\gamma)$, cf. Eq. (A6), where the gate-count reduction originates from avoiding single- and two-body gate decompositions. In addition, it turns out to be much faster to use control pulses that directly implement these multi-qubit gates as opposed to serially applying control pulses to implement the required single- and two-qubit gates. Figure 2(b) summarizes the average run time and gate count reductions (in terms of two- and multi-qubit gates) when replacing the SGS with the QGS within the PM.

In contrast to panels (a) and (b) of Fig. 2, where the gate set changes and the circuits stay in the SGM or PM, panels (c) and (d) examine the opposite scenario, i.e., the gate sets are kept constant but the circuit models change. In detail, Figs. 2(c) and 2(d) show the average circuit run time and gate count reduction when the SGM is replaced by the PM while the gate set is given by the SGS and QGS, respectively. Especially Fig. 2(d) needs to be emphasized as it reveals to which extent the PM allows to reduce circuit run times and gate counts compared to representations of the same circuits in the SGM. As mentioned above and detailed in Appendix A 2, the PM-specific improvements come at the only expense of requiring more qubits. For the current state of NISQ hardware and independent from the platform, we believe the SGM to be better suited for QFTs, as more problem instances seem to be feasible, judging by circuit run time and required qubits. In contrast, for QAOA using the QGS, we find the described parity representation of circuits (despite its overhead-induced inferior success probabilities compared to direct SGM-QAOA [50]) to be an advantageous option, as the run time and gate count is drastically reduced for each QAOA step compared to the SGM. In particular, the need for more PM-QAOA steps compared to SGM-QAOA steps might be compensatable given the resource reduction per PM-QAOA step. Given the ongoing upscaling of quantum computers in terms of qubit numbers, see, e.g., IBM’s quantum roadmap [48], the PM seems to be a viable option for QAOA.

As a final remark, it should be noted that we only discussed optimization problems with all-to-all connectivity, cf. Eq. (A1). However, many realistic optimization problems have rather sparse connectivity, which would lead to a

significant reduction in required qubits [28] while maintaining its strength of constant circuit depth.

III. MODELING NEUTRAL ATOM AND SUPERCONDUCTING CIRCUIT PLATFORMS

In Sec. II we have presented the main result of our paper—namely a calculation and comparison of circuit run times and gate counts using neutral atoms and superconducting circuits as quantum computing platform. In this section, we now introduce the detailed physical models used for both platforms. Since our focus is the description of dynamics, we focus on the respective Hamiltonians including the various control knobs typically available to steer the systems and implement quantum gates.

A. Neutral atoms

Arrays of trapped neutral atoms laser coupled to highly excited Rydberg states are a promising platform for quantum computing [51,52] and quantum simulation [53] as qubits can for example be encoded in long-lived hyperfine ground states. High-fidelity single-qubit gates can be achieved using microwave fields [54], two-photon Raman transitions [55], or a combination of microwaves and gradient fields for individual-qubit addressing [56–58]. In contrast, entangling operations between atoms, i.e., many-body gates, are typically realized via strongly interacting Rydberg levels [19,59] and various control schemes for two- and multi-qubit gates have been experimentally demonstrated [20,60–63].

Such gates have been used in recent experiments [39,40] with up to several hundreds of atoms arranged in a planar geometry. If we consider the smallest building block of such a 2D array, it consists of $N = 4$ atoms with each atom described by three relevant levels. In the rotating frame, its Hamiltonian reads ($\hbar = 1$)

$$\begin{aligned} H(t) = & - \sum_{n=1}^N \Delta_n(t) |r_n\rangle \langle r_n| + \sum_{\substack{n,m=1 \\ n < m}}^N V_{nm} |r_n r_m\rangle \langle r_n r_m| \\ & + \frac{1}{2} \sum_{n=1}^N \sum_{l=\uparrow,\downarrow} [\Omega_{l,n}(t) e^{i\varphi_{l,n}} |r_n\rangle \langle l_n| + \text{H.c.}], \quad (1) \end{aligned}$$

where $|\downarrow_n\rangle$ and $|\uparrow_n\rangle$ denote the qubit levels of the n th atom and $|r_n\rangle$ its Rydberg level. The Rabi frequencies of the two laser fields coupling the $|\downarrow_n\rangle$ and $|\uparrow_n\rangle$ states of the n th atom to their respective Rydberg level $|r_n\rangle$ are denoted by $\Omega_{\downarrow,n}(t)$ and $\Omega_{\uparrow,n}(t)$ [64]. The phases of the respective laser fields are given by $\varphi_{\downarrow,n}(t)$ and $\varphi_{\uparrow,n}(t)$ while $\Delta_n(t)$ denotes the detuning of the laser field from the exact transition frequency. We assume $\Delta_n(t)$ to be identical for both fields. The van der Waals (vdW) interaction between atom n and m is denoted by V_{nm} . Since V_{nm} is proportional to $1/R_{nm}^6$ with R_{nm} the distance between atoms n and m , the interaction strength is quickly vanishing between non-neighboring atoms. All considered gates are assumed to work in the Rydberg blockade regime (i.e., $\Omega/V \ll 1$) in order to avoid errors from populating states featuring multiple Rydberg excitations. In addition, assuming a fixed geometrical atom placement together with fixed maximal laser Rabi frequencies and fixed choice of Rydberg state throughout

algorithm execution, the two-qubit gate duration slows down drastically with increasing interqubit distance for reasonable experimental parameters (see below). As a consequence, in this scenario it is always faster to execute several gates among nearest neighbors as opposed to include non-nearest-neighbor two-qubit gates. Moreover, as our main objective is to find lower bounds on the algorithm run times we focus solely on two-qubit gates between nearest neighbors in Sec. II. However, we want to emphasize, that larger connectivities offered by neutral atom devices can be exploited also in this framework by dedicated multi-qubit gates like the considered four-qubit gate. Note that in Hamiltonian (1) we have omitted fields responsible for single-qubit gates, i.e., fields coupling the two qubit levels, since our focus is on multi-qubit gates for which the Rydberg levels and their interaction are the primary gate mechanism.

If not stated otherwise, we assume the laser fields to act identically on all atoms, i.e., $\Omega_{\downarrow,n}(t) = \Omega_{\downarrow}(t)$, $\Omega_{\uparrow,n}(t) = \Omega_{\uparrow}(t)$ and $\varphi_{\downarrow,n}(t) = \varphi_{\downarrow}(t)$, $\varphi_{\uparrow,n}(t) = \varphi_{\uparrow}(t)$ and $\Delta_n(t) = \Delta(t)$ for all n . In the following, we moreover assume by default a pseudo-2D architecture with $V_{nm} = V$ as described in more detail in Ref. [65]. Effectively, this corresponds to an architecture with equally strong nearest-neighbor (NN) and next-to-nearest-neighbor (NNN) coupling. However, it should be noted that we do not exploit the NNN couplings when constructing the quantum circuits in Sec. II. It thus does not affect any of the single- or two-qubit gate times presented in Table I as these times would be identical for an actual, planar 2D architecture where V is simply the coupling strength between nearest neighbors. The assumption of the pseudo-2D architecture does only affect the multi-qubit constraint gates $ZZZ(\gamma)$ and $ZZZZ(\gamma)$, cf. Eq. (A6), which are—within our study—only relevant for the PM using the QGS, cf. the crosses in Fig. 1(c).

B. Superconducting circuits

Qubits encoded in the lowest energy levels of superconducting circuits are another promising platform for quantum computing [66]. Since their parameters can be to some extent chosen during their fabrication process, superconducting circuits come in various variants and parameter regimes with transmon qubits [67] being currently the most prominent ones for quantum computing. Their qubit levels can be manipulated via microwave fields to either implement single-qubit gates [68,69] or two-qubit gates [70,71] with high fidelity. In contrast to neutral atom qubits, which interact directly via their Rydberg levels, qubits encoded in superconducting circuits often interact indirectly via intermediate coupling elements [72]. In an architecture, where such couplers are made tunable [73], the effective interaction strength between the qubits becomes tunable as well. Such qubit architectures have been successfully used in recent quantum advantage experiments [8,9] and thus are a prototypical NISQ quantum computing platform.

The qubits in such a tunable coupler architecture are arranged on a 2D lattice with nearest-neighbor couplings. In the following, we take the architecture from Ref. [8] as a reference. The smallest building block within such a system consists of $N = 4$ qubits. The Hamiltonian for this subsystem

in a frame rotating with frequency ω_{rot} reads [8]

$$\begin{aligned} \mathbf{H}(t) = & \sum_{n=1}^N \left[(\omega_n(t) - \omega_{\text{rot}}) \mathbf{b}_n^\dagger \mathbf{b}_n - \frac{\alpha_n}{2} \mathbf{b}_n^\dagger \mathbf{b}_n^\dagger \mathbf{b}_n \mathbf{b}_n \right] \\ & + \sum_{n=1}^N \frac{1}{2} [\mathbf{b}_n (\bar{\Omega}_{n,\text{re}}(t) + i\bar{\Omega}_{n,\text{im}}(t)) \\ & + \mathbf{b}_n^\dagger (\bar{\Omega}_{n,\text{re}}(t) - i\bar{\Omega}_{n,\text{im}}(t))] + \mathbf{H}_{\text{int}}(t), \end{aligned} \quad (2)$$

with

$$\begin{aligned} \mathbf{H}_{\text{int}}(t) = & g_{12}(t) (\mathbf{b}_1^\dagger \mathbf{b}_2 + \mathbf{b}_1 \mathbf{b}_2^\dagger) + \frac{g_{12}^2(t)}{|\eta|} \mathbf{b}_1^\dagger \mathbf{b}_1 \mathbf{b}_2^\dagger \mathbf{b}_2 \\ & + g_{23}(t) (\mathbf{b}_2^\dagger \mathbf{b}_3 + \mathbf{b}_2 \mathbf{b}_3^\dagger) + \frac{g_{23}^2(t)}{|\eta|} \mathbf{b}_2^\dagger \mathbf{b}_2 \mathbf{b}_3^\dagger \mathbf{b}_3 \\ & + g_{34}(t) (\mathbf{b}_3^\dagger \mathbf{b}_4 + \mathbf{b}_3 \mathbf{b}_4^\dagger) + \frac{g_{34}^2(t)}{|\eta|} \mathbf{b}_3^\dagger \mathbf{b}_3 \mathbf{b}_4^\dagger \mathbf{b}_4 \\ & + g_{41}(t) (\mathbf{b}_4^\dagger \mathbf{b}_1 + \mathbf{b}_4 \mathbf{b}_1^\dagger) + \frac{g_{41}^2(t)}{|\eta|} \mathbf{b}_4^\dagger \mathbf{b}_4 \mathbf{b}_1^\dagger \mathbf{b}_1, \end{aligned} \quad (3)$$

and

$$\bar{\Omega}_{n,\text{re}}(t) = \Re\{\Omega_n(t) e^{-i(\omega_{\text{rot}} - \bar{\omega}_n(t))t}\}, \quad (4a)$$

$$\bar{\Omega}_{n,\text{im}}(t) = \Im\{\Omega_n(t) e^{-i(\omega_{\text{rot}} - \bar{\omega}_n(t))t}\}. \quad (4b)$$

In the above equations, $\omega_n(t)$ is the frequency-tunable level splitting of transmon n , α_n its anharmonicity and \mathbf{b}_n its annihilation operator. The tunable coupling between transmons n and m is denoted by $g_{nm}(t)$ and η is the nonlinearity of the involved transmons, which is roughly constant $\alpha_n \approx \eta$ for all n . The time-dependent amplitude and frequency of a local X -type control field on transmon n , e.g., some microwave field, is denoted by $\Omega_n(t)$ and $\bar{\omega}_n(t)$, respectively.

While $\Omega_n(t)$ and $\bar{\omega}_n(t)$ are the actual physical control fields, we may take $\bar{\Omega}_{n,\text{re}}(t)$ and $\bar{\Omega}_{n,\text{im}}(t)$ as auxiliary control fields, which capture the time-dependent nature of $\Omega_n(t)$ and $\bar{\omega}_n(t)$ in the rotating frame and the latter can be reobtained from $\bar{\Omega}_{n,\text{re}}(t)$ and $\bar{\Omega}_{n,\text{im}}(t)$.

IV. DETERMINING QUANTUM SPEED LIMITS VIA QUANTUM OPTIMAL CONTROL

After having introduced the physical model for neutral atoms and superconducting circuits in Sec. III, we now review the basic notion of QSLs in Sec. IV A and of quantum optimal control theory in Sec. IV B since both build the theoretical, respectively methodical, foundation for the results presented in this work. Our method is described in Sec. IV C.

A. Quantum speed limits

The notion of quantum speed limits (QSLs) naturally arises in the context of quantum control problems. To this end, let us consider a quantum system described by the Hamiltonian $\mathbf{H}(t) = \mathbf{H}(\{\mathcal{E}_k(t)\})$, which depends on a set of control fields, $\{\mathcal{E}_k(t)\}$, that can be externally tuned, e.g., by the time-dependent amplitudes, phases or detunings in Eqs. (1) or (2). A quantum control problem is then defined by a set of initial states, $\{|\psi_l^{\text{in}}\rangle\}$, that should be transferred into a set of target

states, $\{|\psi_l^{\text{tgt}}\rangle\}$,

$$|\psi_l^{\text{tgt}}\rangle = \mathbf{U}(T, 0; \{\mathcal{E}_k(t)\}) |\psi_l^{\text{in}}\rangle, \quad \forall l, \quad (5)$$

where $\mathbf{U}(T, 0; \{\mathcal{E}_k(t)\})$ is the system's time-evolution operator and T the total time. Any choice of $\{\mathcal{E}_k(t)\}$, which fulfills Eq. (5) is considered a solution to the control problem. It is important to note that solutions to quantum control problems are usually not unique. Even for a fixed protocol duration T there typically exist many and often infinite many solutions.

The QSL for a given control problem is defined by the shortest protocol duration T_{QSL} for which at least one solution exists, i.e., for which at least one set of control fields $\{\mathcal{E}_k(t)\}$ exists that fulfills Eq. (5). In the context of quantum computing and the NISQ era, where time is a limited resource due to decoherence, it is desirable to implement quantum gates at the QSL. In order to calculate T_{QSL} analytically, $\mathbf{U}(T, 0; \{\mathcal{E}_k(t)\})$ must be analytically calculable for *any* set $\{\mathcal{E}_k(t)\}$ of conceivable control fields—a requirement that typically limits an analytical calculation of T_{QSL} to simple systems [74]. We would like to stress that in some cases, it is nevertheless possible to solve the entire control problem analytically, yielding not only T_{QSL} but also the time-optimal control fields $\{\mathcal{E}_k(t)\}$. A popular method to achieve that is, for instance, the quantum brachistochrone formalism [75], although its application has been mostly limited to simple systems [76–78].

Besides an analytical determination of T_{QSL} , there are various methods to approximate it. One prominent method is to calculate a lower bound $T_{\text{bound}} \leq T_{\text{QSL}}$ in order to get an estimate for T_{QSL} itself. For the simplest case of a state-to-state control problem, lower bounds can be calculated analytically [24]. In contrast, in case of multiple pairs of initial and final states, which describe the implementation of quantum gates, such lower bounds only exist for very simple systems [79]. In most cases, one needs to resort to numerical tools for estimating T_{QSL} . In that context, quantum optimal control theory has proven to be very useful [80] as it not only estimates T_{QSL} quite accurately but additionally yields the control fields that implement the desired dynamics, i.e., realizing the transition from initial to target states, cf. Eq. (5). Since this is our method of choice, we introduce it in more detail in the following.

B. Quantum optimal control theory

Quantum optimal control theory (OCT) [81] is a toolbox providing analytical and numerical tools that allow to derive optimized control fields, which solve a given control problem, e.g., in shortest time or with minimal error. Mathematically, an optimal control problem is formulated by introducing the cost functional

$$\begin{aligned} J[\{\psi_l\}, \{\mathcal{E}_k\}, T] \\ = \varepsilon_T[\{\psi_l(T)\}] + \int_0^T J_l[\{\psi_l(t)\}, \{\mathcal{E}_k(t)\}, t] dt, \end{aligned} \quad (6)$$

where $\{\psi_l(t)\}$ is a set of time-evolved states and $\{\mathcal{E}_k(t)\}$ a set of control fields to be optimized. The error measure ε_T quantifies the distance between the time-evolved states $|\psi_l(T)\rangle = \mathbf{U}(T, 0; \{\mathcal{E}_k(t)\}) |\psi_l^{\text{in}}\rangle$ and the desired target states $|\psi_l^{\text{tgt}}\rangle$ at the protocol's final time T , cf. Eq. (5). The term J_l in Eq. (6) captures time-dependent running costs. In most cases,

the error measure ε_T is the crucial figure of merit. In order to optimize for quantum gates, we use the error measure [82]

$$\varepsilon_T[\{\psi_l(T)\}] = 1 - \frac{1}{N_{\text{trgt}}} \sum_{l=1}^{N_{\text{trgt}}} \Re\{\langle \psi_l^{\text{trgt}} | \psi_l(T) \rangle\} \quad (7)$$

and take $|\psi_l^{\text{trgt}}\rangle = \mathcal{O}|\psi_l^{\text{in}}\rangle$ as the desired target states for the target gate \mathcal{O} . The set $\{\psi_l^{\text{in}}\}$ runs over the N_{trgt} logical basis states affected by \mathcal{O} with $N_{\text{trgt}} = 4, 8, 16$ for two-, three- or four-body gates, respectively.

Since the cost functional J is formulated such that smaller values correspond to better solutions of the control problem, solving an optimal control problem becomes essentially a minimization task, i.e., to find a set of control fields $\{\mathcal{E}_k^{\text{opt}}(t)\}$ that minimizes J , respectively ε_T . This is an optimization problem for which several numerical algorithms have been developed [83–87]. Many of them are readily available in open-source software packages [88–92].

C. Optimization procedure and method

In the context of determining T_{QSL} numerically via OCT, we search for the shortest protocol duration T for which the error measure ε_T , cf. Eq. (7), is still sufficiently small. In mathematical terms, we therefore define an error threshold ε_{max} and search for the shortest time T for which

$$\min_{\{\mathcal{E}_k\}}[\varepsilon_T[\{\psi_l(T)\}]] \leq \varepsilon_{\text{max}} \quad (8)$$

has a solution. However, the minimization over all conceivable control fields can not be done numerically—as there are infinitely many fields to check—and thus must be replaced by a sampling over finitely many fields in practice. In order to explore the function space efficiently by finite sampling, optimization algorithm as described in Sec. IV B can be used. Since we are interested in the fundamental QSL, we put only minimal limitations—apart from physically motivated limitations on amplitudes—on the form of each control field $\mathcal{E}_k(t)$. Hence, we need an optimization algorithm that is capable of exploring a function space of almost arbitrary field shapes.

As our method of choice we use Krotov’s method [93], a gradient-based optimization algorithms for time-continuous control fields. Due to the method’s built-in monotonous convergence [85], we deem it fitting for determining the QSL, although any other gradient-based optimization algorithm might work similarly well. While a more detailed description of Krotov’s method is given in Appendix B, its basic working principle is outlined in the following. It consists of an iterative update of the control fields $\{\mathcal{E}_k(t)\}$. Starting from a set of guess control fields $\{\mathcal{E}_k^0(t)\}$, Krotov’s method updates them until either $\varepsilon_T \leq \varepsilon_{\text{max}}$ or a maximum number of iterations is reached. This procedure can be viewed as a local but structured search within the space of all conceivable sets of control fields—the so-called control landscape. The locally searched area is thereby determined by the choice of the guess fields $\{\mathcal{E}_k^0(t)\}$, which set the initial starting point of the search. While the local nature of this search might appear contradicting to the global search required for evaluating Eq. (8), it can be turned into an approximate global search by using various sets of randomized guess fields. The combined effect of all these

local searches “cover” a larger fraction of the control landscape. The total procedure for approximating T_{QSL} using this method is thus to start with a protocol duration T for which the optimization algorithm finds solutions, i.e., optimized fields giving rise to $\varepsilon_T \leq \varepsilon_{\text{max}}$, and then to consecutively lower T until none of the various sets of randomized guess fields finds a solution anymore. Appendix C summarizes the details regarding the generation of random guess fields as well as each field’s parametrization within Krotov’s method.

Similar application of numerical optimal control techniques have previously shown excellent agreement with analytically provable QSLs [94,95]. At worst, this method could overestimate the actual QSLs, in which case the actual QSLs would be even smaller and the circuit and gate times in Sec. II that use the QGS would be even better.

V. BENCHMARKING QUANTUM GATE TIMES ON 2D ARCHITECTURES

In this final section, we now present the detailed results regarding the QSLs obtained via methods described in Sec. IV B for the various gates listed in Table I and have been used to calculate the circuit run times in Sec. II.

A. Neutral atoms

In this section, we determine the QSLs for different gates on neutral atoms by utilizing the control knobs available in Hamiltonian (1). To this end, we set the vdW interaction strength between Rydberg levels to $V/2\pi = 40$ MHz and assume a maximally achievable Rabi frequency of $\Omega_{\text{max}}/2\pi = 0.1V = 4$ MHz for both $\Omega_{\downarrow}(t)$ and $\Omega_{\uparrow}(t)$ as well as $\Delta_{\text{max}}/2\pi = 0.3V = 12$ MHz for $|\Delta(t)|$. These parameters are in the same regime than those reported in recent experiments [20,39,96].

The markers in Fig. 3 (left column) show the achievable gate error ε_T , cf. Eq. (7), for various gates and various gate times T on neutral atoms. Each individual marker thereby indicates the result of a single optimization with Krotov’s method [cf. Appendix (B)] i.e., the final error ε_T after 1500 iterations when starting from random guess fields generated via Eq. (C1). In the following we set $\varepsilon_{\text{max}} = 10^{-3}$ for all gates and stop any optimization as soon as this threshold is reached.

1. Two-qubit gates

Figure 3(a) shows the results for a CZ gate for three different, paradigmatic configurations of control fields. The circles correspond to the “parallel” configuration where all five possible control fields $\Omega_{\downarrow}(t)$, $\Omega_{\uparrow}(t)$, $\varphi_{\downarrow}(t)$, $\varphi_{\uparrow}(t)$, and $\Delta(t)$ have been optimized. In contrast, the crosses correspond to the “phase” configuration, where only $\varphi_{\uparrow}(t)$ has been optimized while $\Omega_{\uparrow}(t) = \Omega_{\text{max}}$, $\Delta(t) = \Delta_{\text{max}}$, and $\Omega_{\downarrow}(t) = \varphi_{\downarrow}(t) = 0$ have been kept fixed. In both cases, we obtain a QSL of $T_{\text{QSL}}^{\text{CZ}} = 350$ ns. The third field configuration, which we called “sequential” configuration (diamonds), consists of a sequential use of $\Omega_{\downarrow}(t)$, $\varphi_{\downarrow}(t)$ and $\Omega_{\uparrow}(t)$, $\varphi_{\uparrow}(t)$, i.e., the first half of the protocol we have $\Omega_{\uparrow}(t) = \varphi_{\uparrow}(t) = 0$ and in the second half $\Omega_{\downarrow}(t) = \varphi_{\downarrow}(t) = 0$. This configuration is inspired by an adiabatic protocol for implementing $\text{ZZZZ}(\gamma)$ gates, cf.

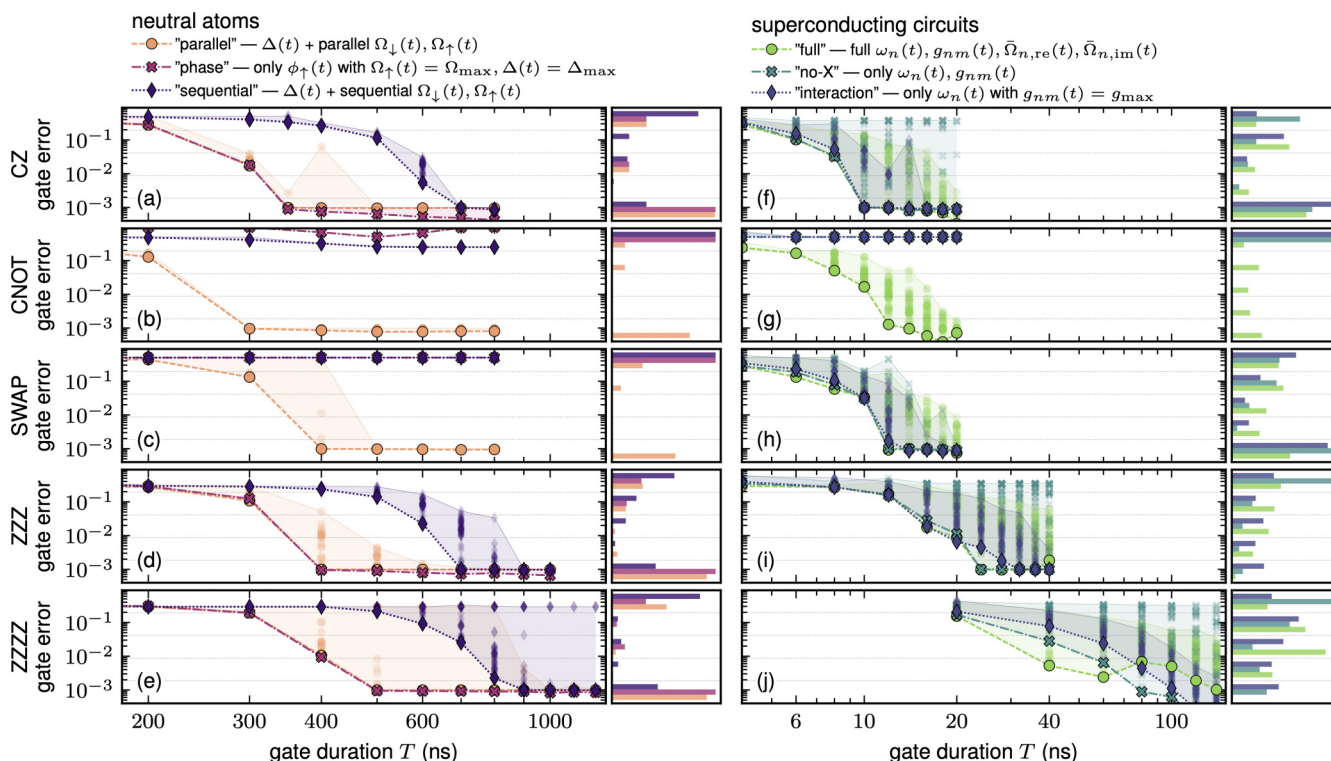


FIG. 3. Overview of the QSLs for the various gates of the “QSL gate set” (QGS) in Table I. The left (right) column shows the results for neutral atoms (superconducting circuits) for three different configurations of control fields (specified in the main text), respectively. Each marker represents the gate error ε_T , cf. Eq. (7), after either reaching $\varepsilon_T \leq \varepsilon_{\max} = 10^{-3}$ or 1500 iterations of Krotov’s method, cf. Appendix B, after starting from a set of random guess fields generated via Eq. (C1). While the lines connect the lowest errors reached for each gate time T within a given field configuration, the shaded background color indicates the range between the lowest and highest error. The marker density is shown at the right side of each panel as a histogram. The parameters for neutral atoms are $V/2\pi = 40$ MHz, $0 \leq \Omega_{\downarrow}(t)$, $\Omega_{\uparrow}(t) \leq \Omega_{\max} = 0.1V$ and $|\Delta(t)| \leq \Delta_{\max} = 0.3V$. The parameters for superconducting circuits are -40 MHz $\leq g_{nm}(t)/2\pi \leq 5$ MHz, 6700 MHz $\lesssim \omega_n(t) \lesssim 7100$ MHz (exact values depending on n and taken from Ref. [8]) and -50 MHz $\leq \bar{\Omega}_{n,\text{re}}(t)$, $\bar{\Omega}_{n,\text{im}}(t) \leq 50$ MHz. The anharmonic ladder for each transmon has been truncated after five levels with population in the highest level suppressed during optimization. The optimization results for $ZZZ(\gamma)$ and $ZZZZ(\gamma)$, cf. Eq. (A6), have been obtained for the maximally entangling gate at $\gamma = \pi/4$. The results for the CNOT gate on neutral atoms have been obtained using site-dependent control fields for each atom.

Eq. (A6) and Ref. [65]. Its QSL $T_{\text{QSL}}^{\text{CZ}} = 700$ ns is twice as long compared to the other two configurations.

In order to compare the results from the three configurations in terms of how successful the optimization has been in finding solutions, the histogram on the right side of Fig. 3(a) provides the probability density for obtaining final errors ε_T within certain ranges. For obtaining a solution with errors $\varepsilon_T \leq \varepsilon_{\max}$, we find the lowest probability for the “sequential” configuration (diamonds)—coinciding with the highest QSL—and the highest and almost identical probability for the other two configurations. Among those two, the “phase” configuration (crosses) needs to be emphasized in particular. From a physical perspective, setting $\Omega_{\downarrow}(t)$ and $\varphi_{\downarrow}(t)$ to zero automatically ensures that $|\downarrow\downarrow\rangle$ is mapped onto itself—as required by the CZ gate. This is not automatically guaranteed by the other two configurations and the optimization needs to explicitly ensure it and therefore needs to solve a slightly more complex optimization problem. However, the advantage of having one basis state automatically mapped correctly does not translate into an advantage regarding the QSL of $T_{\text{QSL}}^{\text{CZ}} = 350$ ns or the reachable error in general. Interestingly, both the “parallel” and “phase” configurations yield the same

achievable lowest errors for each gate time T . This is visually highlighted by the lines connecting the lowest errors per T in Fig. 3(a). Moreover, it should be stressed that these errors ε_T are reached for almost every set of initial guess fields, i.e., independent of the initial starting point within the problem’s control landscape, and obtained independently for both configurations. This supports the conjecture that $T_{\text{QSL}}^{\text{CZ}} = 350$ ns is the actual QSL for a CZ gate and generally validates our method in determining the QSL. From an optimal control perspective, it is interesting to see that the flexibility originating from the extended set of available control fields in the “parallel” configuration can not be turned into an advantage in error or time compared to the “phase” configuration. From a practical perspective, the latter is advantageous for experimental realizations as it requires fewer physical resources.

While the three configurations discussed so far should only be viewed as examples, we did not find any configuration giving rise to faster CZ gates. Hence, we assume $T_{\text{QSL}}^{\text{CZ}} = 350$ ns to be the fundamental QSL across all configurations. A natural comparison for $T_{\text{QSL}}^{\text{CZ}}$ with a value from the literature would be the gate time from the analytical protocol introduced in Ref. [20], especially because it uses the same control fields as

the “phase” configuration to implement the gate. For our parameters, we find $T_{\text{lit}}^{\text{CZ}} \approx 340$ ns as analytical gate time, which we assume identical with our QSL $T_{\text{QSL}}^{\text{CZ}} = 350$ ns given the rather coarse sampling of gate times T in Fig. 3(a). However, it should be noted that the analytical protocol of Ref. [20] implements a CZ gate only up to local operations—operations that are already contained in our optimized gate protocols at the QSL. Nevertheless, for a fair comparison of analytical and QSL gate times, as needed in Sec. II, as well as for simplicity, we set both times to 350 ns in Table I.

The remaining panels (b)–(e) in the left column of Fig. 3 show the results for other gates from the QGS of Table I. The results for a CNOT gate are shown in panel (b). It is the only gate for neutral atoms (among those we considered) that requires individual instead of global fields, i.e., it is the only gate for which we did not assume $\Omega_{\downarrow,n}(t) = \Omega_{\downarrow}(t)$, $\Omega_{\uparrow,n}(t) = \Omega_{\uparrow}(t)$, $\varphi_{\downarrow,n}(t) = \varphi_{\downarrow}(t)$, $\varphi_{\uparrow,n}(t) = \varphi_{\uparrow}(t)$, and $\Delta_n(t) = \Delta(t)$ for all n but actually assume individual fields with unique field shapes directed at each atom. We nevertheless consider the same three field configurations as for the CZ gate in panel (a) but now applied to the individual fields $\Omega_{\downarrow,n}(t)$, $\Omega_{\uparrow,n}(t)$, $\varphi_{\downarrow,n}(t)$, $\varphi_{\uparrow,n}(t)$, and $\Delta_n(t)$ instead of their global versions. Even with this more general setting of control fields, we find only the “parallel” configuration (circles) to allow for the realization of a CNOT gate with a QSL of $T_{\text{QSL}}^{\text{CNOT}} = 300$ ns. In contrast, the “phase” and “sequential” configurations do not allow to realize a CNOT gate at all. These results demonstrate that CNOT gates can be implemented using exclusively the site-dependent laser couplings of the qubit and Rydberg levels and no control knobs for single-qubit gates. However, an experimentally more convenient option, requiring no full site-dependent control of coupling qubit to Rydberg states, is to realize CZ gates with global laser pulses and convert them into CNOTs via local operations. We nevertheless include the CNOT gate and its QSL for completeness in our analysis as well as in the QGS in Table I but exclude it from any quantum circuit for neutral atoms in Sec. II for the reasons just mentioned.

Figure 3(c) shows the results for a SWAP gate. Like for the CNOT gate, we find the “parallel” configuration—assuming again global fields that are identical for each atom—to be the only one capable of realizing a SWAP gate. We obtain $T_{\text{QSL}}^{\text{SWAP}} = 400$ ns as its QSL. The other two configurations are not capable of realizing SWAP gates. However, since the SWAP gate, in contrast to the CNOT gate, can be realized with global control fields, we believe it to be experimentally feasible and thus include it as a viable gate in the QGS in Table I.

2. Three- and four-qubit constraint gates

So far, we have discussed the results for two-qubit gates in Figs. 3(a)–3(c). These gates and their respective QSLs have been used in determining the quantum circuits and calculating the corresponding circuit run times in Sec. II—especially for those circuits in the SGM discussed in Sec. II B. In contrast, for QAOA circuits in the PM [29], circuit representations without two-qubit gates exist, e.g., when the required three- and four-qubit constraint gates $ZZZ(\gamma)$ and $ZZZZ(\gamma)$, cf. Eq. (A6), are available natively and thus must not be decom-

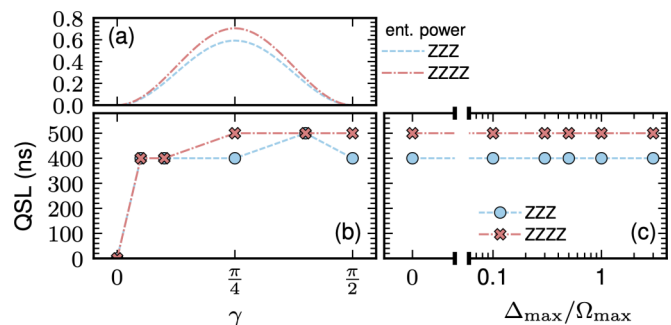


FIG. 4. Panel (a) shows the entanglement power [97] of the three- and four-qubit constraint gates $ZZZ(\gamma)$ and $ZZZZ(\gamma)$, cf. Eq. (A6), as a function of γ . In contrast, panels (b) and (c) examine the QSLs for these gates under various conditions on neutral atoms. In panel (b), the dependence of the QSL on the parameters γ is shown. In panel (c), the impact of $\Delta_{\text{max}}/\Omega_{\text{max}}$ is visualized. In the latter case, we have $\gamma = \pi/4$ and a fixed Ω_{max} while Δ_{max} is modified. The QSLs in panels (b) and (c) have been determined using the parameters and “phase” configuration described in Fig. 3.

posed into single- and two-qubit gates. In the following, we determine and discuss their QSLs.

It should first be noted that, from an algorithmic point of view, it is irrelevant whether, e.g., $ZZZ(\gamma)$ or $e^{i\alpha}ZZZ(\gamma)$, with α some arbitrary phase, is realized in experiments. The latter just changes the global phase of the quantum state during circuit execution. The same holds for the $ZZZZ(\gamma)$ gate. In both cases, we may choose α arbitrarily but, for practical reasons, choose it such that the states $|\downarrow\downarrow\downarrow\rangle$ and $|\downarrow\downarrow\downarrow\downarrow\rangle$ do not acquire a phase from the respective constraint gates. In the following, we therefore consider the phase-shifted constraint gates

$$e^{i\gamma}ZZZ(\gamma), \quad e^{-i\gamma}ZZZZ(\gamma), \quad (9)$$

instead of the ones from Eq. (A6).

In the context of QAOA circuits in the PM, these constraint gates need to be realized for various γ , cf. Eq. (A3). However, since we can not determine their QSL for each value of γ , we first analyze the gate’s entangling power [97] as a function of γ in Fig. 4(a). We observe maximal entangling power for $\gamma = \pi/4$ for both $ZZZ(\gamma)$ and $ZZZZ(\gamma)$ and thus decide to first benchmark their QSLs for that particular value as we expect the control problem in that case to be the hardest to solve and consequently the QSLs to be the largest.

Figures 3(d) and 3(e) show the optimization results for the phase-shifted constraint gates $ZZZ(\gamma)$ and $ZZZZ(\gamma)$, cf. Eq. (9), for $\gamma = \pi/4$, respectively. We use the same three configurations of control fields as for the two-qubit gates in panels (a)–(c) and find all configurations to be capable of implementing the constraint gates but observe the best performance for the “parallel” and “phase” configuration. While both configurations yield the same QSLs of $T_{\text{QSL}}^{\text{ZZZ}} = 400$ ns and $T_{\text{QSL}}^{\text{ZZZZ}} = 500$ ns for $ZZZ(\gamma)$ and $ZZZZ(\gamma)$, respectively, the “phase” configuration exhibits the better convergence behavior. Like for the CZ gate in Fig. 3(a), we observe every set of guess fields for this configuration to reliably converge towards the same final error ε_T for every T . Since the “parallel” configuration yields the same achievable errors as a

function of gate time T and uses by definition a different control strategy than the “phase” configuration, we believe to have reliably identified the QSLs for the constraint gates. In terms of experimental feasibility, the “phase” configuration is advantageous as it requires fewer hardware and control resources.

The fact that using $\varphi_{\uparrow}(t)$ as the only time-dependent control field suffices for implementing $ZZZ(\gamma)$ or $ZZZZ(\gamma)$ originates from considering the gates’ phase-shifted versions of Eq. (9). In detail, since the states $|\downarrow\downarrow\downarrow\rangle$ and $|\downarrow\downarrow\downarrow\downarrow\rangle$ are the only states among the 8 or 16 basis states of $ZZZ(\gamma)$ or $ZZZZ(\gamma)$ that technically require nonzero $\Omega_{\downarrow}(t)$ and $\varphi_{\downarrow}(t)$ in order to be phase configurable, we simply avoid this requirement by considering the gates’ phase-shifted version. For the remaining 7 or 15 basis states, which all have at least one atom initially in the $|\uparrow\rangle$ state, the phase $\varphi_{\uparrow}(t)$ together with a constant $\Omega_{\uparrow}(t) = \Omega_{\max}$ is sufficient to correctly adjust all phases. The “phase” configuration thus represents a hardware-efficient control scheme to realize fast $ZZZ(\gamma)$ and $ZZZZ(\gamma)$ gates.

Our results furthermore reveal that the “sequential” configuration, which was recently introduced in Ref. [65] and designed to implement high-fidelity $ZZZZ(\gamma)$ gates, seems not to be ideal when it comes to gate time as its configuration-specific QSL is roughly twice as long as the QSLs for the other configurations. However, it should be noted that a comparison of the protocol from Ref. [65] with protocols at the QSL is not a fair comparison. On the one hand, the control scheme of Ref. [65] is based on adiabaticity—a regime that we are far away from in our numerical calculations. On the other hand, while the parameter γ is a tunable variable in the adiabatic control scheme, the results in Figs. 3(d) and 3(e) are only valid for $\gamma = \pi/4$. If gates with a different γ are required, one explicitly needs to optimize control fields for that purpose. While it is beyond the scope of this paper to examine whether there exists an analytical control scheme with configurable γ at the QSL, in the following, we nevertheless provide an analysis of the QSLs for $ZZZ(\gamma)$ and $ZZZZ(\gamma)$ beyond $\gamma = \pi/4$. In detail, after having identified the “phase” configuration as the most reliable configuration to determine a gate’s QSL, we provide the QSLs for other γ values in Fig. 4(b). We find the QSLs to be almost constant for $\gamma > 0$ and only zero for $\gamma = 0$, in which case $ZZZ(\gamma)$ and $ZZZZ(\gamma)$ coincide with the identity operation. Interestingly, we do not observe a decrease in the QSLs for $\gamma = \pi/2$ in which case the constraint gates are no longer entangling, cf. Fig. 4(a), and should therefore theoretically be implementable with local operations only. We suspect that we do not see a decrease of the QSLs since we do not consider any control fields for local operations in Hamiltonian (1) and thus need to implement the local gates by means of the Rydberg levels.

We moreover analyze the dependence of the QSLs on the ratio $\Delta_{\max}/\Omega_{\max}$ in Fig. 4(c). Surprisingly, we observe the QSLs for both $ZZZ(\gamma)$ and $ZZZZ(\gamma)$ to be independent on this ratio and find $\Delta_{\max} = 0$ to be a viable option.

In general, we observe that the QSLs for $ZZZ(\gamma)$ and $ZZZZ(\gamma)$ are only slightly larger than those of the two-qubit gates. In view of quantum circuits for PM-QAOA, where such constraint gates are required, it is thus advantageous to have these gates natively available, since their representation via

single- and two-qubit gates [29] consumes significantly more time. This effect can be seen in Fig. 1(c), where the plus signs illustrate the data for constraint gates expanded in single- and two-qubit gates and the crosses for the usage of native constraint gates. One possible explanation for the short QSLs for $ZZZ(\gamma)$ and $ZZZZ(\gamma)$ compared to those of the two-qubit gates might be that the permutation symmetry of atoms within the pseudo-2D architecture [65], i.e., $V_{nm} = V$, matches the permutation symmetry of the gate operation itself.

At last, we therefore examine the impact of the pseudo-2D architecture onto the constraint gates $ZZZ(\gamma)$ and $ZZZZ(\gamma)$. For the three- and four-qubit constraint gates, the change to an actual, planar 2D architecture implies that while the couplings between nearest neighbors remain V , diagonal couplings, i.e., couplings between next-to-nearest neighbors or, in other words, qubits on opposite edges of a 2×2 square plaquette, are replaced by $V/8$. The stars in Figs. 5(a) and 5(b) show the corresponding results for $ZZZ(\gamma)$ and $ZZZZ(\gamma)$ gates and $\gamma = \pi/4$, respectively. While their QSLs within the pseudo-2D architecture are $T_{\text{QSL}}^{\text{ZZZ}} = 400$ ns and $T_{\text{QSL}}^{\text{ZZZZ}} = 500$ ns, they become $T_{\text{QSL,2D}}^{\text{ZZZ}} = T_{\text{QSL,2D}}^{\text{ZZZZ}} = 600$ ns in the actual, planar 2D architecture. Interestingly, this corresponds only to a relatively small increase in the QSLs for both gates. A possible explanation might be that the gate speed for neutral atoms is primarily determined by the maximal Rabi frequency Ω_{\max} , which is identical for both examples, and not so much by the interatomic interaction strength, which is different for both architectures.

Although we believe the pseudo-2D architecture to be viable in experiments due to the great flexibility to arrange neutral atoms [98], we nevertheless take the QSLs for the actual, planar 2D architecture to be the reference gate times within the QGS in Table I and Sec. II. However, recall that the constraint gates are—within our study—only relevant for the QAOA circuits in the PM, cf. Fig. 1(c). In order to nevertheless allow for a comparison of the run times in the actual, planar 2D architecture (orange crosses) with those using the pseudo-2D architecture, we add the latter as orange pentagons to Fig. 1(c).

B. Superconducting circuits

Similar to neutral atoms in Sec. V A, we now determine and analyze the QSLs for the same quantum gates but for superconducting circuits. The available control knobs to implement these gates are the tunable qubit frequencies $\omega_n(t)$, the tunable coupling strength $g_{nm}(t)$ between qubits and the (auxiliary) X-type local control fields $\bar{\Omega}_{n,\text{re}}(t)$ and $\bar{\Omega}_{n,\text{im}}(t)$ in Hamiltonian (2). In order to remain experimentally realistic, we take parameters from Ref. [8]. To this end, we single out a 2×2 plaquette consisting of four qubits from the generally larger 2D architecture. We take the qubit frequencies and their tunable range to be given by $6700 \text{ MHz} \lesssim \omega_n(t) \lesssim 7100 \text{ MHz}$ and their anharmonicities given by $\alpha_n \approx 200 \text{ MHz}$, with exact values depending on n . The tunable coupling strength is given by $-40 \text{ MHz} \leq g_{nm}(t) \leq 5 \text{ MHz}$, as reported in Ref. [8]. Moreover, we assume the (auxiliary) X-type control fields $\bar{\Omega}_{n,\text{re}}(t)$ and $\bar{\Omega}_{n,\text{im}}(t)$, which encode the physical X-type control fields $\Omega_n(t)$ and their tunable driving frequencies $\bar{\omega}_n(t)$, to satisfy $-50 \text{ MHz} \leq \bar{\Omega}_{n,\text{re}}(t), \bar{\Omega}_{n,\text{im}}(t) \leq 50 \text{ MHz}$.

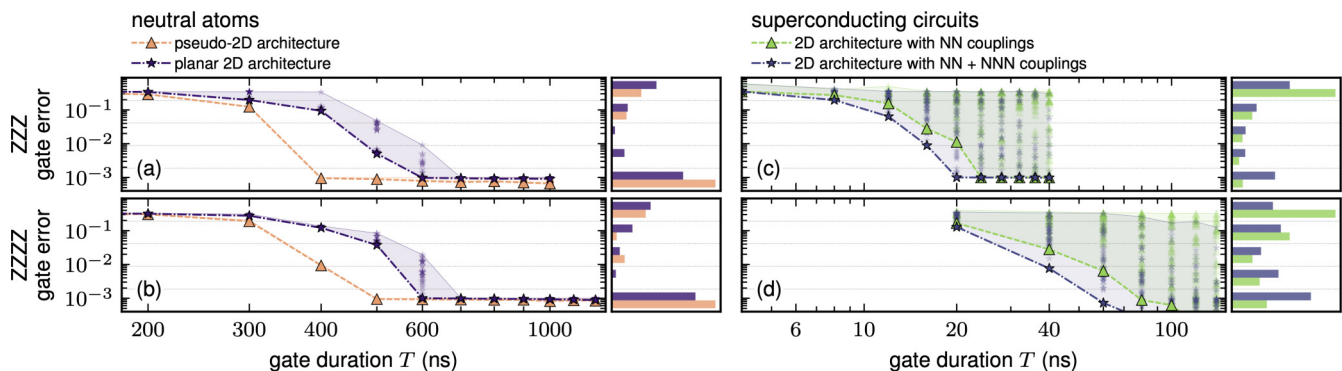


FIG. 5. Overview of different QSLs similar to Fig. 3 but exclusively for the three- and four-qubit constraint gates $ZZZ(\gamma)$ and $ZZZZ(\gamma)$, cf. Eq. (A6). The left column shows results for neutral atoms, obtained using the “phase” configuration of Fig. 3, for the pseudo-2D architecture (triangles) and an actual, planar 2D architecture (stars). The right column compares the results for superconducting circuits using the “no-X” configuration from Fig. 3. The data correspond to the physical architecture of Ref. [8] where only nearest-neighbor (NN) couplings between qubits are present (triangles) and where diagonal couplings, i.e., next-to-nearest-neighbor (NNN) couplings, are added (stars).

1. Two-qubit gates

Figure 3(f) shows results for a CZ gate on superconducting circuits using three different configurations of control fields—different, of course, from those used for neutral atoms. In the “full” configuration (circles), all available control fields, $\omega_n(t)$, $g_{nm}(t)$, $\bar{\Omega}_{n,\text{re}}(t)$, and $\bar{\Omega}_{n,\text{im}}(t)$, are time dependent and being optimized. In the “no-X” configuration (crosses), only $\omega_n(t)$ and $g_{nm}(t)$ are optimized while the X-type control fields are set to zero, $\bar{\Omega}_{n,\text{re}}(t) = \bar{\Omega}_{n,\text{im}}(t) = 0$. At last, in the “interaction” configuration only $\omega_n(t)$ is time dependent and optimized while $g_{nm}(t) = -40 \text{ MHz} = g_{\text{max}}$ is set to its maximum magnitude and $\bar{\Omega}_{n,\text{re}}(t) = \bar{\Omega}_{n,\text{im}}(t) = 0$. For the CZ gate in Fig. 3(f), we observe all three configurations to indicate the same QSL of $T_{\text{QSL}}^{\text{CZ}} = 10 \text{ ns}$. In terms of convergence behavior, the “interaction” configuration shows the best performance as indicated by the probability density on the right side of Fig. 3(f). In general, the three configurations show slightly worse convergence behavior than the three configurations for the neutral atoms, cf. Fig. 3(a). Nevertheless, since all three configurations indicate the same QSL and, in general, yield the same achievable error ε_T depending on T , we believe our method to determine the QSL to yield reliable results also for superconducting circuits and conjecture $T_{\text{QSL}}^{\text{CZ}} = 10 \text{ ns}$ to be the fundamental QSL for CZ gates. While there are reference implementations for CZ gates on similar architectures with tunable couplers [71,73,99,100], none of these architectures matches our architecture and parameter regime. Hence, we compare our QSL to the gate time of the fastest two-qubit gate, the Sycamore gate, reported in Ref. [8]. We find $T_{\text{QSL}}^{\text{CZ}} = 10 \text{ ns}$ to be slightly faster than $T_{\text{lit}}^{\text{Syc}} = 12 \text{ ns}$ [37].

In Fig. 3(g), the results for a CNOT gate are shown using the same three configuration as for the CZ gate. We find only the “full” configuration to be capable of realizing a CNOT gate, yielding the QSL $T_{\text{QSL}}^{\text{CNOT}} = 14 \text{ ns}$, while the other two configuration are not capable of it. Among the available control knobs, the X-type control fields $\bar{\Omega}_{n,\text{re}}(t)$ and $\bar{\Omega}_{n,\text{im}}(t)$ are crucial for a CNOT gate to be feasible. For the SWAP gate, we again find all three configuration to converge, cf. Fig. 3(h), with the “no-X” and “interaction” configuration showing the best convergence behavior. We find a QSL of $T_{\text{QSL}}^{\text{SWAP}} = 12 \text{ ns}$.

2. Three- and four-qubit constraint gates

We now turn towards the three- and four-qubit constraint gates $ZZZ(\gamma)$ and $ZZZZ(\gamma)$. However, note that in the following and in contrast to neutral atoms, we do not consider their phase-shifted versions, cf. Eq. (9), but their original versions, cf. Eq. (A6).

Figures 3(i) and 3(j) show the results for the constraint gates $ZZZ(\gamma)$ and $ZZZZ(\gamma)$, respectively. We use the same three configurations of control fields as for the two-qubit gates of panels (f)–(h). We observe the “interaction” configuration to have the best convergence behavior while the “no-X” configuration gives in both cases rise to the shortest QSLs of $T_{\text{QSL}}^{\text{ZZZ}} = 24 \text{ ns}$ and $T_{\text{QSL}}^{\text{ZZZZ}} = 80 \text{ ns}$ for $ZZZ(\gamma)$ and $ZZZZ(\gamma)$, respectively. Both QSLs have been determined for $\gamma = \pi/4$. The QSL for $ZZZ(\gamma)$ has thereby been confirmed independently by both the “full” and “no-X” configurations as both yield almost identical achievable errors ε_T as a function of gate time T . We thus assume the QSL $T_{\text{QSL}}^{\text{ZZZ}} = 24 \text{ ns}$ to be well backed up. The situation is different for $ZZZZ(\gamma)$, for which we observe very different convergence behaviors for the three configurations, cf. Fig. 3(j). While the “no-X” configuration yields the shortest QSL of $T_{\text{QSL}}^{\text{ZZZZ}} = 80 \text{ ns}$, the “full” configuration shows slightly better performance for $T < T_{\text{QSL}}^{\text{ZZZZ}}$, which might suggest that even shorter gate protocols for $ZZZZ(\gamma)$ may exist but our method did not find them due, e.g., limited numbers of guess fields for exploring the control landscape. In the following, as well as for the calculation of circuit run times in Sec. II, we nevertheless assume $T_{\text{QSL}}^{\text{ZZZZ}} = 80 \text{ ns}$ to be the QSL for the $ZZZZ(\gamma)$ gate as it is the fastest gate time T among the three configurations of control fields for which Krotov’s method was able to find a solution with $\varepsilon_T \leq \varepsilon_{\text{max}}$.

Interestingly, while we observe the QSLs for $ZZZ(\gamma)$ and $ZZZZ(\gamma)$ to be almost identical for neutral atoms and only slightly longer than the QSLs for the two-qubit gates, we observe the same only for the $ZZZ(\gamma)$ gate for superconducting circuits. The $ZZZZ(\gamma)$ gate has a much longer QSL compared to the other QSLs on that platform. To rigorously decide whether this is due to the nonideal convergence behavior observed in Fig. 3(j) or has some deeper physical origin

is beyond the scope of this study. In an attempt to tackle this question nevertheless, we consider the scenario of having additional diagonal couplings, i.e., next-to-nearest-neighbor couplings, among the transmons in the superconducting circuit architecture. For Hamiltonian (3), this implies adding two additional rows with couplings $g_{13}(t)$ and $g_{24}(t)$ in the same form as the already present couplings $g_{12}(t)$, $g_{23}(t)$, $g_{34}(t)$, and $g_{41}(t)$. This scenario is inspired by the pseudo-2D architecture for neutral atoms, which also exhibits identical nearest-neighbor (NN) and next-to-nearest-neighbor (NNN) couplings and where having these couplings is advantageous. Figures 5(c) and 5(d) show the results for the two cases with only NN couplings (triangles) and with NN plus NNN couplings (stars). Besides observing much better convergence properties for the latter case, we also obtain improved QSLs of $T_{\text{QSL,NNN}}^{\text{ZZZ}} = 20$ ns and $T_{\text{QSL,NNN}}^{\text{ZZZZ}} = 60$ ns for the $\text{ZZZ}(\gamma)$ and $\text{ZZZZ}(\gamma)$ gate, respectively. However, despite improvements in this scenario, the QSL of the $\text{ZZZZ}(\gamma)$ gate does not get close to the QSL of the $\text{ZZZ}(\gamma)$ gate as it does for neutral atoms. The presence of NNN couplings might therefore be just a partial explanation of the QSL differences for the constraint gates between neutral atoms and superconducting circuits. Despite the scenario with NNN couplings not reflecting the actual architecture for superconducting circuits, we nevertheless add the circuit run times for the QAOA circuits in the PM using these faster constraint gates to Fig. 1(c) for reference purposes as purple pentagons.

VI. CONCLUSIONS

In this study, we have determined the QSLs for several common two-qubit and two specific multi-qubit quantum gates for two promising quantum computing platforms that allow for a 2D arrangement of qubits—neutral atoms and superconducting circuits. We have used OCT to determine the QSLs, as it provides a generally applicable tool that warrants a fair comparison of both platforms.

On the level of individual quantum gates, our study allows assessing how close gate protocols from the literature are compared to their fundamental QSLs or, in other words, how much can (theoretically) still be gained in time if known gate protocols are replaced by numerically optimized ones. We find the QSLs for all investigated two-qubit gates, encompassing CNOT, CZ, and SWAP, to be very similar within each platform and close to the reference gate times for a CZ gate in case of neutral atoms [20] and a Sycamore gate in case of superconducting circuits [8].

On the level of quantum algorithms, our study has moreover allowed us to determine the “QSLs” for entire quantum circuits. To this end, we have assumed a 2D grid architecture for both platforms and qubit connectivity that allows physical quantum gates only between neighboring qubits. However, we have assumed these gates to be executable at the QSL. This has allowed us to calculate the circuit run times at the “QSL” for two paradigmatic quantum algorithms—the QFT and a single step of the QAOA. We find that the corresponding weighted circuit run times scale comparably with respect to the system size. Furthermore, we observe this to be independent of the chosen gate set used to translate the quantum algorithms into executable quantum circuits, i.e., independent

of whether the SGS or QGS is used. We observe platform independently that the QGS yields circuit run times and gate counts that are roughly half compared to those in the SGS. On the one hand, this demonstrates that further speedup of circuit run times is theoretically possible on both platforms. On the other hand, it also shows that both platforms perform equally well when running prototypical quantum circuits using typical, present-day NISQ hardware.

Besides a representation of the quantum circuits in the SGM, we have also explored the representation of the same quantum algorithms in the PM [26,27,29]. We observe a reduction in circuit run times as well as in gate counts in most cases. This reduction comes at the expense of requiring more physical qubits but without the need to change the geometrical layout or the control hardware. In this context, we want to specifically emphasize the circuit run times of a single QAOA step in the PM. Compared to its representation in the SGM, it offers a constant circuit depth independent on the problem complexity but requires the implementation of local three- and four-qubit constraint gates [29]. For superconducting circuits we find their gate times at the QSL to be roughly similar to the run times of their decompositions into single- and two-qubit gates. In contrast, for neutral atoms we find the direct implementation of the constraint gates to be only slightly slower than any single two-qubit gate and especially much faster than their decompositions into single- and two-qubit gates. We nevertheless observe for both platforms run times for a single QAOA step on the order of 2–4% of the platform’s intrinsic coherence time. While this corresponds to an improvement of one order of magnitude for a problem size with $N = 9$ logical qubits, this grows to an improvement of three orders of magnitude for $N = 121$. Since the gate counts, when the native implementations of the constraint gates are used, are also smaller in the PM compared to the SGM, we believe the PM to be advantageous in terms of the described resources for a single QAOA step. It minimizes errors due to finite coherence time and allows for large-depth QAOA. However, we want to emphasize that deeper PM-QAOA circuits are in general necessary to reach similar success probabilities compared to lower-depth SGM-QAOA implementations.

While we want to emphasize that the feasibility of a quantum circuit depends on more than just its run time and gate count, our benchmark study demonstrates that at least these two factors can be theoretically further improved using optimized gate protocols. This holds for both neutral atoms and superconducting circuits.

The data that supports the findings of this study are available from the corresponding author upon reasonable request.

ACKNOWLEDGMENTS

We would like to thank Hannes Pichler and Kilian Ender for helpful discussions and Michael Fellner for help with the parity QFT circuits. This work was supported by the Austrian Science Fund (FWF) through a START Grant under Projects No. Y1067-N27 and No. I 6011. This research was funded in whole, or in part, by the Austrian Science Fund (FWF) SFB BeyondC Project No. F7108-N38. For the purpose of open access, the author has applied a CC BY

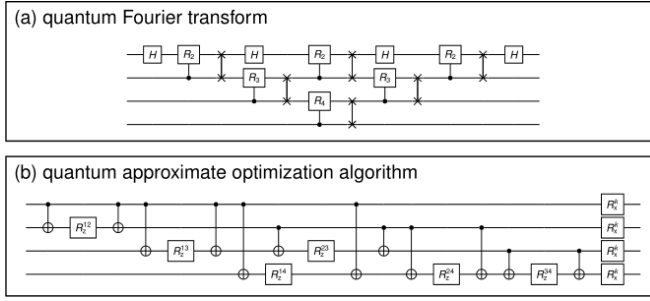


FIG. 6. Illustrations of the quantum circuits for a QFT (a) and a single QAOA step (b) for $N = 4$ qubits, respectively.

public copyright licence to any Author Accepted Manuscript version arising from this submission. This project was funded within the QuantERA II Programme that has received funding from the European Union’s Horizon 2020 research and innovation programme under Grant Agreement No. 101017733.

W.L. and D.B. conceived the project. D.B. and C.D. developed the models for neutral atoms and D.B. for superconducting circuits. D.B. performed all numerical simulations. All authors analyzed and discussed the data and contributed to the writing of the manuscript.

The authors declare no competing interests.

APPENDIX A: QUANTUM CIRCUITS FOR QFT AND QAOA

1. Standard gate model

The quantum Fourier transform (QFT) is a key ingredient in Shor’s algorithm for integer factorization [1] and thus a prototypical application for quantum computers. While a typical circuit representation of a QFT contains exclusively Hadamard gates, H , and controlled-phase gates, $R_n = \text{diag}\{1, 1, 1, \exp\{2\pi i/2^n\}\}$, a more efficient representation with fewer gates and lower circuit depths can be constructed using SWAP gates [45]. Figure 6(a) illustrates the circuit for a QFT with $N = 4$ qubits.

The quantum approximate optimization algorithm (QAOA) aims at finding approximate solution to combinatorial optimization problems [4]. For instance, let us consider the task to find the ground state of the N qubit spin glass Hamiltonian

$$H_z = \sum_{\substack{n,m=1 \\ n < m}}^N J_{nm} \sigma_z^{(n)} \sigma_z^{(m)}, \quad (\text{A1})$$

where J_{nm} denotes the interaction strength between qubits n and m and $\sigma_z^{(i)}$ the Pauli-Z operator on qubit i . The QAOA allows to find an approximate solution, i.e., an approximate ground state for Hamiltonian (A1), by applying the procedure

$$|\psi_{\text{out}}\rangle = \prod_{k=1}^p e^{-i\alpha_k H_x} e^{-i\beta_k H_z} |\psi_{\text{in}}\rangle, \quad H_x = \sum_{n=1}^N \sigma_x^{(n)}, \quad (\text{A2})$$

where $|\psi_{\text{in}}\rangle$ is the ground state of the so-called mixing Hamiltonian H_x and $\alpha_k, \beta_k \in [0, 2\pi)$ are angles—physically corresponding to evolution times—that are iteratively optimized via a classical, closed-loop feedback optimization and

the energy expectation value of $|\psi_{\text{out}}\rangle$ being the objective to minimize. The number of steps is denoted by p . A single step in the QAOA is thus given by the application of the spin glass or problem Hamiltonian H_z followed by the mixing Hamiltonian H_x . While the latter corresponds to a parallel application of α_k -dependent single-qubit X rotations, R_x^k , in the associated quantum circuit—and is thus negligible time-wise—the circuit implementation of the former requires multiple CNOT gates as well as single-qubit phase gates, $R_z^{nm} = \exp\{-i\beta_k J_{nm} \sigma_z\}$, containing information about J_{nm} and β_k . Figure 6(b) shows the quantum circuit for a single QAOA step for the Hamiltonian of Eq. (A1) and $N = 4$ qubits.

While Fig. 6 illustrates the ideal quantum circuits for the QFT and QAOA, they need to be converted into circuits, which are capable to run on existing quantum hardware—in our study either neutral atoms or superconducting circuits—and their connectivity and native gate sets. To this end, we use the *pytket* compiler [16] to convert the quantum circuits. The core steps of this process are the initial replacement of any non-native two-qubit gate with a circuit identity consisting of native gates, including gates between nonconnected qubits, and a subsequent reduction and optimization procedure. The latter squashes gates together wherever possible and uses known gate identities to further reduce the number of two-qubit gates. For details, we refer to Ref. [16].

2. Parity mapping

In the SGM, as described in Appendix A1, every logical qubit is given by exactly one physical qubit and gates on logical qubits are equivalent to gates on physical qubits. This is different for the PM, where $K > N$ physical qubits are required for a problem of N logical qubits and gates on logical qubits become different gates or even gate sequences on the physical qubits. However, due to the arrangement of physical qubits according to the PM, all gates between physical qubits are strictly local and thus require only nearest-neighbor connectivity.

For the QFT, we need $K = N(N + 1)/2$ physical qubits in the PM and find that Hadamard gates, H , on logical qubits become equivalent to several single- and two-qubit gates on neighboring physical qubits. In contrast, the logical two-qubit controlled-phase gates R_n between any pair for logical qubits is given by exactly three parallel single-qubit gates on physical qubits [26,49]. Hence, while the logical Hadamard gates require more resources in the PM, the logical controlled-phase gates require significantly less resources—especially for those gates where the logical qubits are far away from each other.

For the QAOA, we need $K = N(N - 1)/2$ physical qubits in the PM and Eq. (A2) becomes [29]

$$|\psi_{\text{out}}\rangle = \prod_{k=1}^p e^{-i\alpha_k H_x^{\text{phys}}} e^{-i\beta_k H_z^{\text{phys}}} e^{-i\gamma_k H_c} |\psi_{\text{in}}\rangle, \quad (\text{A3})$$

where

$$H_x^{\text{phys}} = \sum_{k=1}^K \tilde{\sigma}_x^{(k)}, \quad H_z^{\text{phys}} = \sum_{k=1}^K \tilde{J}_k \tilde{\sigma}_z^{(k)} \quad (\text{A4})$$

are the modified mixing and spin glass Hamiltonians in the PM, respectively. The local field strengths \tilde{J}_k run over the

$N(N-1)/2$ interactions J_{nm} , cf. Eq. (A1), and $\tilde{\sigma}_z^{(k)}$ encodes the parity of the corresponding two-qubit interaction $\sigma_z^{(n)}\sigma_z^{(m)}$ [27]. In order to solve optimization problems using the PM, we additionally need to constraint the dynamics to the 2^{N-1} dimensional subspace within the 2^K dimensional physical Hilbert space that corresponds to the 2^{N-1} eigenstates of Eq. (A1) that have unique eigenvalues [101]. Hence, since not every eigenstates of H_z^{phys} has a logical counterpart in H_z , it needs to be energetically penalized as it would not be a valid solution to the optimization problem. This is achieved by realizing $C = K - N + 1$ local three- and four-qubit constraints via [27]

$$H_c = \sum_{c=1}^C \tilde{\sigma}_z^{(k_1)} \tilde{\sigma}_z^{(k_2)} \tilde{\sigma}_z^{(k_3)} (\tilde{\sigma}_z^{(k_4)}), \quad (\text{A5})$$

where ‘‘local’’ refers to k_1, \dots, k_4 being nearest-neighbor physical qubits. For every QAOA step in Eq. (A3), this requires C constraint gates of the form (neglecting tildes and indices)

$$\text{ZZZ}(\gamma) = \exp\{-i\gamma\sigma_z\sigma_z\sigma_z\}, \quad (\text{A6a})$$

$$\text{ZZZZ}(\gamma) = \exp\{-i\gamma\sigma_z\sigma_z\sigma_z\sigma_z\}, \quad (\text{A6b})$$

with γ the effective constraint strengths, which—like in the original QAOA scheme of Eq. (A2)—are optimized in a classical, closed-loop feedback optimization. The gates in Eq. (A6) can be either realized directly [65] or by decomposing them into single- and two-qubit gates, e.g., by using four or six CNOT gates plus one γ -dependent single-qubit phase-gate for ZZZ(γ) or ZZZZ(γ), respectively [29].

It is important to note that the constraint gates are the only multi-qubit gates in Eq. (A3). Independent of N , their implementation can be parallelized with at most nine [65] or four [29] consecutive layers of constraint gates for neutral atoms or superconducting circuits, respectively. However, note that shallower circuits may be feasible by now [102]. The difference between the two platforms originates from their different qubit-qubit coupling mechanism. The tunable coupler architecture of superconducting circuits [8] allows to switch off the coupling between any pair of qubits. As a consequence,

all constraint plaquettes that do not share a common qubit, i.e., next-to-nearest-neighbor plaquettes, can be implemented in parallel. A single QAOA step thus requires at most four layers of constraint gates to realize all of them. Neutral atoms, in contrast, interact via their Rydberg levels and therefore require additional spatial separation between atoms that are in their Rydberg levels but are not supposed to interact, i.e., to suppress unwanted interactions. Assuming that a single line of atoms in non-Rydberg levels suffices as a buffer between plaquettes that are to be implemented in parallel, this corresponds to two lines of plaquettes as a buffer in each spatial direction. This yields a maximal number of nine layers of constraint gates. Despite the platform-dependent differences, all quantum circuits corresponding to Eq. (A3) have constant circuit depth.

Note that the parity quantum circuits, in contrast to the circuits of the standard gate model as illustrated in Appendix A 1 and Fig. 6, are not modified or optimized by the *pytket* compiler [16] since their gates obey the hardware connectivity by construction.

APPENDIX B: KROTOV’S METHOD FOR QUANTUM OPTIMAL CONTROL

Krotov’s method [93] is an iterative, gradient-based optimization algorithm for time-continuous control fields featuring a build-in monotonic convergence [85]. To achieve the latter, Krotov’s method requires a specific choice of the total optimization functional J , cf. Eq. (6). In detail, while the error measure ε_T at final time T remains the relevant figure of merit that we want to minimized, Krotov’s method achieves its minimization only indirectly by minimizing the total functional J where the time-dependent running costs J_t are give by [82]

$$J_t[\{\psi_l(t)\}, \{\mathcal{E}_k(t)\}, t] = \sum_k \frac{\lambda_k}{S_k(t)} (\mathcal{E}_k(t) - \mathcal{E}_k^{\text{ref}}(t))^2, \quad (\text{B1})$$

where $\mathcal{E}_k^{\text{ref}}(t)$ is a reference field for the control field $\mathcal{E}_k(t)$ that is to be optimized, $S_k(t) \in (0, 1]$ is a shape function and $\lambda_k > 0$ a numerical parameter. With the choice of Eq. (B1), the update equation for field $\mathcal{E}_k(t)$ becomes [85]

$$\mathcal{E}_k^{(i+1)}(t) = \mathcal{E}_k^{\text{ref}}(t) + \frac{S_k(t)}{\lambda_k} \text{Im} \left\{ \sum_l \langle \chi_l^{(i)}(t) | \frac{\partial H[\{\mathcal{E}_k\}]}{\partial \mathcal{E}_k} \Big|_{\{\mathcal{E}_k^{(i+1)}(t)\}} | \psi_l^{(i+1)}(t) \rangle \right\}, \quad (\text{B2})$$

where $|\psi_l^{(i+1)}(t)\rangle$ are forward-propagated states and solutions to the Schrödinger equations

$$\frac{d}{dt} |\psi_l^{(i+1)}(t)\rangle = -i\mathbf{H}^{(i+1)}(t) |\psi_l^{(i+1)}(t)\rangle \quad (\text{B3a})$$

with boundary conditions given by the initial states

$$|\psi_l^{(i+1)}(0)\rangle = |\psi_l(0)\rangle. \quad (\text{B3b})$$

In contrast, $|\chi_l^{(i)}(t)\rangle$ are backward-propagated costates and solutions to the equations

$$\frac{d}{dt} |\chi_l^{(i)}(t)\rangle = i\mathbf{H}^{(i)}(t) |\chi_l^{(i)}(t)\rangle \quad (\text{B4a})$$

with boundary conditions

$$|\chi_l^{(i)}(T)\rangle = -\frac{\partial \varepsilon_T}{\partial \langle \psi_l |} \Big|_{\{\psi_l^{(i)}(T)\}}. \quad (\text{B4b})$$

The superscripts i and $i+1$ in Eqs. (B2)–(B4) indicate whether the corresponding quantity is calculated using the ‘‘old’’ fields from iteration i or the updated fields from iteration $i+1$, respectively.

In order to turn Eq. (B2) into a proper update equation, the reference field $\mathcal{E}_k^{\text{ref}}(t)$ is taken to be the field $\mathcal{E}_k^{(i)}(t)$ from the previous iteration in which case the second term of the left-hand side of Eq. (B2) becomes its update. This choice

causes the time-dependent costs J_t , cf. Eq. (B1), to gradually vanish as the iterative procedure converges. Hence, the error measure ε_T at final time T becomes the dominant term within the total optimization functional J , cf. Eq. (6), and is thus predominantly minimized. Equation (B2) also reveals that while λ_k can be used to control the general size of the update, $S_k(t)$ can be used to suppress updates at certain times.

We refer to Ref. [85] for a more detailed introduction of Krotov’s method and to Ref. [88] for a detailed discussion about its numerical implementation.

APPENDIX C: FIELD GENERATION AND PARAMETRIZATION

In this Appendix, we specify some of the technicalities for the method described in Sec. IV C.

We generate randomized, albeit smooth, guess fields via the formula [103]

$$f(t) = a_0 + \sqrt{2} \sum_{j=1}^m \left[a_j \cos\left(\frac{2\pi jt}{t_1 - t_0}\right) + b_j \sin\left(\frac{2\pi jt}{t_1 - t_0}\right) \right] \quad (C1)$$

where a_0, a_j, b_j are chosen randomly from a normal distribution $N(\mu, \sigma)$ with center $\mu = 0$ and variance $\sigma = 1/(2m+1)$.

The integer m thereby not only determines the number of frequency components in $f(t)$ but also the frequencies of each component. Since we expect frequencies on the same time scale than those contained in the Hamiltonian to be of greater relevance for finding solutions, we chose m randomly from an interval matching each Hamiltonian’s frequencies.

Besides the generation of randomized guess fields, we also specify the internal parametrization of the fields. To this end, note that some of the physical or auxiliary control fields in Eqs. (1) or (2) are experimentally limited in range, i.e., have a lower and upper bound that should not be violated by the optimization algorithm. Let $\mathcal{E}_{\min} \leq \mathcal{E}(t) \leq \mathcal{E}_{\max}$ be bounded. In order to restrict $\mathcal{E}(t)$ to its bounds and to avoid manual truncation in case of violation of the bounds, we internally parametrize $\mathcal{E}(t)$ via

$$u(t) = \operatorname{arctanh}\left(\frac{2\mathcal{E}(t) - \mathcal{E}_{\max} - \mathcal{E}_{\min}}{\mathcal{E}_{\max} - \mathcal{E}_{\min}}\right), \quad (C2a)$$

$$\mathcal{E}(t) = \frac{\mathcal{E}_{\max} - \mathcal{E}_{\min}}{2} \tanh(u(t)) + \frac{\mathcal{E}_{\max} + \mathcal{E}_{\min}}{2}. \quad (C2b)$$

The auxiliary field $-\infty < u(t) < \infty$ is the one optimized in practice and can be optimized boundless since, by construction, it can never violate the boundaries of the actual field $\mathcal{E}(t)$, which it encodes.

-
- [1] P. W. Shor, Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer, *SIAM J. Comput.* **26**, 1484 (1997).
 - [2] A. J. Daley, I. Bloch, C. Kokail, S. Flannigan, N. Pearson, M. Troyer, and P. Zoller, Practical quantum advantage in quantum simulation, *Nature (London)* **607**, 667 (2022).
 - [3] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. D. Sawaya *et al.*, Quantum chemistry in the age of quantum computing, *Chem. Rev.* **119**, 10856 (2019).
 - [4] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, [arXiv:1411.4028](https://arxiv.org/abs/1411.4028).
 - [5] R. Orús, S. Mugel, and E. Lizaso, Quantum computing for finance: Overview and prospects, *Rev. Phys.* **4**, 100028 (2019).
 - [6] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
 - [7] B. M. Terhal, Quantum error correction for quantum memories, *Rev. Mod. Phys.* **87**, 307 (2015).
 - [8] F. Arute *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature (London)* **574**, 505 (2019).
 - [9] Y. Wu *et al.*, Strong quantum computational advantage using a superconducting quantum processor, *Phys. Rev. Lett.* **127**, 180501 (2021).
 - [10] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu *et al.*, Quantum computational advantage using photons, *Science* **370**, 1460 (2020).
 - [11] H.-S. Zhong, Y.-H. Deng, J. Qin, H. Wang, M.-C. Chen, L.-C. Peng, Y.-H. Luo, D. Wu, S.-Q. Gong, H. Su, Y. Hu, P. Hu, X.-Y. Yang, W.-J. Zhang, H. Li, Y. Li, X. Jiang, L. Gan, G. Yang, L. You *et al.*, Phase-programmable Gaussian boson sampling using stimulated squeezed light, *Phys. Rev. Lett.* **127**, 180502 (2021).
 - [12] M. Amy, D. Maslov, and M. Mosca, Polynomial-time t-depth optimization of Clifford + t circuits via matroid partitioning, *IEEE TCAD* **33**, 1476 (2014).
 - [13] L. E. Heyfron and E. T. Campbell, An efficient quantum compiler that reduces t count, *Quantum Sci. Technol.* **4**, 015004 (2018).
 - [14] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, Quantum circuit learning, *Phys. Rev. A* **98**, 032309 (2018).
 - [15] T. Fösel, M. Y. Niu, F. Marquardt, and L. Li, Quantum circuit optimization with deep reinforcement learning, [arXiv:2103.07585](https://arxiv.org/abs/2103.07585).
 - [16] S. Sivarajah, S. Dilkes, A. Cowtan, W. Simmons, A. Edgington, and R. Duncan, t|ket): A retargetable compiler for NISQ devices, *Quantum Sci. Technol.* **6**, 014003 (2021).
 - [17] A. Kissinger and J. van de Wetering, Reducing the number of non-Clifford gates in quantum circuits, *Phys. Rev. A* **102**, 022406 (2020).
 - [18] H. Nagarajan, O. Lockwood, and C. Coffrin, Quantumcircuitopt: An open-source framework for provably optimal quantum circuit design, [arXiv:2111.11674](https://arxiv.org/abs/2111.11674).
 - [19] D. Jaksch, J. I. Cirac, P. Zoller, S. L. Rolston, R. Côté, and M. D. Lukin, Fast quantum gates for neutral atoms, *Phys. Rev. Lett.* **85**, 2208 (2000).
 - [20] H. Levine, A. Keesling, G. Semeghini, A. Omran, T. T. Wang, S. Ebadi, H. Bernien, M. Greiner, V. Vuletić, H. Pichler, and M. D. Lukin, Parallel implementation of high-fidelity multi-qubit gates with neutral atoms, *Phys. Rev. Lett.* **123**, 170503 (2019).

- [21] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, Simple pulses for elimination of leakage in weakly nonlinear qubits, *Phys. Rev. Lett.* **103**, 110501 (2009).
- [22] J. J. García-Ripoll, P. Zoller, and J. I. Cirac, Speed optimized two-qubit gates with laser coherent control techniques for ion trap quantum computing, *Phys. Rev. Lett.* **91**, 157901 (2003).
- [23] V. M. Schäfer, C. J. Ballance, K. Thirumalai, L. J. Stephenson, T. G. Ballance, A. M. Steane, and D. M. Lucas, Fast quantum logic gates with trapped-ion qubits, *Nature (London)* **555**, 75 (2018).
- [24] S. Deffner and S. Campbell, Quantum speed limits: from Heisenberg's uncertainty principle to optimal quantum control, *J. Phys. A: Math. Theor.* **50**, 453001 (2017).
- [25] Because of the restriction to 2D architectures, we do not consider trapped ion platforms in this study.
- [26] M. Fellner, A. Messinger, K. Ender, and W. Lechner, Universal parity quantum computing, *Phys. Rev. Lett.* **129**, 180503 (2022).
- [27] W. Lechner, P. Hauke, and P. Zoller, A quantum annealing architecture with all-to-all connectivity from local interactions, *Sci. Adv.* **1**, e1500838 (2015).
- [28] K. Ender, R. ter Hoeven, B. E. Niehoff, M. Drieb-Schön, and W. Lechner, Parity quantum optimization: Compiler, *Quantum* **7**, 950 (2023).
- [29] W. Lechner, Quantum approximate optimization with parallelizable gates, *IEEE Trans. Quantum Eng.* **1**, 1 (2020).
- [30] K. Ender, A. Messinger, M. Fellner, C. Dlaska, and W. Lechner, Modular parity quantum approximate optimization, *PRX Quantum* **3**, 030304 (2022).
- [31] A. Pagano, S. Weber, D. Jaschke, T. Pfau, F. Meinert, S. Montangero, and H. P. Büchler, Error budgeting for a controlled-phase gate with strontium-88 Rydberg atoms, *Phys. Rev. Res.* **4**, 033019 (2022).
- [32] G. Pelegrí, A. J. Daley, and J. D. Pritchard, High-fidelity multiqubit Rydberg gates via two-photon adiabatic rapid passage, *Quantum Sci. Technol.* **7**, 045020 (2022).
- [33] S. Jandura and G. Pupillo, Time-optimal two- and three-qubit gates for Rydberg atoms, *Quantum* **6**, 712 (2022).
- [34] J. Kelly, R. Barends, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, I.-C. Hoi, E. Jeffrey, A. Megrant, J. Mutus, C. Neill, P. J. J. O'Malley, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White *et al.*, Optimal quantum control using randomized benchmarking, *Phys. Rev. Lett.* **112**, 240504 (2014).
- [35] M. H. Goerz, F. Motzoi, K. B. Whaley, and C. P. Koch, Charting the circuit QED design landscape using optimal control theory, *npj Quantum Inf.* **3**, 37 (2017).
- [36] M. Werninghaus, D. J. Egger, F. Roy, S. Machnes, F. K. Wilhelm, and S. Filipp, Leakage reduction in fast superconducting qubit gates via optimal control, *npj Quantum Inf.* **7**, 14 (2021).
- [37] Google Quantum AI, Quantum computer datasheet, www.quantumai.google/hardware/datasheet/weber.pdf, access: November 2022.
- [38] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Elementary gates for quantum computation, *Phys. Rev. A* **52**, 3457 (1995).
- [39] D. Bluvstein, H. Levine, G. Semeghini, T. T. Wang, S. Ebadi, M. Kalinowski, A. Keesling, N. Maskara, H. Pichler, M. Greiner *et al.*, A quantum processor based on coherent transport of entangled atom arrays, *Nature (London)* **604**, 451 (2022).
- [40] T. M. Graham, Y. Song, J. Scott, C. Poole, L. Phuttitarn, K. Jooya, P. Eichler, X. Jiang, A. Marra, B. Grinkemeyer *et al.*, Multi-qubit entanglement and algorithms on a neutral-atom quantum computer, *Nature (London)* **604**, 457 (2022).
- [41] Note that in practice, gate sets are typically tailored to match the quantum algorithm in question. However, this goes beyond the scope of this paper.
- [42] Assuming that the qubit is encoded in the atomic ground state manifold, such as hyperfine states.
- [43] Note that we do not consider the lifetime of the intermediate levels between qubit and Rydberg states.
- [44] C. Wang, X. Li, H. Xu, Z. Li, J. Wang, Z. Yang, Z. Mi, X. Liang, T. Su, C. Yang *et al.*, Towards practical quantum computers: transmon qubit with a lifetime approaching 0.5 milliseconds, *npj Quantum Inf.* **8**, 3 (2022).
- [45] A. Holmes, S. Johri, G. G. Guerreschi, J. S. Clarke, and A. Y. Matsuura, Impact of qubit connectivity on quantum algorithm performance, *Quantum Sci. Technol.* **5**, 025009 (2020).
- [46] M. P. Harrigan *et al.*, Quantum approximate optimization of non-planar graph problems on a planar superconducting processor, *Nat. Phys.* **17**, 332 (2021).
- [47] K. Barnes, P. Battaglino, B. J. Bloom, K. Cassella, R. Coxe, N. Crisosto, J. P. King, S. S. Kondov, K. Kotru, S. C. Larsen *et al.*, Assembly and coherent control of a register of nuclear spin qubits, *Nat. Commun.* **13**, 2779 (2022).
- [48] IBM Quantum, IBM quantum roadmap, www.ibm.com/quantum/roadmap, access: November 2022.
- [49] M. Fellner, A. Messinger, K. Ender, and W. Lechner, Applications of universal parity quantum computation, *Phys. Rev. A* **106**, 042442 (2022).
- [50] Note that this statement does not include additional PM-QAOA-performance-increasing strategies such as described in Ref. [104].
- [51] M. Saffman, Quantum computing with atomic qubits and Rydberg interactions: progress and challenges, *J. Phys. B: At. Mol. Opt. Phys.* **49**, 202001 (2016).
- [52] L. Henriët, L. Beguin, A. Signoles, T. Lahaye, A. Browaeys, G.-O. Reymond, and C. Jurczak, Quantum computing with neutral atoms, *Quantum* **4**, 327 (2020).
- [53] A. Browaeys and T. Lahaye, Many-body physics with individually controlled Rydberg atoms, *Nat. Phys.* **16**, 132 (2020).
- [54] D. Schrader, I. Dotsenko, M. Khudaverdyan, Y. Miroshnychenko, A. Rauschenbeutel, and D. Meschede, Neutral atom quantum register, *Phys. Rev. Lett.* **93**, 150501 (2004).
- [55] D. D. Yavuz, P. B. Kulatunga, E. Urban, T. A. Johnson, N. Proite, T. Henage, T. G. Walker, and M. Saffman, Fast ground state manipulation of neutral atoms in microscopic optical traps, *Phys. Rev. Lett.* **96**, 063001 (2006).
- [56] T. Xia, M. Lichtman, K. Maller, A. W. Carr, M. J. Piotrowicz, L. Isenhower, and M. Saffman, Randomized benchmarking of single-qubit gates in a 2D array of neutral-atom qubits, *Phys. Rev. Lett.* **114**, 100503 (2015).
- [57] Y. Wang, X. Zhang, T. A. Corcovilos, A. Kumar, and D. S. Weiss, Coherent addressing of individual neutral atoms in a 3D optical lattice, *Phys. Rev. Lett.* **115**, 043003 (2015).

- [58] Y. Wang, A. Kumar, T.-Y. Wu, and D. S. Weiss, Single-qubit gates based on targeted phase shifts in a 3D neutral atom array, *Science* **352**, 1562 (2016).
- [59] M. D. Lukin, M. Fleischhauer, R. Cote, L. M. Duan, D. Jaksch, J. I. Cirac, and P. Zoller, Dipole blockade and quantum information processing in mesoscopic atomic ensembles, *Phys. Rev. Lett.* **87**, 037901 (2001).
- [60] K. M. Maller, M. T. Lichtman, T. Xia, Y. Sun, M. J. Piotrowicz, A. W. Carr, L. Isenhower, and M. Saffman, Rydberg-blockade controlled-not gate and entanglement in a two-dimensional array of neutral-atom qubits, *Phys. Rev. A* **92**, 022336 (2015).
- [61] Y.-Y. Jau, A. M. Hankin, T. Keating, I. H. Deutsch, and G. W. Biedermann, Entangling atomic spins with a Rydberg-dressed spin-flip blockade, *Nat. Phys.* **12**, 71 (2016).
- [62] H. Levine, A. Keesling, A. Omran, H. Bernien, S. Schwartz, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, High-fidelity control and entanglement of Rydberg-atom qubits, *Phys. Rev. Lett.* **121**, 123603 (2018).
- [63] T. M. Graham, M. Kwon, B. Grinkemeyer, Z. Marra, X. Jiang, M. T. Lichtman, Y. Sun, M. Ebert, and M. Saffman, Rydberg-mediated entanglement in a two-dimensional neutral atom qubit array, *Phys. Rev. Lett.* **123**, 230501 (2019).
- [64] In practice, the coupling between qubit states and the Rydberg level is typically achieved via a two-photon process and an intermediate level [96], which we neglect for simplicity.
- [65] C. Dłaska, K. Ender, G. B. Mbeng, A. Kruckenhauser, W. Lechner, and R. van Bijnen, Quantum optimization via four-body Rydberg gates, *Phys. Rev. Lett.* **128**, 120503 (2022).
- [66] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver, Superconducting qubits: Current state of play, *Annu. Rev. Condens. Matter Phys.* **11**, 369 (2020).
- [67] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Charge-insensitive qubit design derived from the Cooper pair box, *Phys. Rev. A* **76**, 042319 (2007).
- [68] S. Gustavsson, O. Zwiernik, J. Bylander, F. Yan, F. Yoshihara, Y. Nakamura, T. P. Orlando, and W. D. Oliver, Improving quantum gate fidelities by using a qubit to measure microwave pulse distortions, *Phys. Rev. Lett.* **110**, 040502 (2013).
- [69] M. A. Rol, C. C. Bultink, T. E. O'Brien, S. R. de Jong, L. S. Theis, X. Fu, F. Luthi, R. F. L. Vermeulen, J. C. de Sterke, A. Bruno, D. Deurloo, R. N. Schouten, F. K. Wilhelm, and L. DiCarlo, Restless tuneup of high-fidelity qubit gates, *Phys. Rev. Appl.* **7**, 041001(R) (2017).
- [70] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell *et al.*, Superconducting quantum circuits at the surface code threshold for fault tolerance, *Nature (London)* **508**, 500 (2014).
- [71] Y. Sung, L. Ding, J. Braumüller, A. Vepsäläinen, B. Kannan, M. Kjaergaard, A. Greene, G. O. Samach, C. McNally, D. Kim, A. Melville, B. M. Niedzielski, M. E. Schwartz, J. L. Yoder, T. P. Orlando, S. Gustavsson, and W. D. Oliver, Realization of high-fidelity CZ and zz -free i SWAP gates with a tunable coupler, *Phys. Rev. X* **11**, 021058 (2021).
- [72] J. Majer, J. M. Chow, J. M. Gambetta, J. Koch, B. R. Johnson, J. A. Schreier, L. Frunzio, D. I. Schuster, A. A. Houck, A. Wallraff *et al.*, Coupling superconducting qubits via a cavity bus, *Nature (London)* **449**, 443 (2007).
- [73] Y. Chen, C. Neill, P. Roushan, N. Leung, M. Fang, R. Barends, J. Kelly, B. Campbell, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, A. Megrant, J. Y. Mutus, P. J. J. O'Malley, C. M. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White *et al.*, Qubit architecture with high coherence and fast tunable coupling, *Phys. Rev. Lett.* **113**, 220502 (2014).
- [74] U. Boscain, M. Sigalotti, and D. Sugny, Introduction to the Pontryagin maximum principle for quantum optimal control, *PRX Quantum* **2**, 030203 (2021).
- [75] A. Carlini, A. Hosoya, T. Koike, and Y. Okudaira, Time-optimal quantum evolution, *Phys. Rev. Lett.* **96**, 060503 (2006).
- [76] J. K. Nauth and V. M. Stojanović, Quantum-brachistochrone approach to the conversion from w to Greenberger-Horne-Zeilinger states for Rydberg-atom qubits, *Phys. Rev. A* **106**, 032605 (2022).
- [77] X. Wang, M. Allegra, K. Jacobs, S. Lloyd, C. Lupo, and M. Mohseni, Quantum brachistochrone curves as geodesics: Obtaining accurate minimum-time protocols for the control of quantum systems, *Phys. Rev. Lett.* **114**, 170501 (2015).
- [78] A. M. Frydryszak and V. M. Tkachuk, Quantum brachistochrone problem for a spin-1 system in a magnetic field, *Phys. Rev. A* **77**, 014103 (2008).
- [79] K. Svoboda, L. B. Levitin, T. Toffoli, and Z. Walton, Maximum speed of quantum gate operation, *Int. J. Theor. Phys.* **44**, 965 (2005).
- [80] T. Caneva, M. Murphy, T. Calarco, R. Fazio, S. Montangero, V. Giovannetti, and G. E. Santoro, Optimal control at the quantum speed limit, *Phys. Rev. Lett.* **103**, 240501 (2009).
- [81] C. P. Koch, U. Boscain, T. Calarco, G. Dirr, S. Filipp, S. J. Glaser, R. Kosloff, S. Montangero, T. Schulte-Herbrüggen, D. Sugny, and F. K. Wilhelm, Quantum optimal control in quantum technologies. Strategic report on current status, visions and goals for research in Europe, *EPJ Quantum Technol.* **9**, 19 (2022).
- [82] J. P. Palao and R. Kosloff, Optimal control theory for unitary transformations, *Phys. Rev. A* **68**, 062308 (2003).
- [83] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, Optimal control of coupled spin dynamics: Design of NMR pulse sequences by gradient ascent algorithms, *J. Magn. Reson.* **172**, 296 (2005).
- [84] T. Caneva, T. Calarco, and S. Montangero, Chopped random-basis quantum optimization, *Phys. Rev. A* **84**, 022326 (2011).
- [85] D. M. Reich, M. Ndong, and C. P. Koch, Monotonically convergent optimization in quantum control using Krotov's method, *J. Chem. Phys.* **136**, 104103 (2012).
- [86] J. J. W. H. Sørensen, M. O. Aramburu, T. Heinzel, and J. F. Sherson, Quantum optimal control in a chopped basis: Applications in control of Bose-Einstein condensates, *Phys. Rev. A* **98**, 022119 (2018).
- [87] S. Machnes, E. Assémat, D. Tannor, and F. K. Wilhelm, Tunable, flexible, and efficient optimization of control pulses for practical qubits, *Phys. Rev. Lett.* **120**, 150401 (2018).
- [88] M. H. Goerz, D. Basilewitsch, F. Gago-Encinas, M. G. Krauss, K. P. Horn, D. M. Reich, and C. P. Koch, Krotov: A Python implementation of Krotov's method for quantum optimal control, *SciPost Phys.* **7**, 080 (2019).

- [89] N. Wittler, F. Roy, K. Pack, M. Werninghaus, A. S. Roy, D. J. Egger, S. Filipp, F. K. Wilhelm, and S. Machnes, Integrated tool set for control, calibration, and characterization of quantum devices applied to superconducting qubits, *Phys. Rev. Appl.* **15**, 034080 (2021).
- [90] M. H. Goerz, S. C. Carrasco, and V. S. Malinovsky, Quantum optimal control via semi-automatic differentiation, *Quantum* **6**, 871 (2022).
- [91] J. D. Teske, P. Cerfontaine, and H. Bluhm, qopt: An experiment-oriented software package for qubit simulation and quantum optimal control, *Phys. Rev. Appl.* **17**, 034036 (2022).
- [92] M. Rossignolo, T. Reisser, A. Marshall, P. Rembold, A. Pagano, P. J. Vetter, R. S. Said, M. M. Müller, F. Motzoi, T. Calarco, F. Jelezko, and S. Montangero, QuOCS: The quantum optimal control suite, *Comput. Phys. Commun.* **291**, 108782 (2023).
- [93] A. I. Konnov and V. F. Krotov, On global methods of successive improvement of controlled processes, *Autom. Rem. Contr.* **60**, 1427 (1999).
- [94] I. Brouzos, A. I. Streltsov, A. Negretti, R. S. Said, T. Caneva, S. Montangero, and T. Calarco, Quantum speed limit and optimal control of many-boson dynamics, *Phys. Rev. A* **92**, 062110 (2015).
- [95] D. Basilewitsch, J. Fischer, D. M. Reich, D. Sugny, and C. P. Koch, Fundamental bounds on qubit reset, *Phys. Rev. Res.* **3**, 013110 (2021).
- [96] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner *et al.*, Probing many-body dynamics on a 51-atom quantum simulator, *Nature (London)* **551**, 579 (2017).
- [97] T. Linowski, G. Rajchel-Mieldzióć, and K. Życzkowski, Entangling power of multipartite unitary gates, *J. Phys. A: Math. Theor.* **53**, 125303 (2020).
- [98] D. Barredo, V. Lienhard, S. de Léséleuc, T. Lahaye, and A. Browaeys, Synthetic three-dimensional atomic structures assembled atom by atom, *Nature (London)* **561**, 79 (2018).
- [99] Y. Xu, J. Chu, J. Yuan, J. Qiu, Y. Zhou, L. Zhang, X. Tan, Y. Yu, S. Liu, J. Li, F. Yan, and D. Yu, High-fidelity, high-scalability two-qubit gate scheme for superconducting qubits, *Phys. Rev. Lett.* **125**, 240503 (2020).
- [100] X. Li, T. Cai, H. Yan, Z. Wang, X. Pan, Y. Ma, W. Cai, J. Han, Z. Hua, X. Han, Y. Wu, H. Zhang, H. Wang, Y. Song, L. Duan, and L. Sun, Tunable coupler for realizing a controlled-phase gate with dynamically decoupled regime in a superconducting circuit, *Phys. Rev. Appl.* **14**, 024070 (2020).
- [101] Note that every eigenstate of Eq. (A1) appears at least twice due to a global spin flip symmetry.
- [102] J. Unger, A. Messinger, B. E. Niehoff, M. Fellner, and W. Lechner, Low-depth circuit implementation of parity constraints for quantum optimization, [arXiv:2211.11287](https://arxiv.org/abs/2211.11287).
- [103] S. Filip, A. Javeed, and L. N. Trefethen, Smooth random functions, random ODEs, and Gaussian processes, *SIAM Rev.* **61**, 185 (2019).
- [104] A. Weidinger, G. B. Mbeng, and W. Lechner, Error mitigation for quantum approximate optimization, [arXiv:2301.05042](https://arxiv.org/abs/2301.05042).