

Guided quantum walk

Sebastian Schulz^{1,*}, Dennis Willsch¹, and Kristel Michielsen^{1,2,3}

¹Jülich Supercomputing Centre, Institute for Advanced Simulation, Forschungszentrum Jülich, 52425 Jülich, Germany

²AIDAS, 52425 Jülich, Germany

³RWTH Aachen University, 52056 Aachen, Germany



(Received 24 August 2023; accepted 4 January 2024; published 21 March 2024)

We utilize the theory of local amplitude transfer (LAT) to gain insights into quantum walks (QWs) and quantum annealing (QA) beyond the adiabatic theorem. By representing the eigenspace of the problem Hamiltonian as a hypercube graph, we demonstrate that probability amplitude traverses the search space through a series of local Rabi oscillations. We argue that the amplitude movement can be systematically guided towards the ground state using a time-dependent hopping rate based solely on the problem's energy spectrum. Building upon these insights, we extend the concept of multistage QW by introducing the guided quantum walk (GQW) as a bridge between QW-like and QA-like procedures. We assess the performance of the GQW on exact cover, traveling salesperson, and garden optimization problems with 9 to 30 qubits. Our results provide evidence for the existence of optimal annealing schedules, beyond the requirement of adiabatic time evolutions. These schedules might be capable of solving large-scale combinatorial optimization problems within evolution times that scale linearly in the problem size.

DOI: [10.1103/PhysRevResearch.6.013312](https://doi.org/10.1103/PhysRevResearch.6.013312)

I. INTRODUCTION

Combinatorial optimization is a fundamental problem in computer science that has a wide range of important applications in finance [1], scheduling [2], machine learning [3,4], database search [5], computational biology [6], and operations research [7]. However, finding optimal solutions to such problems can be challenging and computationally expensive, which often makes them intractable for classical computers today. Recent advances in quantum hardware are raising the expectations for demonstrating useful quantum computation in the upcoming years. In particular, solving large-scale combinatorial optimization problems is considered one of the great application areas of quantum computation, driving the need for quantum optimization algorithms suitable for near-term quantum devices [8].

Quantum walks (QW) [5,9–12] and quantum annealing (QA) [13–22] have emerged as two promising candidates for continuous-time quantum optimization algorithms in this context. QWs, introduced by Aharonov *et al.* [9], model the search space as a graph and govern the walker's transitions between vertices using a time-independent Hamiltonian. On the other hand, QA employs a time-dependent Hamiltonian to adiabatically evolve from an initial Hamiltonian to a problem Hamiltonian. Both algorithms have been extensively

studied for various problems with tens of qubits, including Sherrington-Kirkpatrick spin glasses [11,23], Max-Cut [12], 2-SAT [24–28], and exact cover [2] problem instances. While these algorithms exhibit distinct dynamics, Morley *et al.* [29] have argued that they can be seen as extreme cases of annealing schedules in the context of search problems. However, the dynamics occurring in the intermediate region between QA and QW for combinatorial optimization problems have yet to be fully explored. We argue that these intermediate evolutions, which go beyond the scope of the adiabatic theorem, are a very promising regime for effective quantum computation.

In this paper, we utilize the theory of local amplitude transfer (LAT) to investigate continuous-time quantum algorithms. LAT theory focuses on the local energy structure of the problem Hamiltonian on the hypercube graph. It provides insights into the movement of probability amplitude between individual elements of the search space through a series of local Rabi oscillations.

Building upon these insights, we extend the concept of multistage QW [11,12] by introducing the guided quantum walk (GQW). The GQW combines multiple QWs through a time-dependent hopping rate, effectively guiding the transfer of probability amplitude throughout the graph. This approach relates to the problem of finding optimal annealing schedules [30–32], but goes beyond the requirement of adiabatic time evolutions, placing the GQW in between QW-like and QA-like procedures.

To evaluate the performance of the GQW, we numerically simulate its application to exact cover (EC) [2,33,34], traveling salesperson (TSP) [7,35–38], and garden optimization (GO) [39] problems using the JUWELS Booster supercomputer [40]. We extensively study the GQW on problem instances with up to 30 qubits. Our results provide evidence

*Corresponding author: se.schulz@fz-juelich.de

for the existence of optimal annealing schedules capable of solving large-scale combinatorial optimization problems with evolution times that scale linearly in the problem size.

The paper is organized as follows: In Sec. II, we introduce the types of combinatorial optimization problems that are used in our research. Section III provides a review of QWs and explores their dynamics on search and optimization problems using the LAT theory. We discuss how probability amplitude can be effectively guided through the hypercube graph, leading to the development of the GQW. Furthermore, we examine the relationship between the GQW, QA, and QWs at different evolution times. In Sec. IV, we present a comprehensive performance analysis of the GQW. Finally, in Sec. V, we summarize our findings and their implications for future research.

II. COMBINATORIAL OPTIMIZATION PROBLEMS

This section presents the combinatorial optimization problems studied in our research. First, we describe the problem types and their encoding into a quantum setting in Sec. II A. Subsequently, we introduce a benchmarking metric, the solution quality, which ensures fair comparisons of quantum optimization algorithms across different problem types and sizes in Sec. II B.

A. Definition

A broad class of real-world problems can be framed as combinatorial optimization problems [7]. These are problems defined on N -bit binary strings $z = z_{N-1} \dots z_0 = \{0, 1\}^{\otimes N}$, with the objective of finding a string z_{opt} that minimizes a given classical cost function $C(z) : \{0, 1\}^{\otimes N} \rightarrow \mathbb{R}_{\geq 0}$. A natural way to express C in a quantum setting is to encode it into the energy spectrum of the computational basis states $|z\rangle$ of a quantum cost Hamiltonian

$$\hat{H}_C = \sum_{z \in \{0,1\}^{\otimes N}} C(z) |z\rangle \langle z|. \quad (1)$$

We consider the case in which $C(z)$ consists only of linear and quadratic terms, $C(z) = \sum_{ij} z_i Q_{ij} z_j$. Such problems are known as quadratic unconstrained binary optimization (QUBO) problems, with $\{Q_{ij}\}$ denoting the QUBO coefficients. For these problems, $\hat{H}_C$ can be written in the form of an Ising Hamiltonian by substituting $z_i \mapsto (1 - \hat{\sigma}_i^z)/2$, where $\hat{\sigma}_i^z$ denotes the Pauli-Z operator applied to the i th qubit with the identity operator acting on the remaining qubits. One obtains

$$\hat{H}_C = - \sum_{i=0}^{N-1} h_i \hat{\sigma}_i^z + \sum_{i \neq j} J_{ij} \hat{\sigma}_i^z \hat{\sigma}_j^z, \quad (2)$$

with the coefficients $h_i = \sum_j (Q_{ij} + Q_{ji})/2$ and $J_{ij} = Q_{ij}/4$ representing the optimization problem. In our study, we investigate three types of combinatorial optimization problems, which can be expressed in Ising form and represent two categories of cost functions.

The first category (constraint-only) concerns cost functions consisting entirely of constraining terms, meaning that all states describing valid solutions to the optimization problem (called valid states henceforth) are assigned to the same cost

value. Here, we focus on EC problem instances with $N \in \{12, 15, 18, 21, 24, 30\}$. Generically, these problem instances exhibit numerous distinct energy levels with low degeneracies and a nondegenerate ground state.

The second category (constraint+optimization) includes cost functions that involve both constraining and optimization terms, such that the set of valid states spans multiple energy levels. We consider TSP and GO problem instances with $N \in \{9, 16, 25\}$ and $N \in \{12, 15, 18, 21, 24, 30\}$, respectively. It is worth noting that the GO instances generally exhibit higher degeneracies among their energy levels compared to the TSP instances.

The explicit cost functions for the three types of combinatorial optimization problems can be found in Appendix A. For each problem type and size, we investigate 10 randomly generated problem instances. Note that all cost functions have been rescaled and shifted, such that $C(z_{\text{opt}}) = 0$ and $\max_z C(z) = 100$.

B. Solution quality

Quantum optimization algorithms are designed to find the optimal solution to a given combinatorial optimization problem [Eq. (1)] with high success probability $P_{z_{\text{opt}}} = |\langle z_{\text{opt}} | \Psi \rangle|^2$, where $|\Psi\rangle = \sum_z \psi_z |z\rangle$ denotes the final quantum state. However, in many cases, approximate solutions, where not solely the solution state but also other valid states with a slightly larger cost function value are obtained with high probability, are also of interest, especially if they can be found significantly faster. Since $P_{z_{\text{opt}}}$ does not take these approximate solutions into account, a common approach is to evaluate the performance of quantum optimization algorithms based on the energy expectation value

$$E_\Psi = \langle \Psi | \hat{H}_C | \Psi \rangle, \quad (3)$$

often in the form of the approximation ratio $r(\Psi) = E_\Psi / E_{\text{max}}$, where $E_{\text{max}} = \max_z E_z$. Note that smaller approximation ratios are supposed to represent better solutions.

We emphasize, however, that the approximation ratio r in Eq. (3) is often not able to capture the “quality” of the produced quantum state $|\Psi\rangle$, that is the distribution of the measurement probabilities with respect to the energies of the valid states. This is because r considers not only the set of valid states (which we are generally interested in) but also the set of invalid states (i.e., states that violate at least one constraint and may shift E_{max} arbitrarily). For the problems under investigation, the latter correspond to the majority of states, covering $\geq 95\%$ of the total energy range. Hence, comparing the approximation ratio between different quantum optimization algorithms does not necessarily compare their ability to produce “good” approximate solutions, as a final state with a large value of r might still provide valid states near the optimal solution with higher probability than a different state of smaller r value.

In order to address this issue of the approximation ratio, we propose the use of a theoretical benchmarking metric termed the *solution quality*

$$S_q = \sum_{z \in \text{valid states}} [1 - r_{\text{valid}}(z)] \cdot P_z, \quad (4)$$

where $P_z = |\langle \Psi | z \rangle|^2$ represents the measurement probability of the state $|z\rangle$, and $r_{\text{valid}}(z) = E_z / \max_{k \in \text{valid states}}(E_k)$ is used to give a weight among all valid states in the energy spectrum (note that $|z_{\text{opt}}\rangle$ is at $E_{z_{\text{opt}}} = 0$ by definition). S_q is designed to capture both the characteristics of the success probability and the approximation ratio, restricted to the set of valid states. This provides a theoretical benchmarking metric that is comparable across different problem instances by focusing on the practicality of the solutions obtained. Note that $S_q \in [0, 1]$, with $S_q = 1$ corresponding to the solution state $|\Psi\rangle = |z_{\text{opt}}\rangle$, and $S_q = 0$ indicates a final state $|\Psi\rangle$ that consists solely of invalid states and highest-energy valid states with $r_{\text{valid}} = 1$. The latter are effectively excluded from increasing the solution quality S_q , because obtaining any valid solution to the constraint+optimization problems (TSP and GO) can be done in polynomial time. Hence, we do not consider an information gain from these states. Furthermore, in the case of constraint-only problem instances (EC), we set $r_{\text{valid}}(z) = 0$, such that $S_q = P_{z_{\text{opt}}}$.

III. CONTINUOUS-TIME QUANTUM WALK

In this section, we provide a concise review of the continuous-time QW (Sec. III A) and investigate its dynamics on search (Sec. III B) and combinatorial optimization problems (Sec. III C) using LAT theory. Our analysis shows that the QW can be systematically guided on the hypercube graph based on the energy spectrum of the problem Hamiltonian. Building upon these insights, we introduce the GQW in Sec. III D. In Sec. III E we examine the relationship between the GQW, QA, and QWs.

A. Definition

The continuous-time QW is a quantum algorithm that assigns the computational basis states $|j\rangle$ of an N -qubit Hilbert space to the set of vertex labels $V = \{j\}_{j=0}^{2^N-1}$ of an undirected graph $G(V, T)$. In this framework, the vertices encode the walker's position, and the set of edges T indicate allowed transitions between label pairs (j, k) . The latter is described through an adjacency matrix A , whose elements satisfy $A_{j,k} = 1$ if an edge in G connects vertices j and k , and $A_{j,k} = 0$ otherwise. As G is undirected, A is symmetric and can be used to define the quantum-walk Hamiltonian, given by

$$\hat{H}_{\text{QW}} = \Gamma \cdot \hat{H}_D = -\Gamma \cdot \sum_{j,k} A_{j,k} |j\rangle \langle k|, \quad (5)$$

where $\hat{H}_D$ denotes the driver Hamiltonian, and Γ is the hopping rate. It is important to note that Eq. (5) is not the only possible Hamiltonian for a QW. In the literature, the Laplacian $\hat{L} = -\hat{H}_D - \hat{D}$ of G is often used instead of $\hat{H}_D$, with the diagonal matrix $\hat{D} = \sum_j \deg(j) |j\rangle \langle j|$ encoding the degree $\deg(j)$ of each vertex j , i.e., the number of edges incident to j [5]. However, for the regular graphs G considered in this paper, where $\deg(j)$ is constant with respect to j , both formulations are equivalent up to an unobservable global phase factor. We chose to use the adjacency operator form of $\hat{H}_{\text{QW}}$ for consistency with other quantum optimization algorithms.

Given the QW Hamiltonian [Eq. (5)], the state of the walker evolves from some initial state $|\Psi_0\rangle$ according to the

time-dependent Schrödinger equation, which yields the state of the system at a time T as

$$|\Psi(T)\rangle = e^{-i\hat{H}_{\text{QW}}T} |\Psi_0\rangle, \quad (6)$$

where we have used units with $\hbar = 1$. The quantum dynamics implemented by this evolution clearly depend on the connectivity in the graph G . In the past, QWs have been studied on a variety of graph layouts, including the complete graph, which couples every vertex to every other, and the N -dimensional hypercube, which connects only vertices of Hamming distance one [5]. In this paper, we focus on a hypercube, as it provides a natural encoding for a QW into qubits. Specifically, transitioning from one vertex to a neighboring vertex corresponds to flipping the computational state of one qubit. As such, the driving Hamiltonian $\hat{H}_D$ is composed of N single-body terms,

$$\hat{H}_D = -\sum_{j=1}^N \hat{\sigma}_j^x, \quad (7)$$

where $\hat{\sigma}_j^x$ denotes the Pauli- X operator applied to the j th qubit with the identity operator acting on the remaining qubits. The corresponding QW Hamiltonian for the hypercube is given by $\hat{H}_{\text{QW}} = -\Gamma \cdot \sum_{j=1}^N \hat{\sigma}_j^x$. A primary feature of this Hamiltonian is its ability to rapidly explore the vertices in the graph G , thereby providing high dynamics in the computational basis. By introducing a secondary problem Hamiltonian that is diagonal in the computational basis, the graph G becomes directed, leading to a concentration of amplitude in the ground state of the problem Hamiltonian (see Appendix B for further information). In the following sections, we investigate this ability of QWs to find ground states, starting with the search problem—a well-studied toy problem that can be analytically solved using the QW.

B. Quantum walk search

In the search problem, we aim to find a specific bit string $z_{\text{opt}} \in \{0, 1\}^N$ from a set of 2^N possible strings. The problem can be mapped to a quantum setting using an oracle Hamiltonian

$$\hat{H}_O = -|z_{\text{opt}}\rangle \langle z_{\text{opt}}| \quad (8)$$

that assigns one unit of energy less to the solution state $|z_{\text{opt}}\rangle$ compared to the rest of basis states. Solving the search problem is then equivalent to finding the ground state of $\hat{H}_O$. The QW provides a means of solving the search problem by combining $\hat{H}_O$ with the driving Hamiltonian $\hat{H}_D$ [Eq. (7)], and adjusting their relative strength via the dimensionless hopping rate Γ . The computation is performed by evolving the quantum system, initialized in the equal superposition state

$$|\Psi_0\rangle = |+\rangle^{\otimes N} = \frac{1}{\sqrt{2^N}} \sum_{j=0}^{2^N-1} |j\rangle, \quad (9)$$

under the QW search Hamiltonian

$$\hat{H}_{\text{QWS}} = \Gamma \cdot \hat{H}_D + \hat{H}_O \quad (10)$$

for a time T and measuring the qubit register in the computational basis afterward.




FIG. 1. Application of the QW to an $N = 12$ qubit search problem using the Hamiltonian $\hat{H}_{QWS}$ in Eq. (10). (a) Energy spectrum of $\hat{H}_{QWS}$ as a function of the hopping rate Γ . The blue and red curve denote the ground and first excited state, respectively. The data is shifted by the ground state energy. (b) Overlaps of the ground state (blue) and the first excited state (red) with the solution state (solid) and the initial state (dashed) as a function of Γ . (c) Solution quality S_q as a function of the evolution time T for $\Gamma = \Gamma_{opt} = 1/2^N \sum_{r=1}^N \binom{N}{r}/2r$ [5] (solid blue), $\Gamma = 95\% \cdot \Gamma_{opt}$ (dashed red) and $\Gamma = 90\% \cdot \Gamma_{opt}$ (dash-dotted green).

Childs and Goldstone have solved the QW search problem analytically for various graph layouts [5], including the complete and hypercube graphs. For each layout, they have calculated optimal values $\Gamma = \Gamma_{opt}$ (see also [10]) for which the performance of the QW search achieves the same optimal quadratic speed up as Grover's search algorithm [41].

Figure 1 presents the application of the QW to an $N = 12$ qubit search problem, giving insights into the characteristics of the QW's dynamics. Specifically, the entire system evolves periodically, with the individual measurement probabilities of the vertices oscillating as a function of the evolution time T [see Fig. 1(c)]. This behavior occurs because the QW performs Rabi oscillations between the initial state $|\Psi_0\rangle$ and the

solution state $|z_{opt}\rangle$. Figure 1(a) shows the two lowest energy levels corresponding to the states $|E_0(\Gamma)\rangle$ and $|E_1(\Gamma)\rangle$ as a function of the hopping rate Γ . When $\Gamma = \Gamma_{opt}$, the relative strengths of the two contributing Hamiltonians in $\hat{H}_{QWS}$ are balanced equally, and the two energy levels undergo an avoided level crossing. If N is large enough, $|E_{0,1}(\Gamma_{opt})\rangle$ is approximately equal to the uniform superposition of the initial and the solution state, i.e., $|E_{0,1}(\Gamma_{opt})\rangle \approx (|\Psi_0\rangle \pm |z_{opt}\rangle)/\sqrt{2}$ [see Fig. 1(b)]. Hence, $\hat{H}_{QWS}$ drives transitions between $|\Psi_0\rangle$ and $|z_{opt}\rangle$ with a frequency $\propto [E_1(\Gamma_{opt}) - E_0(\Gamma_{opt})]$. Consequently, the overlap with the solution state $|\langle z_{opt} | \Psi(T) \rangle|$ depends on the hopping rate Γ and the evolution time T , which both require high precision in order to obtain accurate results [see. Fig. 1(c)].

C. Guiding a quantum walk

Combinatorial optimization problems can also benefit from the application of QWs. We can address optimization problems within the QW by augmenting the driver Hamiltonian $\hat{H}_D$ with the cost Hamiltonian $\hat{H}_C$ defined in Eq. (1). The latter induces complex phase gradients between connected vertices based on their assigned cost value, thereby defining a direction of propagation for the walker in the graph. The QW optimization Hamiltonian on the hypercube mapping is given by

$$\hat{H}_{QWO} = \Gamma \cdot \hat{H}_D + \hat{H}_C, \quad (11)$$

with Γ balancing the relative strength of the two contributing parts. The QW is performed analogously to Sec. III B, by initializing the qubits in the equal superposition state $|\Psi_0\rangle$ [see Eq. (9)], evolving the system under $\hat{H}_{QWO}$ for a time T , and then measuring it in the computational basis. However, unlike for the search problem, the evolution under $\hat{H}_{QWO}$ cannot be efficiently calculated analytically. This makes it impractical to predict optimal parameter sets Γ_{opt} and T_{opt} , which maximize the final overlap with the solution state for an arbitrary optimization problem. This is because the energy spectrum of $\hat{H}_C$ typically features numerous distinct levels with unknown energy gaps, in contrast to the almost completely degenerate spectrum of $\hat{H}_O$. Consequently, the energy levels of the combined Hamiltonian split as a function of Γ and thereby undergo numerous avoided level crossings [see Fig. 2(a)]. Since the system is initialized in a superposition across multiple energy levels, $\hat{H}_{QWO}$ drives transitions between various eigenstates, and the simple two-level description of the walker's dynamics used previously is no longer applicable [see Fig. 2(b)]. As a result, the oscillation of the solution quality S_q becomes highly complex, as multiple streams of amplitude transfers at different energy levels interfere with each other [see Fig. 2(c)].

Previous studies have explored heuristic approaches to obtain near-optimal hopping rates Γ for the QW within polynomial time. For instance, Callison *et al.* proposed estimating Γ_{opt} from the overall energy scale of $\hat{H}_C$ by matching the total energy spreads of the two Hamiltonians in $\hat{H}_{QWO}$ [23]. Later, the authors extended this strategy by sampling Γ_{opt} based on a maximization of the average dynamics on Sherrington-Kirkpatrick spin glass problems [11]. Recently, Banks *et al.* investigated the link between time-independent Hamiltonians




FIG. 2. Application of the QW to solve an EC problem for $N = 12$ qubits using the Hamiltonian $\hat{H}_{QWO}$ in Eq. (11). (a) Energy spectrum of $\hat{H}_{QWO}$ as a function of the hopping rate Γ . The blue and red curve denote the ground and first excited state, respectively. The data is shifted by the ground state energy. (b) Overlaps of the ground state (blue) and the first excited state (red) with the solution state (solid) and the initial state (dashed) as a function of Γ . (c) Solution quality S_q as a function of the evolution time T for $\Gamma = \Gamma_{opt}$ (solid blue), $\Gamma = 1.2 \cdot \Gamma_{opt}$ (dashed red) and $\Gamma = 1.5 \cdot \Gamma_{opt}$ (dash-dotted green). The value of Γ_{opt} has been computed numerically.

and thermalization, leading to an estimate of Γ_{opt} through the eigenstate thermalization hypothesis on Max-Cut problem instances [12].

While these strategies demonstrate the general ability of QWs to solve combinatorial optimization problems, they necessitate additional adjustments for each problem type (e.g., estimating energy gaps) and focus solely on the average dynamics in the hypercube graph. However, local variations in these dynamics across different regions of the graph are the primary reason for the observed distortions in the oscillation of the solution quality in Fig. 2(c). These distortions not only make it challenging to estimate T_{opt} from a few samples, but also limit the maximally achievable solution quality for any T

as the system approaches a stationary state (cf. Fig. 4 below). The optimal solution quality typically scales exponentially with the problem size N [cf. Fig. 8(b) below] because the QW can only drive amplitude transfers within a fixed energy range, neglecting the amplitude originating from exponentially many states outside this range. Consequently, the QW as defined in Eq. (11) is not well suited for large combinatorial optimization problems.

Inspired by these strategies, but being interested in achieving practical quantum computation (as measured by a large solution quality) for any given evolution time T , problem size N and problem type, we investigate the dynamics of QWs by applying the LAT theory to the hypercube graph G shown in Fig. 3(a). By analyzing the transfer of probability amplitude in local subspaces spanned by basis-state pairs with Hamming distance 1 (i.e., states connected by an edge in G), we aim to derive a mechanism to control the movement of the walker locally in the graph, such that backpropagation of amplitude into (undesired) high-energy states can be suppressed. Thus, instead of maximizing transitions in the entire graph collectively, as proposed in prior study, our approach is to maximize them only locally at a time in order to guide the walker towards the solution state more effectively.

The LAT theory focuses on two-dimensional subspaces spanned by pairs of basis states, $|j\rangle$ and $|k\rangle$, that are connected in the hypercube graph G , i.e., $\langle j|\hat{H}_D|k\rangle \neq 0$ [see inset of Fig. 3(a)]. The effective two-level subspace Hamiltonian is given by

$$\hat{H}_{QWO}^{(j,k)} = \Gamma \hat{H}_D^{(j,k)} + \hat{H}_C^{(j,k)}, \quad (12)$$

assuming the rest of the system remains in its initial state. Here, $\hat{H}_C^{(j,k)}$ denotes the local cost Hamiltonian and $\hat{H}_D^{(j,k)}$ is the local driving Hamiltonian, which are defined as

$$\hat{H}_C^{(j,k)} = \begin{pmatrix} \langle j|\hat{H}_C|j\rangle & 0 \\ 0 & \langle k|\hat{H}_C|k\rangle \end{pmatrix} \quad \hat{H}_D^{(j,k)} = -\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}. \quad (13)$$

If the system starts in the local equal superposition state, $|+\rangle^{(j,k)} = (|j\rangle + |k\rangle)/\sqrt{2}$, with the local energy gaps $\Gamma \Delta^D = -2\Gamma$ of $\Gamma \hat{H}_D^{(j,k)}$ and $E_j - E_k = \Delta_{j,k}^C = 2\delta$ of $\hat{H}_C^{(j,k)}$, the measurement probability of the desired lower-energy state $|k\rangle$, is given by

$$P_k(t, \Gamma) = \frac{1}{2} + \frac{\Gamma \cdot \delta}{\Gamma^2 + \delta^2} \sin(\sqrt{\Gamma^2 + \delta^2} \cdot t). \quad (14)$$

Equation (14) represents a sinusoidal oscillation with a Rabi frequency of $\omega = \sqrt{\Gamma^2 + \delta^2}$ similar to Sec. III B (note that this property holds for any initial state). Specifically, if $\Gamma \gg 1$, the driving Hamiltonian dominates $\hat{H}_{QWO}^{(j,k)}$, and the cost Hamiltonian has almost no influence on the system's evolution. Since $|+\rangle^{(j,k)}$ corresponds to the ground state of $\hat{H}_D^{(j,k)}$, the system remains primarily in its initial state. Conversely, if $\Gamma \ll 1$, $\hat{H}_C^{(j,k)}$ governs the evolution. Since $\hat{H}_C^{(j,k)}$ is diagonal and only induces phase rotations in the computational basis, no amplitude transfer occurs. Only for $|\Gamma \Delta^D| \approx \Delta_{j,k}^C$, the two Hamiltonians' relative strengths are balanced, and transitions in the local subspace are maximized.

As a result, amplitude transfers can only occur efficiently among specific subsets of vertex pairs with an energy gap




FIG. 3. Energy spectrum analysis of an $N = 21$ qubit EC problem. (a) Hypercube graph representation of $\hat{H}_C$, where each point corresponds to a computational basis state $|j\rangle$, ordered by increasing energy $E_j = \langle j | \hat{H}_C | j \rangle$ from left to right. The color of each point represents the total Hamming distance to the ground state z_{opt} . Two states $|j\rangle$ and $|k\rangle$ are connected by an edge if the bitstrings j and k have a Hamming distance of 1, meaning that the driver Hamiltonian $\hat{H}_D$ drives transitions between them. The inset in panel (a) provides a zoomed-in view near z_{opt} , highlighting the energy gaps $\Delta^C_{j,k}$ between connected states $|j\rangle$ and $|k\rangle$. (b) Distribution of the largest energy gaps $\Delta^C_{j,k}$ from a vertex $|j\rangle$ to a lower energy vertex $|k\rangle$ as a function of E_j . The data reveals an increasing trend of $\Delta^C_{j,k}$ with respect to E_j . The solid black curve represents a fit to the average energy gap $\langle \Delta^C \rangle(E)$ at energy level E (using the scale of the top axis), used by the GQW to locally adjust the relative strength of the driving and problem Hamiltonian in Eq. (11). The dashed black curve illustrates the optimal relative strength for a QW. The bottom axis denotes the relative time in both algorithms, with time progressing from right to left, as the GQW progresses from high-energy to low-energy amplitude transfers.

$\Delta^C_{j,k} \approx |\Gamma \Delta^D|$ in $\hat{H}_C$. Transitions between states with significantly larger or smaller energy gaps are suppressed for fixed Γ . Since the energy gaps of vertex pairs typically vary throughout the graph, we can steer the walker's movement effectively by selecting Γ to activate only the desired transitions in the graph.

For all EC, TSP, and GO problems under investigation, we noticed empirically that the distribution of the energy levels acquires a characteristic “onion shape” [see Fig. 3(a)] as soon as the problem instance is sufficiently complex (i.e., the number of qubits is sufficiently large). Additionally, the largest energy gap $\Delta^C_{j,k}$ from vertex $|j\rangle$ to a lower energy vertex $|k\rangle$ increases approximately monotonically as a function of its energy level E_j [see Fig. 3(b)]. Consequently, large values of Γ are usually optimal in the regime of high-energy states, while small values are preferable near the solution state. Choosing a fixed value for Γ has the disadvantage of limiting the maximally achievable success probability because not all edges can sufficiently contribute to the amplitude transport. In particular, only amplitude transfers within a fixed energy range can be addressed, generally prohibiting states located at high energies from transporting amplitude to the solution state.

The strategy we propose in this paper is based on an energy-dependent hopping rate, where $|\tilde{\Gamma}(E) \cdot \Delta^D|$ corresponds to the average $\langle \Delta^C \rangle(E)$ of the largest energy gaps of $\hat{H}_C$ at energy level E in the graph. The approach is to confine the walker to a gradually shrinking energy region around the solution state by progressing monotonically from high-energy to low-energy optimal hopping rates $\tilde{\Gamma}$. Therefore, we set $\Gamma(t) = \tilde{\Gamma}(E(t))$, where $E(t)$ is a monotonic sweep, and define

$$\hat{H}_{\text{GQW}} = \Gamma(t) \cdot \hat{H}_D + \hat{H}_C. \quad (15)$$

Equation (14) shows that the frequency of the local Rabi oscillations depends on the size of the energy gaps, resulting in faster amplitude transfers at larger gaps [see Fig. 3(b)]. Thus, the GQW needs to spend less time driving transitions at high energies than at small energies to avoid amplitude going back to high energy state. Consequently, the energy sweep must be rescaled according to $\langle \Delta^C \rangle(E)$, yielding the rescaled time

$$s(t) = \frac{\int_{E_{\text{max}} + (E_{\text{min}} - E_{\text{max}}) \cdot t/T}^{E_{\text{max}}} \langle \Delta^C \rangle(E) dE}{\int_{E_{\text{min}}}^{E_{\text{max}}} \langle \Delta^C \rangle(E) dE} \quad (16)$$

with $E(t) = E_{\text{max}} + (E_{\text{min}} - E_{\text{max}}) \cdot s(t)$, where E_{min} and E_{max} denote the lowest and highest energy levels of $\hat{H}_C$, respectively, and T is the total evolution time. At $t = 0$, the algorithm starts with a relatively high hopping rate [$\Gamma = \tilde{\Gamma}(E_{\text{max}})$] that maximizes amplitude transfers only at high-energy levels and suppresses dynamics at low and intermediate energies. As the hopping rate is decreased over time, the system starts to perform amplitude transfers at lower energy levels. Simultaneously, subspaces at higher energies become gradually detuned again, suppressing the action of the driving Hamiltonian. This is essential as it prevents the backpropagation of probability into high-energy states and enables us to actively guide the walker towards low-energy states.

The primary advantage of this approach lies in the establishment of a continuous amplitude flow from all computational basis states towards the solution state, consequently overcoming the previous limitation on the maximally achievable solution quality. Furthermore, the latter is no longer subject to the complex oscillations, which required a precisely chosen evolution time T [cf. Fig. 2(c)]. Instead, $S_q(T)$ follows a collective sinusoidal oscillation of Eq. (14), where the




FIG. 4. Performance comparison between the GQW (blue) and the QW (red) on the $N = 21$ qubit EC problem depicted in Fig. 3. The GQW dynamically adjusts the relative strength Γ of the driving and problem Hamiltonian in Eq. (15) based on the average $\langle \Delta^C \rangle(E)$ of the largest energy gaps in $\hat{H}_C$. In contrast, the QW employs a fixed hopping rate Γ that has been pre-optimized [see Fig. 3(b)]. The performance is evaluated using the solution quality S_q , as defined in Eq. (4). Note that the results presented here are indicative and require knowledge of the complete energy spectrum.

selection of $\Gamma(t)$ ensures that active subspaces oscillate with similar Rabi frequencies while high-energy subspaces become gradually detuned, hence suppressing the backpropagation of probability amplitude. By varying the evolution time T , the speed at which $\tilde{\Gamma}(E)$ is swept can be adjusted, thereby altering the duration that groups of subspaces remain active. As a result, $S_q(T)$ depends solely on the order of magnitude of T , rather than its precise value.

To evaluate the efficiency of a GQW, we compare its performance to a conventional QW on an EC problem instance comprising $N = 21$ qubits. Figure 3(b) shows the average $\langle \Delta^C \rangle(E)$ of the largest energy gaps between connected vertices at each energy level E . We obtain $|\tilde{\Gamma}(E) \cdot \Delta^D|$ by fitting a polynomial of degree 6 to the data points [see the black arrows in Fig. 3(b)]. The bottom axis indicates the nonlinear sweep of $E(t)$. In comparison, for the QW, we determined Γ_{opt} by sampling a classical optimizer (see dashed black curve). Γ_{opt} corresponds to the optimal hopping rate that yields the highest success probability over $T \in [0.1, 10.0]$.

Figure 4 compares the performance of the GQW and the QW as a function of the total evolution time T . For $T \leq 2$, both strategies exhibit similar behavior, showing a rapid increase in solution quality, with the QW achieving its peak at $T = 1.5$. Notably, for $T \leq 2.5$, the GQW yields lower solution qualities compared to the QW. This discrepancy arises from the relatively small optimal hopping rate of the QW, enabling it to focus its evolution on amplitude transitions near and into the solution state. In contrast, the GQW considers the entire graph and therefore spends the initial part of its evolution at high and intermediate energy levels [the right part of Fig. 3(b)]. This approach leads to very short evolutions at each energy level for small T , causing only fractions of the probability amplitudes to be transported towards z_{opt} . However, as T increases, the situation changes, and the GQW obtains superior solution qualities to the QW for $T \geq 2.5$.

Here, the QW exhibits oscillatory behavior between 10^{-5} to 10^{-2} . In contrast, the GQW follows a monotonic increase, where $S_q(T)$ saturates at approximately 21%, outperforming the QW by approximately one order of magnitude.

The observed saturation results from imbalances within the local subspaces at small E . As a large amount of amplitude accumulates in the ground state, the system progressively deviates from the state of equal local superposition. Consequently, driving these subspaces with their respective optimal hopping rate $\tilde{\Gamma}(E)$ eventually redirects amplitude back into the higher energy state, thus limiting $S_q(T)$ (see also Appendix B).

D. Practical guided quantum walk

The previous section has shown the potential of GQWs for solving combinatorial optimization problems, by adjusting the relative strength of the two Hamiltonians in $\hat{H}_{\text{GQW}}$ based on the walker's position in the graph. The main benefits are an increased maximum solution quality and a suppression of complex oscillations in $S_q(t)$, providing good results for arbitrary T and N .

Of course, obtaining the optimal function $\tilde{\Gamma}(E(t))$ is generally not efficiently possible, because it requires knowledge of the entire energy spectrum of $\hat{H}_C$. In order to still make use of the promising methodology described above, we propose a variational ansatz, in analogy to other quantum optimization algorithms [42–44]. The idea is to imitate the optimal distribution $\tilde{\Gamma}(E)$ by a function $\tilde{\Gamma}(E, \lambda)$, which is tuned using a set of M hyperparameters $\lambda = (\lambda_1, \dots, \lambda_M)$. Note that henceforth we consider only linear sweeps of the energy spectrum, i.e., $E(t) = (E_{\min} - E_{\max}) \cdot t/T + E_{\max}$, thus encoding the sampling speed in the shape of $\tilde{\Gamma}$. The hope is that as long as $\tilde{\Gamma}(E, \lambda)$ describes $\tilde{\Gamma}(E)$ closely enough, similar dynamics to Sec. III C can be obtained. In fact, introducing hyperparameters into the algorithm even enables the guided quantum walk to overcome its limitations at small and large T . For instance, at small T , the GQW could selectively model $\tilde{\Gamma}(E)$ only up to $E < E_{\max}$, thereby operating solely on a subspace of the entire graph near $|z_{\text{opt}}\rangle$. Conversely, at large T , the GQW can prevent the transfer of amplitude into higher energy states when operating at small E by compensating the imbalances within the local subspaces through a reduction of the final hopping rate $\tilde{\Gamma}(E = E_{\min})$.

The proposed algorithm employs a hybrid quantum-classical ansatz, in which a classical optimizer adjusts the set of hyperparameters λ based on the minimization of the energy expectation value E_Ψ [see Eq. (3)] of the final quantum state $|\Psi\rangle$ obtained by a quantum device performing the GQW. Note that the number M of hyperparameters is fixed, and we investigate the impact of this optimization phase on the total run time in Sec. IV C.

We propose a function $\tilde{\Gamma}(E, \lambda)$ based on cubic Bézier curves. We chose Bézier curves instead of simple polynomials because we expect the optimal hopping rate to be smooth and monotonically decreasing in E . Although polynomials can produce such functions for $E \in [E_{\min}, E_{\max}]$, their parameters are generally hard to tune, as small changes can lead to substantially different functions. In contrast, Bézier curves are much easier to optimize since their general shape can

be predetermined. Moreover, the latter varies continuously and sufficiently slowly in its parameters, resulting in a sufficiently smooth search space for λ . Due to these properties, we strongly encourage their use in other fields of quantum computing, such as optimizing annealing schedules [30–32] or deriving optimal parameter sets for the quantum approximate optimization algorithm (QAOA) [2,42,45–47].

Cubic Bézier curves are based on Bernstein polynomials and are defined through four control points $C_i = (x_i, y_i)^T$ in a two-dimensional plane as

$$x(\tau) = x_0(1 - \tau)^3 + 3x_1(1 - \tau)^2\tau + 3x_2(1 - \tau)\tau^2 + x_3\tau^3, \quad (17)$$

$$y(\tau) = y_0(1 - \tau)^3 + 3y_1(1 - \tau)^2\tau + 3y_2(1 - \tau)\tau^2 + y_3\tau^3. \quad (18)$$

Choosing $x_0 < x_1, x_2 < x_3$, the curve can be mapped into one dimension, by solving $t = x(\tau)$ for $\tau(t)$ and inserting it into y . Based on the observations made in Sec. III C, we set $C_0 = (0, 1)$ and $C_3 = (1, 0)$ to ensure a monotonically decreasing function, with four hyperparameters $\lambda_{1,2,3,4} \in (0, 1)$ controlling its shape, i.e., $C_1 = (\lambda_1, \lambda_2)^T$ and $C_2 = (\lambda_3, \lambda_4)^T$. The hopping rate is then given by $\tilde{\Gamma}(t, \lambda) = y(t) \times 10^{2 \cdot \lambda_5} + (1 - y(t)) \times 10^{-3 \cdot \lambda_6}$, where we introduced two additional hyperparameters, $\lambda_5, \lambda_6 \in [0, 1]$, that describe the boundary conditions $\tilde{\Gamma}(t = 0, \lambda) = 10^{2 \cdot \lambda_5}$ and $\tilde{\Gamma}(t = T, \lambda) = 10^{-3 \cdot \lambda_6}$. λ_5 determines the energy range in which the GQW operates in the graph (see discussion for small T) and λ_6 controls the detuning of the local subspaces near the solution state (see discussion for large T). Since the optimal scale of $\tilde{\Gamma}(t, \lambda)$ can vary significantly between different problem instances, we chose an exponential rescaling to simplify the optimization landscape. Note that the parameters -2 and 3 were found suitable for the problems under investigation, but generally depend on the energy scale of the cost Hamiltonian $\hat{H}_C$. In Fig. 5 we present $\tilde{\Gamma}(t, \lambda)$ for the aforementioned EC problem and $T = 10$.

E. From quantum walk to quantum annealing

The LAT theory, used in the derivation of the GQW (see Sec. III C), offers a perspective on the working principle of quantum annealing (QA) and its relationship to QWs. QA belongs to the class of continuous-time quantum optimization algorithms that rely on an adiabatic transition from the driver Hamiltonian $\hat{H}_D$ [i.e., $\Gamma(t = 0) \rightarrow \infty$] to the problem Hamiltonian $\hat{H}_C$ [i.e., $\Gamma(t = T) = 0$] throughout the time evolution. According to the adiabatic theorem of quantum mechanics, if this transition occurs sufficiently slowly, and the system is initially prepared in the ground state of $\hat{H}_D$, it will remain in the instantaneous ground state of the combined Hamiltonian $\hat{H}_{GQW}$ in Eq. (15), ultimately reaching the solution state $|z_{\text{opt}}\rangle$.

To explore the relationship between QWs and QA, we use the GQW to examine the strategies governing optimal quantum evolutions across different time intervals T . Figure 7 presents simulation results of the GQW applied to an $N = 15$ qubit EC problem, covering short ($T = 0.5$), intermediate ($T = 2.0$), and long ($T = 12$) evolutions. Panels (a)–(c) depict the average measurement probabilities of the




FIG. 5. Optimal hopping rate $\tilde{\Gamma}(t, \lambda)$ (blue) for the $N = 21$ qubit EC problem shown in Fig. 3, with a total evolution time of $T = 10$. The hopping rate is represented by a cubic Bézier curve [see Eq. (18)] and is defined by four control points $C_{0,1,2,3}$ based on the hyperparameters λ_{1-4} , along with the boundary conditions $\tilde{\Gamma}(t = 0, \lambda) = 10^{2 \cdot \lambda_5}$ and $\tilde{\Gamma}(t = T, \lambda) = 10^{-3 \cdot \lambda_6}$. The curve is obtained by optimizing a set of six hyperparameters λ using the classical Nelder-Mead optimizer [48,49], employing $N_{\text{opt}} = 100$ optimization steps starting from a random initial configuration.

energy levels E_C within the problem Hamiltonian $\hat{H}_C$. Panels (d)–(f) showcase the evolution of the instantaneous energy levels E_{GQW} of the combined Hamiltonian $\hat{H}_{GQW}$. Furthermore, Fig. 6 presents the optimal range of $\tilde{\Gamma}(t, \lambda)$ as a function of the total evolution time T for the same EC problem.




FIG. 6. Ranges of the optimal hopping rate $\tilde{\Gamma}(t, \lambda)$ (solid lines) for an $N = 15$ qubit EC problem as a function of the total evolution time T . $\tilde{\Gamma}(t = 0, \lambda)$ and $\tilde{\Gamma}(t = T, \lambda)$ denote the initial and final hopping rate, respectively, as determined by the hyperparameters λ_5 and λ_6 (see Sec. III D). Note that $\tilde{\Gamma}(t, \lambda)$ is scaled by Δ^D to map to the energy gaps of the local subspaces in the hypercube graph. The heatmap presents the evolution of the measurement probabilities averaged over states with equal maximum energy gaps $\Delta_{j,k}^C$ [see Fig. 3(b)]. Dotted lines mark the range of energy gaps of local subspaces involving the solution state.




FIG. 7. Optimized quantum evolutions performed by the GQW on the $N = 15$ qubit EC problem shown in Fig. 6 for short [$T = 0.5$, (a) and (d)], intermediate [$T = 2.0$, (b) and (e)], and long [$T = 12$, (c) and (f)] evolution times. Panels (a)–(c) illustrate the average measurement probabilities of the energy levels E_C of the problem Hamiltonian $\hat{H}_C$. Panels (d)–(f) display the evolution of the instantaneous energy levels E_{GQW} of the combined Hamiltonian $\hat{H}_{GQW}$ [Eq. (15)]. The blue and red curves represent the instantaneous ground state and the first excited state, respectively. Dashed lines are included in all panels to indicate the instantaneous energy expectation values of the system with respect to $\hat{H}_C$ [panels (a)–(c)] and $\hat{H}_{GQW}$ [panels (d)–(f)]. The final solution qualities are $S_q = 0.7\%$ ($T = 0.5$), $S_q = 13.6\%$ ($T = 2.0$), and $S_q = 91.3\%$ ($T = 12.0$).

Short evolutions. In the case of short evolutions, the optimal hopping rate schedule $\tilde{\Gamma}(t, \lambda)$ determined by the GQW maintains a near-constant relationship between $\hat{H}_D$ and $\hat{H}_C$ in Eq. (15). Consequently, the instantaneous energy levels mostly remain unchanged during the evolution, and the system starts and ends in a superposition of the instantaneous basis states [see the dashed line in Fig. 7(d)]. For $T = 0.5$, $\tilde{\Gamma}(t, \lambda)$ decreases from 0.6, approximately equivalent to $\tilde{\Gamma}(E_{min})$, to 0.3. This transition initially drives the subspaces involving $|z_{opt}\rangle$, hence resembling a QW-like procedure, before progressively detuning them to prevent the backpropagation of amplitude into higher energy states (see Fig. 6). The latter is necessary as these subspaces gradually move away from a close-to-equal superposition state (see Appendix B). Figure 7(a) demonstrates that this strategy results in a guided movement of amplitude towards $|z_{opt}\rangle$ for $E_C \leq 12$, evident from the emerging gradient in the measurement probabilities. Probability amplitude originating from intermediate and high energy levels, on the other hand, fails to reach the solution state and instead becomes trapped at these energies within the graph [see horizontal dark and bright stripes in Fig. 7(a)]. This highlights the inherent limitations of QWs.

Intermediate evolutions. In the case of intermediate evolution times, the optimal hopping rate schedule operates over

a wider range of values, with the initial hopping rate $\tilde{\Gamma}(t = 0, \lambda)$ increasing monotonically as a function of T (see Fig. 6). In doing so, the GQW extends the subset of the hypercube graph in which amplitude is actively guided through the local subspaces. This is evident from the increased final success probability and the gradual transport of amplitude from high-to low-energy states for $t \leq 0.5$ [see Fig. 7(b)]. The presence of bright stripes at $E_C \geq 10$ indicates that the GQW is still operating on a subset of the hypercube graph, causing amplitude from high-energy states to become trapped at intermediate energy levels. This demonstrates that even for intermediate evolution times, it is more advantageous to neglect amplitude at high-energy states and focus on amplitude transfers near the solution state. Interestingly, also the final hopping rate $\tilde{\Gamma}(t = T, \lambda)$ increases for longer evolution times T up to a maximum at $T \approx 2$. Subsequently, for $T > 2$, $\tilde{\Gamma}(t = T, \lambda)$ declines, again detuning the local subspaces at small E . The shape of $\tilde{\Gamma}(t = T, \lambda)$ is likely influenced by the density of states of $\hat{H}_C$. As Fig. 3(a) illustrates for an $N = 21$ qubit EC problem, the state density of $\hat{H}_C$ for our EC instances rapidly increases for smaller E , reaching a peak at E_p , followed by an exponential decline. Consequently, as $\tilde{\Gamma}(t = 0, \lambda)$ initially increases, the accessible amplitude likely grows faster than the transport of amplitude into $|z_{opt}\rangle$ within T . This keeps the local subspaces surrounding $|z_{opt}\rangle$ closer to an equal superposition state, re-

quiring less final detuning. The peak of $\tilde{\Gamma}(t = T, \lambda)$ aligns with the point where $\tilde{\Gamma}(t = 0, \lambda) \approx \tilde{\Gamma}(E_p)$. As $\tilde{\Gamma}(t = 0, \lambda)$ further increases, the growth of accessible amplitude slows down, causing the local subspaces at lower E to move away from an equal superposition state. Consequently, a larger final detuning is required, resulting in a smaller $\tilde{\Gamma}(t = T, \lambda)$.

Long evolutions. For long evolutions, the optimal hopping rate schedule determined by the GQW resembles a QA-like schedule by transitioning nearly entirely from $\hat{H}_D$ to $\hat{H}_C$ in Eq. (15) (see Fig. 6). Consequently, the system follows the instantaneous ground state of $\hat{H}_{\text{GQW}}$ throughout the evolution, indicated by the dashed line in Fig. 7(f). Throughout this process, the GQW effectively drives amplitude transfers across the entire graph, initiating from high-energy levels and moving towards low energies, evident by the absence of horizontal stripes in Fig. 3(c). Notably, the confinement of amplitude occurs exponentially in t , indicating a nonlinear sweep through the energy spectrum [cf. Eq. (16)]. Furthermore, the propagation of amplitude follows a wave-like pattern, as illustrated by the dashed line in Fig. 7(c). This pattern arises due to the confinement of amplitude in a decreasing number of vertices, leading to deviations from the equal superposition states in the local subspaces. This results in a temporary backpropagation of amplitude into higher energy states (see Appendix B). Nevertheless, the continuous decrease of the hopping rate $\tilde{\Gamma}(t, \lambda)$ ensures, on average, the transportation of amplitude into the solution state.

The hopping rate schedules derived from the GQW not only emphasize the intrinsic connection between QWs and QA but also highlight the existence of optimal quantum evolutions that extend beyond the scopes of these two algorithms. While a QW-like strategy proves optimal for short evolutions and a QA-like procedure is favored for long evolutions, our investigations reveal that intermediate values of T necessitate a combination of both strategies to maximize the solution quality.

This observation can be explained through LAT theory (see Sec. IIIC and Appendix B), as QWs and QA can be viewed as two distinct formulations of the same underlying concept. Both approaches aim for the optimal transfer of probability amplitude within local subspaces of the graph. QWs achieve this by employing a constant, small hopping rate Γ , focusing exclusively on direct transfers into the solution state during short evolutions. However, this strategy becomes less effective for long evolutions, where sufficient time is available to guide amplitude at higher energy levels as well. Consequently, QA guides amplitude through the entire graph by linking multiple local QWs together, employing a continuously decreasing hopping rate $\tilde{\Gamma}(t)$. Thus, QWs and QA represent the two extremes within the GQW framework, with one concentrating solely on subspaces around the solution state [i.e., $\tilde{\Gamma}(t = 0) \approx \tilde{\Gamma}(t = T) \ll 1$] and the other considering the entire graph [i.e., $\tilde{\Gamma}(t = 0) \gg 1$ and $\tilde{\Gamma}(t = T) \ll 1$, cf. Fig. 5].

The GQW operates in the transition region between these two extremes, striking a balance between the number of guided local subspaces (i.e., the amount of guided amplitude) and the time spent at each energy level (i.e., the amount of amplitude transferred within each local subspace) for a given T . As will be discussed in Sec. IV B, these intermediate

evolutions, which surpass the limitations of adiabatic time evolutions, might be capable of effectively solving large-scale optimization problems.

IV. RESULTS

In this section, we present a comprehensive performance analysis of the GQW. We compare the solution quality S_q both as a function of the evolution time T and the system size N to the conventional QW and QA (see Sec. IV A). Section IV B provides scaling results of the three algorithms in the regime of large problem instances (i.e., $N \gg 1$). Our analysis indicates that the GQW achieves significantly higher solution qualities within a linearly growing timespan, rendering it a promising candidate for near-term quantum devices. Finally, in Sec. IV C we address the impact of the classical optimization phase on the total run time, demonstrating that the average time-to-solution scales better by a factor of ≈ 2 (≈ 4) compared to linear QA (QW).

A. Comparison of GQW, QW, and QA

We assess the effectiveness of the GQW through numerical simulations carried out on the JUWELS Booster supercomputer at the Jülich Supercomputing Centre of the Forschungszentrum Jülich [40]. We compare its performance against a conventional QW and QA. Our simulations consider EC and GO instances with problem sizes $N \in \{12, 15, 18, 21, 24, 30\}$, as well as TSP instances with $N \in \{9, 16, 25\}$ qubits. To demonstrate the generality of our findings, we examine 10 randomly generated problem instances for each problem type and size. We remark that the energy spectrum of each problem has been obtained classically, providing the total energy range and the individual energies of the valid states for the calculation of S_q [see Eq. (4)]. Note, however, that this is done only for the purpose of benchmarking, and it is not required to apply the GQW in a practical scenario.

To obtain the final quantum state of the system at the end of each algorithm, we use the second-order Suzuki-Trotter product formula algorithm [51,52] to solve the time-dependent Schrödinger equation, $i \partial_t |\Psi(t)\rangle = H(t) |\Psi(t)\rangle$, with a time step of 10^{-5} and total evolution times T ranging between 0.1 and 12.0. Given the consistency of our findings, we present the results for the EC instances in Fig. 8, showing the obtained solution qualities averaged within each system size, along with the standard deviations. The analogous results for the TSP and GO problems are available in Appendix C and support the same conclusions.

Guided quantum walk. The GQW is optimized for each problem instance and evolution time T individually by tuning its six hyperparameters λ to minimize the energy expectation value E_Ψ [see Eq. (3)] of the final quantum state $|\Psi\rangle$. Each parameter set is thereby selected from a pool of $N_{\text{rep}} = 100$ repetitions of the Nelder-Mead classical optimizer, where in each sample the six parameters are initialized randomly and adjusted a maximum of $N_{\text{opt}} = 100$ times. We choose this approach to ensure that the algorithm converges to a near-optimal minimum in the parameter search space. However, we note that a sufficient set of parameters is typically found




FIG. 8. Performance comparison of (a) the GQW, (b) the QW, and (c) QA as a function of the total evolution time T on EC problems with problem sizes $N \in \{12, 15, 18, 21, 24, 30\}$ (see legend). Panel (c) presents simulation results for both linear QA (solid lines) and optimized locally adiabatic QA [50] (dashed lines). The latter uses a rescaled time based on the size of the instantaneous energy gap during the annealing process, assuming complete knowledge about the optimization problem. The performance of all algorithms is assessed based on the solution quality S_q defined in (4). Each problem size N is evaluated using a set of 10 randomly generated problem instances, and the presented data corresponds to the geometric mean solution qualities (solid curves) along with the geometric standard deviations (shaded areas). Note that the GQW and the QW undergo a prior classical optimization phase. The dotted gray lines indicate horizontal cuts ($S_q \in \{1\%, 10\%, 90\%\}$), which are detailed in Fig. 9. The inset in panel (a) provides a zoomed-in view of the region corresponding to large evolution times T . Additionally, the dashed lines in panel (b) indicate the stationary solution quality that the damped oscillation of $S_q(T)$ approaches for the QW at large T .

within the first $N_{\text{rep}} = 20$ repetitions. In Sec. IV C, we discuss the impact of this optimization phase on the total run time.

Figure 8(a) presents the simulation results of the GQW, showing the scaling of the solution quality $S_q(T)$ as a function of the evolution time T . Across various problem types and sizes, we observe consistent patterns in the scaling of $S_q(T)$, corresponding to the three regimes of evolutions (see Sec. III E).

Initially, the solution quality exhibits rapid growth, matching the results obtained by the QW algorithm for $T \leq 0.5$ [see Fig. 8(b)]. At $S_q \approx 10\%$, however, the scaling decelerates, with solution qualities above 70% for all investigated problem instances at $T = 12$. This scaling behavior is primarily influenced by the energy range considered during the algorithm's evolution. Since the GQW cannot sufficiently transport amplitude from all states towards the solution state at short evolution times T , the algorithm focuses its efforts on a subset of the graph to maximize the accumulation of amplitude into $|z_{\text{opt}}\rangle$. As T increases, the GQW accesses a larger number of states, and consequently, $S_q(T)$ seems to scale according to the amount of accessed amplitude (cf. discussion in Sec. III E).

Quantum walk. For simulating the conventional QW, we employ a procedure similar to that of the GQW. Specifically, we determine the optimal hopping rate Γ for each problem instance and T separately, using the Nelder-Mead optimizer with $N_{\text{rep}} = 100$ repetitions and a maximum of $N_{\text{opt}} = 100$ parameter evaluations, aiming to minimize the final energy expectation value E_Ψ .

Figure 8(b) illustrates the evolution of $S_q(T)$ as a function of T for the QW. Across all problem instances, $S_q(T)$ exhibits a damped oscillation pattern, converging to a stationary solution quality (indicated by dashed lines) below 1 at $T \geq 12$. Notably, this stationary solution quality decreases exponentially in the problem size N , because the QW, with a constant hopping rate Γ , can only drive a few local subspaces sufficiently (cf. discussion in Sec. III C). Consequently, the GQW surpasses the QW in terms of performance even for short evolution times (e.g., $T = 0.5$), highlighting the significance of local adjustments to Γ already at short time scales. Furthermore, it is noteworthy that the QW is the only algorithm investigated that fails to achieve solution qualities greater than 20% for any problem instance and T .

Quantum annealing. In the context of QA, we examine two annealing schemes: a linear annealing scheme represented by $\Gamma(t) = (1 - s(t))/s(t)$, where $s(t) = t/T$, and an optimized locally adiabatic schedule [50] employing a rescaled time $s^{\text{opt}}(t)$. The latter is determined by numerically computing the instantaneous energy gap $\Delta(s)$ between the ground and first excited state of $\hat{H}_{\text{GQW}}$ across various values of s , followed by solving $ds^{\text{opt}}/dt \propto \Delta^2(s^{\text{opt}})$ to derive $s^{\text{opt}}(t)$. This approach yields an optimized schedule that decelerates the annealing process in regions with small energy gaps while accelerating it elsewhere. A comprehensive discussion on this approach is given in [50]. It is important to note that the linear scheme represents the baseline performance of QA, where no prior optimization phase is required. Conversely, the optimized locally adiabatic schedule mirrors the theoretical best performance

of QA for $T \gg 1$, assuming complete knowledge about the optimization problem.

Figure 8(c) presents the solution qualities achieved with linear QA (solid lines) and optimized locally adiabatic QA (dashed lines) as a function of T . As expected from the adiabatic theorem, $S_q(T)$ increases with the evolution time T for both approaches, with optimized locally adiabatic QA achieving up to one order of magnitude higher solution qualities at $T = 12$ than linear QA. Interestingly, for the $N \in \{21, 24\}$ qubit problems, linear QA beats the optimized locally adiabatic schedule for $T \leq 4$. This is likely caused by diadiabatic transitions in the context of fast annealing. When comparing QA to both the GQW and the QW, we observe a significantly steeper increase in solution quality for the latter two, underscoring the importance of focused amplitude transfers near the solution state for short time scales. For intermediate and long evolution times, QA surpasses the QW. The GQW, however, outperforms both algorithms across all investigated problem instances and T . Interestingly, even in the case of long evolutions (e.g., $T = 12$) the optimized locally adiabatic QA fails to match the solution qualities achieved by the GQW, indicating the existence of optimal schedules beyond the adiabatic theorem.

B. Performance on large problem instances

The previous section has demonstrated the efficiency of the proposed guiding procedure for quantum walks in solving combinatorial optimization problems. The GQW outperforms both the QW and QA by providing significantly higher solution qualities across all studied problem instances and evolution times T . However, to determine how these algorithms compare for real-world problem sizes ($N \gg 1$) that exceed the capabilities of our numerical simulations, we analyze the scaling of the evolution time $T_{S_q}(N)$ required to reach a solution quality $S_q \in \{1\%, 10\%, 90\%\}$ as a function of the problem size N . The corresponding data is shown in Fig. 9(a) together with linear ($a + b \cdot N$) and exponential ($a \cdot 2^{b \cdot N}$) fits to the data points obtained by the GQW and QA, respectively. The QW was excluded from this analysis, as it could not reach the required solution qualities. It is worth noting that the QW is not designed as a single-shot algorithm, and we will discuss the multishot QW in the subsequent section.

The data demonstrates that quantum annealing (QA) exhibits exponential scaling for both linear and optimized locally adiabatic schedules, with scaling coefficients $b \in [0.12, 0.26]$ for the three solution quality levels S_q . This is in line with the expectation derived from the adiabatic theorem, stating the instantaneous energy gaps, which shrink exponentially in N , demand an exponentially slow annealing process for the system to remain in its instantaneous ground state [53,54]. Notably, the optimized locally adiabatic QA [50] achieves the best scaling with $b = 0.12$ for $S_q = 0.9$, showing that local adjustments to the annealing speed can significantly reduce the annealing time needed to reach high solution qualities. In contrast, $T_{S_q}(N)$ follows a linear scaling in N for the GQW. This can be explained by the fact that the depth of a hypercube graph scales linearly in the number of qubits N (i.e., the largest Hamming distance between any two states cannot be larger than N). Hence, the GQW must




FIG. 9. Scaling of the total evolution time $T_{S_q}(N)$ required to achieve a specified solution quality S_q on EC problems as a function of the problem size N for the GQW (circles and solid lines), linear QA (triangles and dotted lines), and optimized locally adiabatic QA [50] (squares and dashed lines). Colors indicate results for various target solution qualities: $S_q = 1\%$ (blue), $S_q = 10\%$ (red), and $S_q = 90\%$ (green). The solid lines depict linear fits ($a + b \cdot N$), while the dashed and dotted lines represent exponential fits ($a \cdot 2^{b \cdot N}$) applied to the data points (see legend).

at most drive amplitude transfers within N local subspaces to transport amplitude from any state into the solution state. Since the Hamming distance from z_{opt} to each computational basis state seems to correlate positively with the energy gaps of these states [cf. Fig. 3(b)], these amplitude transfers are performed simultaneously for all states, yielding a linear in N evolution time T . In the regime of large T , the GQW can thus be seen as an optimized annealing schedule, where the algorithm selectively spends more time in critical parts of the graph, while progressing faster elsewhere.

Although more extensive studies are required to verify the observation of a linear scaling, our findings identify the GQW as a highly efficient algorithm that achieves high solution qualities within short time scales. This makes the GQW an attractive choice for near-term quantum devices with limited coherence times.

C. Classical optimization phase

We have examined the performance of the GQW, showcasing its capability to achieve optimal quantum evolutions by guiding amplitude transfers locally in the hypercube graph. Our findings indicate a linear scaling of the total evolution time $T_{S_q}(N)$ with respect to the problem size N , surpassing the performance of both the QW and the QA when provided with an optimal set of hyperparameters λ . However, our investigation has primarily focused on the quantum aspect of this hybrid algorithm, without counting the classical optimization phase responsible for fine-tuning the six hyperparameters λ .

In Fig. 10, we investigate the influence of the classical optimization phase on the GQW by analyzing the scaling of the average solution quality S_q as a function of the total evolution time T and the number of parameter evaluations N_{opt} . The latter refers to the number of iterations performed




FIG. 10. Scaling analysis of the solution quality S_q as a function of the total evolution time T and the number of parameter evaluations N_{opt} performed by the Nelder-Mead algorithm during the classical optimization phase of the hyperparameters λ in the GQW. Panels (a)–(e) depict the results for EC problems of sizes $N \in \{12, 15, 18, 21, 24\}$, respectively. The data is sampled with a granularity of $\Delta T = 1$ and $\Delta N_{\text{opt}} = 5$. Each data point represents the average solution quality obtained from 100 repetitions of the optimization phase. Furthermore, the data is averaged across 10 random problem instances for each problem size N . The blue circles mark the configurations of T and N_{opt} that yield the lowest $TTS_{99.99\%}$ [see Eq. (19)]. The dashed curves indicate exponential fits to the $S_q = 0.5$ contour. In panel (f) we show the scaling of the optimal (smallest) $TTS_{99.99\%}$ achieved by the GQW (blue), linear QA (red), and the QW (green) as a function of the problem size N . The solid lines correspond to exponential fits ($a \cdot 2^{b \cdot N}$) applied to the data points.

by the Nelder-Mead algorithm during the initial optimization phase. Panels (a)–(e) present the averaged results for the EC problems with sizes $N \in \{12, 15, 18, 21, 24\}$, respectively.

The data reveals consistent characteristics across all investigated problem instances. Specifically, we observe that the number of parameter evaluations $N_{\text{opt}}^{S_q}$ necessary to reach a minimum solution quality S_q decreases exponentially as a function of the total evolution time T [see dashed curves in panels (a)–(e)]. Additionally, for a fixed value of T , $N_{\text{opt}}^{S_q}$ scales exponentially in the problem size N . These observations indicate that for short evolutions, where the GQW prepares QW-like hopping rate schedules, the parameter search space tends to be more complex, thereby requiring longer optimization phases. This complexity arises, as the GQW is considering only a few subspaces, making it crucial to precisely set the hyperparameters λ to effectively drive amplitude transfers within these subspaces. Notably, the performance of the GQW is particularly sensitive to the choice of λ_4 and λ_5 in this regime of T . On the other hand, in the case of long evolutions, high-quality solutions can be obtained with just a few iterations from the classical optimizer. This is

because deviations from the optimal schedule have minimal impact on the overall evolution of the quantum system, as the optimal hopping rate schedules approach a QA-like evolution, and success is increasingly guaranteed by the adiabatic theorem.

To incorporate the classical optimization phase into our performance evaluations, we consider the time-to-solution

$$TTS_{P_{\text{target}}} = \frac{\ln(1 - P_{\text{target}})}{\ln(1 - P_{\text{gs}})} \cdot T + N_{\text{opt}} \cdot T, \quad (19)$$

where P_{gs} and T denote the success probability and the evolution time of the algorithm, respectively. $TTS_{P_{\text{target}}}$ represents the total run time required to measure the solution state $|z_{\text{opt}}\rangle$ at least once, with a probability of P_{target} , over multiple runs of the algorithm. Note that the optimization phase is accounted for through the offset $N_{\text{opt}} \cdot T$.

In Fig. 10(f), we show a comparison of the scaling of the optimal (smallest) $TTS_{99.99\%}$ achieved by the GQW, the QW, and linear QA as a function of the problem size N . Note that optimized locally adiabatic QA is excluded from this analysis, as it requires knowledge of the full spectral information about the optimization problem. Exponential fits

$(a \cdot 2^{b \cdot N})$ are included as a reference. The data reveals that for small problem sizes, both the QW and linear QA demonstrate faster convergence to the solution state compared to the GQW, due to the GQW's initial optimization phase. However, for $N \geq 15$ and $N \geq 21$, respectively, the GQW surpasses both algorithms with a scaling factor of $b = 0.12$, which is approximately four (two) times better than $b = 0.47$ ($b = 0.22$) for the QW (linear QA). Although the initial optimization phase in the GQW leads to an exponential scaling of $TT_{S99.99\%}$, the GQW's ability to focus solely on a subset of the graph for intermediate values of T enables a significantly more efficient utilization of computational resources compared to the other algorithms. Furthermore, the fact that the exponential scaling is shifted into the optimization phase, while the single run times T scale at most linearly with N (see Sec. IV B), offers the opportunity to distribute the optimization phase across multiple quantum computing devices, thereby providing an option to parallelize the process. This is a feature generally not feasible for QA due to the exponential scaling of T , but it could potentially allow solving large optimization problems on near-term quantum devices.

V. CONCLUSIONS

We have utilized the theory of local amplitude transfer (LAT), which offers a perspective on the operational principles of quantum annealing (QA) and quantum walks (QWs) beyond the adiabatic theorem, while also providing insights into the design of optimal quantum evolutions. The theory is rooted in the description of a quantum evolution within the eigenspace of the problem Hamiltonian $\hat{H}_C$. In this context, the search space is represented as a graph G where the states are interconnected through the driving Hamiltonian $\hat{H}_D$ [see Fig. 3(a)]. By decomposing G into two-dimensional subspaces spanned by pairs of eigenstates, we have demonstrated that probability amplitude traverses the graph through a sequence of local Rabi oscillations occurring within these subspaces. The amplitude of these oscillations depends on the relative strengths of the local driving and problem Hamiltonians, controlled by the hopping rate Γ .

We have highlighted that for sufficiently complex problems, the average energy gap of the local problem Hamiltonian monotonically increases as a function of the energy level [see Fig. 3(b)], allowing to selectively drive amplitude transfers within distinct regions of the graph. This property provides an understanding of how probability amplitude propagates through the search space during continuous-time quantum algorithms.

In particular, we have identified QWs and QA as two formulations of the same underlying principle (see Fig. 7), with QA corresponding to a sequence of distinct QWs with gradually decreasing hopping rates Γ . We have shown that a QW-like approach employing a small and constant hopping rate is generally optimal for short evolutions, as it allows for localized dynamics near the solution state. However, it becomes suboptimal for long evolutions, as it fails to effectively transfer amplitude from higher energy states. Conversely, a QA-like strategy is preferred for long evolutions, as it guides amplitude throughout the entire graph, but it is suboptimal for short evolutions, since it spends insufficient time in sub-

spaces near the solution state. Based on these insights, we have argued that optimal quantum evolutions must adapt to the total evolution time T by striking a balance between the number of guided local subspaces (representing the amount of guided amplitude) and the time spent at each energy level (reflecting the amount of amplitude transferred within each local subspace).

Within the LAT framework, we have introduced the guided quantum walk (GQW) as a promising approach for solving large-scale combinatorial optimization problems in the transition region between QWs and QA. The GQW progressively drives local subspaces at gradually decreasing energy levels by utilizing a monotonically decreasing hopping rate $\Gamma(t)$. The hopping rate is controlled through a cubic Bézier curve (see Fig. 5) defined by six hyperparameters, which allows for fine-tuning the quantum evolution to each problem instance (i.e., energy spectrum of $\hat{H}_C$) and evolution time T .

We assessed the performance of the GQW in comparison to QA and QWs on exact cover (EC), traveling salesperson (TSP), and garden optimization (GO) problems ranging from 9 to 30 qubits. Across all investigated problem instances and evolution times T , the GQW outperformed both the QW and QA significantly. Specifically, at intermediate timescales, our data reveals an up to four (three) orders of magnitude better performance on 30 qubit problems compared to QA (QW), see Fig. 8. This observation is further supported by the scaling of the minimal evolution time necessary to reach a fixed solution quality as a function of the problem size. In contrast to the exponential scaling observed for QA, the GQW demonstrates a linear scaling, strongly indicating the existence of optimal quantum evolutions that solve combinatorial optimization problems in linear time T , thus surpassing the limitations of adiabatic time evolutions.

It is worth noting that the achieved linear scaling is made possible by shifting the exponential scaling to the classical optimization phase of the hyperparameters. Nonetheless, even when considering the parameter tuning in the total run time, the GQW exhibits a time-to-solution scaling that is approximately two (four) times better than for QA (QWs), see Fig. 10(f). This positions the GQW as a powerful tool for deriving optimal annealing schedules. Furthermore, the presence of the exponential scaling in the classical optimization phase, rather than in the single run times, offers the opportunity to distribute the optimization phase across multiple quantum computing devices, thereby enabling parallelization of the process. Moreover, short evolution times also suggest the possibility of discretizing the GQW into a few time steps, therefore adapting the Bézier curve parametrization into a QAOA-like scheme on gate-based quantum computers. These are features generally not feasible for QA due to the exponential scaling of T , but it could potentially allow solving large optimization problems on near-term quantum devices. While further investigation is needed to determine how these observations translate to real quantum devices in the presence of environmental noise and on problems with noncanonical energy spectra (e.g., with large degeneracies), our results strongly support the practicality of the GQW for real-world optimization problems, and we expect that our strategy is easily applicable to other types of optimization problems, beyond EC, TSP, and GO instances.




FIG. 11. Performance comparison of (a) the GQW, (b) the QW, and (c) QA as a function of the total evolution time T on GO problems with problem sizes $N \in \{12, 15, 18, 21, 24, 30\}$ (see legend). The performance is assessed based on the solution quality S_q defined in (4). Each problem size N is evaluated using a set of 10 randomly generated instances, and the presented data corresponds to the geometric mean solution qualities (solid curves) along with the geometric standard deviations (shaded areas). Note that the GQW and the QW undergo a prior classical optimization phase. The dotted gray lines indicate horizontal cuts ($S_q \in \{1\%, 10\%, 90\%\}$), which are detailed in Fig. 13. The inset in panel (a) provides a zoomed-in view of the region corresponding to large evolution times T . Additionally, the dashed lines in panel (b) indicate the stationary solution quality that the damped oscillation of $S_q(T)$ approaches for the QW at large T .

ACKNOWLEDGMENTS

The authors thank Viv Kendon, Madita Willsch, Carlos Gonzalez, Berat Yenilen, and Fengping Jin for stimulating discussions. D.W. acknowledges support from the project Jülich UNified Infrastructure for Quantum computing (JUNIQ) that has received funding from the German Federal Ministry of Education and Research (BMBF) and the Ministry of Culture and Science of the State of North Rhine-Westphalia. The authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. for funding this project by providing computing time on the GCS Supercomputer JUWELS [40] at Jülich Supercomputing Centre (JSC).

APPENDIX A: PROBLEM COST FUNCTIONS

In this Appendix, we provide an overview of the cost functions $C(z)$ for the EC, TSP, and GO problems under investigation.

1. Exact cover problem

The EC problem [2,33] involves a set $U = \{x_0, x_1, \dots, x_{P-1}\}$ with P distinct elements and N subsets $V_i \subseteq U$, such that $U = \bigcup_i V_i$. An EC is a subset L of the set of sets $\{V_i\}$, such that the elements of L are disjoint sets, and their union is U . The problem can be expressed in matrix form using a Boolean problem matrix $A \in \{0, 1\}^{N \times P}$, where the matrix columns refer to the P elements x_k in U and the matrix rows correspond to the N subsets V_i . If $A_{ik} = 1$, then element x_k is included in subset V_i , otherwise not. Thus, L is the subset of matrix rows such that the entry 1 appears in each

column exactly once, leading to the cost function

$$C(z) = \sum_{k=0}^{P-1} \left[\sum_{i=0}^{N-1} A_{ik} z_i - 1 \right]^2, \quad (\text{A1})$$

where $z \in \{0, 1\}^{\otimes N}$ encodes the selection of subsets V_i . Note that the minimum of $C(z)$ is zero and represents the EC solution. By expanding the square and collecting linear, quadratic, and constant terms in z_i , we obtain the QUBO coefficients

$$Q_{ij} = P + \begin{cases} \sum_{k=0}^P 2A_{ik}A_{jk}, & i < j \\ \sum_{k=0}^P A_{ik} \cdot (A_{ik} - 2), & i = j \end{cases}. \quad (\text{A2})$$

2. Traveling salesperson problem

The TSP problem [7,35–38] involves a list of M locations and a cost matrix c_{kj} between each pair of locations (k, j) . The goal is to find the most optimal route, with the lowest total cost, that visits each location exactly once and returns to its starting location. We represent the problem using a Boolean problem matrix $A \in \{0, 1\}^{M \times M}$, where the matrix rows correspond to the M locations, and the matrix columns denote the order in which the locations are visited. If $A_{k,j} = 1$, then the location k is visited at time step j , otherwise not. The solution to the TSP problem is the lowest-cost arrangement of 0's and 1's in A , such that an 1 is contained in each row and in each column exactly once, leading to the cost




FIG. 12. Performance comparison of (a) the GQW, (b) the QW, and (c) QA as a function of the total evolution time T on TSP problems with problem sizes $N \in \{9, 16, 25\}$ (see legend). The performance is assessed based on the solution quality S_q defined in (4). Each problem size N is evaluated using a set of 10 randomly generated instances, and the presented data corresponds to the geometric mean solution qualities (solid curves) along with the geometric standard deviations (shaded areas). Note that the GQW and the QW undergo a prior classical optimization phase. The dotted gray lines indicate horizontal cuts ($S_q \in \{1\%, 10\%, 90\%\}$), which are detailed in Fig. 14. The inset in panel (a) provides a zoomed-in view of the region corresponding to large evolution times T . Additionally, the dashed lines in panel (b) indicate the stationary solution quality that the damped oscillation of $S_q(T)$ approaches for the QW at large T .

function

$$C(z) = \lambda \sum_{k=0}^{M-1} \sum_{j=0}^{M-1} \sum_{t=0}^{M-1} c_{kj} z_{kM+t} z_{jM+t+1} + \sum_{k=0}^{M-1} \left[\sum_{t=0}^{M-1} z_{kM+t} - 1 \right]^2 + \sum_{t=0}^{M-1} \left[\sum_{k=0}^{M-1} z_{kM+t} - 1 \right]^2, \quad (\text{A3})$$

with $z_i = A_{k,j}$ with $i = kM + j$ and λ denoting a free parameter to scale the cost matrix with respect to the constraints. We can determine the QUBO coefficients $Q_{ii'}$ by stepping through the terms in Eq. (A3) and summing up the respective contributions. It is important to note that we fix the starting location, reducing the number of problem variables z_i to $(M-1)^2$. Moreover, the number of optimal routes is even because each route can be traveled in one direction or the other.

3. Garden optimization problem

The GO problem was first introduced by Gonzalez Calaza *et al.* [39] and refers to the problem of arranging plants in a garden. The garden is represented by M pots randomly distributed on a grid, with a Boolean matrix $J \in \{0, 1\}^{M \times M}$ encoding the adjacency of pots. Given P distinct plant species, with c_j plants per species, the goal of the GO problem is to find the optimal arrangement of plants in the pots with respect to the relationships between different plant species. These relationships are encoded in the companions matrix

$A \in \{-1, 0, 1\}^{P \times P}$, yielding the cost function

$$C(z) = \sum_{k,k'=0}^{M-1} J_{k,k'} \left[1 + \sum_{j,j'=0}^{P-1} z_{kP+j} A_{j,j'} z_{k'P+j'} \right] + \lambda_1 \sum_{k=0}^{M-1} \left[1 - \sum_{j=0}^{P-1} z_{kP+j} \right]^2 + \lambda_2 \sum_{j=0}^{P-1} \left[c_j - \sum_{k=0}^{M-1} z_{kP+j} \right]^2, \quad (\text{A4})$$

where λ_1 and λ_2 are free parameters. The $M \times P$ Boolean problem variables z indicate whether a plant of species j is placed in pot k (i.e., $z_{kP+j} = 1$) or not. By stepping through the terms in Eq. (A4) and summing up the respective contributions, we can determine the QUBO coefficients $Q_{ii'}$ in the same way as for the other problems.

APPENDIX B: LAT THEORY ON GENERAL STATES

In this Appendix, we explore the generalized scenario of the LAT theory, where the computational basis states may not be in an equal superposition state. For an in-depth discussion, we refer the reader to [55].

Considering the GQW Hamiltonian $\hat{H}_{\text{GQW}} = \Gamma(t) \cdot \hat{H}_D + \hat{H}_C$, with the time-dependent hopping rate $\Gamma(t)$, the driving Hamiltonian $\hat{H}_D$ [see Eq. (7)] and the cost Hamiltonian $\hat{H}_C$ [see Eq. (1)], the time evolution operator $\hat{U}$

reads

$$\hat{U}(T) = \exp \left[-i \int_0^T \hat{H}_{\text{GQW}}(t) dt \right] \quad (\text{B1})$$

$$\approx \prod_{k=1}^P \underbrace{\exp[-i \Gamma(k \cdot \Delta t) \Delta t \hat{H}_D]}_{\hat{U}_D(\Gamma(k \cdot \Delta t) \Delta t)} \underbrace{\exp[-i \hat{H}_C \Delta t]}_{\hat{U}_C(\Delta t)}. \quad (\text{B2})$$

Here we have used the first-order Suzuki-Trotter product formula [51,52] to decompose $\hat{U}$ into small-time steps Δt .

Applying one time-step to a quantum state $|\Psi\rangle = \sum_j \psi_j |j\rangle$ expressed in the computational basis $\{|j\rangle\}$, yields

$$\psi_j = r_j e^{-i\alpha_j} \quad (\text{B3})$$

$$\xrightarrow{\hat{U}_C(\gamma)} r_j e^{-i(\alpha_j + \Delta t E_j)} \quad (\text{B4})$$

$$\xrightarrow{\hat{U}_D(\beta)} \sum_{l=0}^N i^l \cdot \cos^{N-l}(\beta) \sin^l(\beta) \cdot \left[\sum_{\forall k: \Delta(k,j)=l} r_k e^{-i(\alpha_k + \Delta t E_k)} \right] \quad (\text{B5})$$

$$\xrightarrow{\beta \ll 1} r_j e^{-i(\alpha_j + \Delta t E_j)} \cdot \left[\cos^N(\beta) + i \cos^{N-1}(\beta) \sin(\beta) \sum_{\forall k: \Delta(k,j)=1} \frac{r_k}{r_j} e^{-i(\Delta\alpha_{k,j} + \Delta t \Delta E_{k,j})} \right], \quad (\text{B6})$$

with $\beta = \Gamma(k \Delta t) \cdot \Delta t$, $\gamma = \Delta t$, $\Delta\alpha_{k,j} = \alpha_k - \alpha_j$, $\Delta E_{k,j} = E_k - E_j = \langle k | \hat{H}_C | k \rangle - \langle j | \hat{H}_C | j \rangle$, and $\Delta(k, j)$ denoting the Hamming distance between the labels k and j . In the limit of $\Delta t \ll 1$, $\hat{H}_D$ simplifies to first-order interactions, enabling amplitude transfers solely between states of Hamming distance one and thus revealing the hypercube graph [e.g., see Fig. 3(a)]. Considering one of these two-dimensional subspaces spanned by the basis states $|j\rangle$ and $|k\rangle$ (i.e., $\Delta(k, j) = 1$), the amplitude transfer during one time-step is given by

$$\psi_j = r_j e^{-i\alpha_j} \quad (\text{B7})$$

$$\xrightarrow{\hat{U}(\Delta t)} r_j e^{-i(\alpha_j + \gamma \cdot E_j)} \left[\cos^N(\beta) + \frac{r_k}{r_j} \cos^{N-1}(\beta) \sin(\beta) \sin(\Delta\alpha_{k,j} + \gamma \cdot \Delta E_{k,j}) + i \frac{r_k}{r_j} \cos^{N-1}(\beta) \sin(\beta) \cos(\Delta\alpha_{k,j} + \gamma \cdot \Delta E_{k,j}) \right], \quad (\text{B8})$$

$$\psi_k = r_k e^{-i\alpha_k} \quad (\text{B9})$$

$$\xrightarrow{\hat{U}(\Delta t)} r_k e^{-i(\alpha_k + \gamma \cdot E_k)} \left[\cos^N(\beta) - \frac{r_j}{r_k} \cos^{N-1}(\beta) \sin(\beta) \sin(\Delta\alpha_{k,j} + \gamma \cdot \Delta E_{k,j}) + i \frac{r_j}{r_k} \cos^{N-1}(\beta) \sin(\beta) \cos(\Delta\alpha_{k,j} + \gamma \cdot \Delta E_{k,j}) \right]. \quad (\text{B10})$$

Equations (B8) and (B10) demonstrate that complex phase shifts caused by $\hat{H}_C$ result in opposing signs of the second terms of the squared brackets, with the lower (higher) energy state having a “+” (“−”). When $r_j/r_k \approx 1$ and β is small, the amplitude of each basis state after one time-step is mainly defined by $\cos^N(\beta)$, which denotes the amount of amplitude that remains at each basis state. The second term increases this amplitude for the low energy state, while decreasing it for the high energy state, resulting in a localized amplitude transfer between the two. As a result, the hypercube graph becomes directed and probability amplitude is steered towards the ground state of $\hat{H}_C$. It is important to note, however, that for a time-independent β (e.g., in a QW) this process will eventually reverse, causing amplitude to flow back into the high energy state. This occurs because the fraction r_j/r_k grows as amplitude is exchanged between the two states. Thus, the higher energy state's second term eventually overwhelms the $\cos^N(\beta)$ resulting in an increased amplitude, while the lower energy state's second term becomes suppressed causing an overall decrease in amplitude

as $\cos^N(\beta) < 1$. This leads to oscillations in the probability amplitude between the two states (see e.g., Fig. 2). To prevent this, the GQW makes use of a time-dependent monotonically decreasing hopping rate $\Gamma(t)$, which counteracts the increasing fraction r_j/r_k through the $\sin(\beta)$. By doing so, the GQW is able to prohibit the back-propagation of amplitude even in the case of large detuning by suppressing the amplitude transfer locally.

APPENDIX C: GO AND TSP RESULTS

In this Appendix, we present the simulation results for the GQW, the QW, and QA applied to GO and TSP problems. We investigate GO problems with problem sizes $N \in \{12, 15, 18, 21, 24, 30\}$ and TSP problems with problem sizes $N \in \{9, 16, 25\}$, considering a total of 10 randomly generated instances for each problem size and type. Similar to the EC results from the main text, the GQW and the QW are tuned for each problem instance and evolution time T separately by optimizing their hyperparameters to mini-




FIG. 13. Scaling behavior of the total evolution time $T_{S_q}(N)$ required to achieve a specified solution quality S_q on GO problems as a function of the problem size N for the GQW (circles) and linear QA (triangles). Colors indicate results for various target solution qualities: $S_q = 1\%$ (blue), $S_q = 10\%$ (red), and $S_q = 90\%$ (green). The solid lines depict linear fits ($a + b \cdot N$), while the dashed lines represent exponential fits ($a \cdot 2^{b \cdot N}$) applied to the data points (see legend).

mize the energy expectation value E_Ψ of the final quantum state $|\Psi\rangle$. Each set of hyperparameters is selected from a pool of $N_{\text{rep}} = 100$ runs of the Nelder-Mead classical optimizer, where for each run, the hyperparameters are initialized randomly and adjusted up to a maximum of $N_{\text{eval}} = 100$ times. The QA results are obtained using a linear annealing scheme.

Figures 11 and 12 depict the scaling of the solution quality S_q as a function of the evolution time T and the problem size N for the three algorithms on GO and TSP problems,




FIG. 14. Scaling behavior of the total evolution time $T_{S_q}(N)$ required to achieve a specified solution quality S_q on TSP problems as a function of the problem size N for the GQW (circles) and linear QA (triangles). Colors indicate results for various target solution qualities: $S_q = 1\%$ (blue), $S_q = 10\%$ (red), and $S_q = 90\%$ (green). The solid lines depict linear fits ($a + b \cdot N$), while the dashed lines represent exponential fits ($a \cdot 2^{b \cdot N}$) applied to the data points (see legend).




FIG. 15. Scaling analysis of the optimal (smallest) time-to-solution $TT_{S_{99.99\%}}$ [see Eq. (19)] for the GQW (blue), QA (red), and the QW (green) on GO problems as a function of the problem size N . The solid lines correspond to exponential fits ($a \cdot 2^{b \cdot N}$) applied to the data points (see legend).

respectively. The data exhibit similar characteristics to the EC problems presented in Fig. 8 and are qualitatively consistent with the discussion in Sec. IV A. Notably, the GQW achieves higher solution qualities for some large GO problems (e.g., $N = 30$) compared to some smaller problem sizes (e.g., $N = 18$) for long evolutions [see inset in Fig. 11(a)]. This indicates that the hardness of the combinatorial optimization problems varies throughout the problem sizes, such that some large problem instances are easier to solve than their small counterparts once the GQW considers the entire graph. Consequently, we expect similar characteristics in the scaling of S_q to appear for QA for long evolutions.

In Figures 13 and 14, we present the scaling of the minimum evolution time $T_{S_q}(N)$ required to achieve a specified solution quality S_q as a function of the problem size N . Similar to the results obtained for the EC problems, $T_{S_q}(N)$ exhibits linear scaling for the GQW, in contrast to the exponential




FIG. 16. Scaling analysis of the optimal (smallest) time-to-solution $TT_{S_{99.99\%}}$ [see Eq. (19)] for the GQW (blue), QA (red), and the QW (green) on TSP problems as a function of the problem size N . The solid lines correspond to exponential fits ($a \cdot 2^{b \cdot N}$) applied to the data points (see legend).

scaling observed for QA. This further supports the existence of optimal hopping rate schedules that can solve combinatorial optimization problems in linear time. Note, however, that for the GO problems, $T_{0.9}(N)$ neither corresponds to an exponential nor a linear scaling, due to the differences in the hardness of the optimization problems.

Figures 15 and 16 show the scaling of the time-to-solution [see Eq. (19)] as a function of the problem size N . For both problem types, the GQW achieves a superior scaling compared to QA and the QW by factors of ≈ 2 (QA on GO problems), ≈ 4 (QW on GO problems), and ≈ 2 (QA and QW on TSP problems).

-
- [1] M. Marzec, Portfolio optimization: Applications in quantum computing, in *Handbook of High-Frequency Trading and Modeling in Finance*, edited by I. Florescu, M. C. Mariani, H. E. Stanley, and F. G. Viens, 73 (John Wiley & Sons, Ltd, Hoboken, New Jersey, 2016).
- [2] D. Willsch, M. Willsch, C. D. G. Calaza, F. Jin, H. De Raedt, M. Svensson, and K. Michielsen, Benchmarking advantage and d-wave 2000Q quantum annealers with exact cover problems, *Quant. Info. Proc.* **21**, 141 (2022).
- [3] M. H. Amin, E. Andriyash, J. Rolfe, B. Kulchytsky, and R. Melko, Quantum Boltzmann machine, *Phys. Rev. X* **8**, 021050 (2018).
- [4] M. Benedetti, J. Realpe-Gómez, R. Biswas, and A. Perdomo-Ortiz, Estimation of effective temperatures in quantum annealers for sampling applications: A case study with possible applications in deep learning, *Phys. Rev. A* **94**, 022308 (2016).
- [5] A. M. Childs and J. Goldstone, Spatial search by quantum walk, *Phys. Rev. A* **70**, 022314 (2004).
- [6] A. Perdomo-Ortiz, N. Dickson, M. Drew-Brook, G. Rose, and A. Aspuru-Guzik, Finding low-energy conformations of lattice protein models by quantum annealing, *Sci. Rep.* **2**, 571 (2012).
- [7] R. J. Boucherie, A. Braaksma, and H. Tijms, *Operations Research* (World Scientific, Singapore, 2021).
- [8] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [9] Y. Aharonov, L. Davidovich, and N. Zagury, Quantum random walks, *Phys. Rev. A* **48**, 1687 (1993).
- [10] A. M. Childs, E. Deotto, E. Farhi, J. Goldstone, S. Gutmann, and A. J. Landahl, Quantum search by measurement, *Phys. Rev. A* **66**, 032314 (2002).
- [11] A. Callison, M. Festenstein, J. Chen, L. Nita, V. Kendon, and N. Chancellor, Energetic perspective on rapid quenches in quantum annealing, *PRX Quantum* **2**, 010338 (2021).
- [12] R. J. Banks, E. Haque, F. Nazef, F. Fethallah, F. Ruqaya, H. Ahsan, H. Vora, H. Tahir, I. Ahmad, I. Hewins *et al.*, Continuous-time quantum walks for MAX-CUT are hot, *Quantum* **8**, 1254 (2024).
- [13] B. Apolloni, C. Carvalho, and D. de Falco, Quantum stochastic optimization, *Stoch. Process. Their Appl.* **33**, 233 (1989).
- [14] A. Finnila, M. Gomez, C. Sebenik, C. Stenson, and J. Doll, Quantum annealing: A new method for minimizing multidimensional functions, *Chem. Phys. Lett.* **219**, 343 (1994).
- [15] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse Ising model, *Phys. Rev. E* **58**, 5355 (1998).
- [16] J. Brooke, D. Bitko, T. F. Rosenbaum, and G. Aeppli, Quantum annealing of a disordered magnet, *Science* **284**, 779 (1999).
- [17] M. W. Johnson, M. H. S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk *et al.*, Quantum annealing with manufactured spins, *Nature (London)* **473**, 194 (2011).
- [18] R. Harris, M. W. Johnson, T. Lanting, A. J. Berkley, J. Johansson, P. Bunyk, E. Tolkacheva, E. Ladizinsky, N. Ladizinsky, T. Oh, F. Cioata, I. Perminov, P. Spear, C. Enderud, C. Rich, S. Uchaikin, M. C. Thom, E. M. Chapple, J. Wang, B. Wilson *et al.*, Experimental investigation of an eight-qubit unit cell in a superconducting optimization processor, *Phys. Rev. B* **82**, 024511 (2010).
- [19] P. I. Bunyk, E. M. Hoskinson, M. W. Johnson, E. Tolkacheva, F. Altomare, A. J. Berkley, R. Harris, J. P. Hilton, T. Lanting, A. J. Przybysz, and J. Whittaker, Architectural considerations in the design of a superconducting quantum annealing processor, *IEEE Trans. Appl. Supercond.* **24**, 1 (2014).
- [20] T. Albash and D. A. Lidar, Adiabatic quantum computation, *Rev. Mod. Phys.* **90**, 015002 (2018).
- [21] T. Albash and D. A. Lidar, Demonstration of a scaling advantage for a quantum annealer over simulated annealing, *Phys. Rev. X* **8**, 031016 (2018).
- [22] A. D. King, J. Raymond, T. Lanting, R. Harris, A. Zucca, F. Altomare, A. J. Berkley, K. Boothby, S. Ejtemaee, C. Enderud *et al.*, Quantum critical dynamics in a 5,000-qubit programmable spin glass, *Nature (London)* **617**, 61 (2023).
- [23] A. Callison, N. Chancellor, F. Mintert, and V. Kendon, Finding spin glass ground states using quantum walks, *New J. Phys.* **21**, 123022 (2019).
- [24] T. Neuhaus, Monte Carlo search for very hard KSAT realizations for use in quantum annealing, [arXiv:1412.5361](https://arxiv.org/abs/1412.5361).
- [25] T. Neuhaus, Quantum searches in a hard 2SAT ensemble, [arXiv:1412.5460](https://arxiv.org/abs/1412.5460).
- [26] V. Mehta, F. Jin, H. De Raedt, and K. Michielsen, Quantum annealing with trigger Hamiltonians: Application to 2-satisfiability and nonstoquastic problems, *Phys. Rev. A* **104**, 032421 (2021).
- [27] V. Mehta, F. Jin, H. De Raedt, and K. Michielsen, Quantum annealing for hard 2-satisfiability problems: Distribution and scaling of minimum energy gap and booth probability, *Phys. Rev. A* **105**, 062406 (2022).
- [28] V. Mehta, F. Jin, K. Michielsen, and H. De Raedt, On the hardness of quadratic unconstrained binary optimization problems, *Front. Phys.* **10**, 956882 (2022).
- [29] J. G. Morley, N. Chancellor, S. Bose, and V. Kendon, Quantum search with hybrid adiabatic-quantum-walk algorithms and realistic noise, *Phys. Rev. A* **99**, 022339 (2019).
- [30] L. Zeng, J. Zhang, and M. Sarovar, Schedule path optimization for adiabatic quantum computing and optimization, *J. Phys. A: Math. Theor.* **49**, 165305 (2016).
- [31] D. Herr, E. Brown, B. Heim, M. Könz, G. Mazzola, and M. Troyer, Optimizing schedules for quantum annealing, [arXiv:1705.00420](https://arxiv.org/abs/1705.00420).

- [32] Y.-Q. Chen and Y. Chen and C.-K. Lee and S. Zhang and C.-Y. Hsieh, Optimizing quantum annealing schedules with Monte Carlo tree search enhanced with neural networks, *Nat. Mach. Intell.* **4**, 269 (2022).
- [33] R. M. Karp, Reducibility among combinatorial problems, in *Complexity of Computer Computations: Proceedings of a symposium on the Complexity of Computer Computations, held March 20–22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, and sponsored by the Office of Naval Research, Mathematics Program, IBM World Trade Corporation, and the IBM Research Mathematical Sciences Department*, edited by R. E. Miller, J. W. Thatcher, and J. D. Bohlinger (Springer US, Boston, 1972) pp. 85–103.
- [34] V. Choi, Adiabatic quantum algorithms for the NP-complete maximum-weight independent set, exact cover and 3SAT problems, *arXiv:1004.2226*.
- [35] G. Dantzig, R. Fulkerson, and S. Johnson, Solution of a large-scale traveling-salesman problem, *J. Oper. Res. Soc. Am.* **2**, 393 (1954).
- [36] R. Bellman, Dynamic programming treatment of the travelling salesman problem, *J. ACM* **9**, 61 (1962).
- [37] M. Held and R. M. Karp, A dynamic programming approach to sequencing problems, *J. Soc. Ind. Appl. Math.* **10**, 196 (1962).
- [38] A. Lucas, Ising formulations of many NP problems, *Front. Phys.* **2**, 5 (2014).
- [39] C. D. Gonzalez Calaza, D. Willsch, and K. Michielsen, Garden optimization problems for benchmarking quantum annealers, *Quant. Info. Proc.* **20**, 305 (2021).
- [40] Jülich Supercomputing Centre, JUWELS cluster and booster: Exascale pathfinder with modular supercomputing architecture at juelich supercomputing centre, *J. Large-Scale Res. Facil.* **7**, A183 (2021).
- [41] L. K. Grover, A fast quantum mechanical algorithm for database search, in *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing, STOC '96* (Association for Computing Machinery, New York, 1996) p. 212–219.
- [42] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, *arXiv:1411.4028*.
- [43] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, *New J. Phys.* **18**, 023023 (2016).
- [44] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth, and J. Tennyson, The variational quantum eigensolver: A review of methods and best practices, *Phys. Rep.* **986**, 1 (2022).
- [45] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, Quantum computation by adiabatic evolution, *arXiv:quant-ph/0001106*.
- [46] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices, *Phys. Rev. X* **10**, 021067 (2020).
- [47] S. H. Sack and M. Serbyn, Quantum annealing initialization of the quantum approximate optimization algorithm, *Quantum* **5**, 491 (2021).
- [48] J. A. Nelder and R. Mead, A simplex method for function minimization, *Comput. J.* **7**, 308 (1965).
- [49] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes 3rd Edition: The Art of Scientific Computing* (Cambridge University Press, New York, 2007).
- [50] J. Roland and N. J. Cerf, Quantum search by local adiabatic evolution, *Phys. Rev. A* **65**, 042308 (2002).
- [51] M. Suzuki, General decomposition theory of ordered exponentials, *Proc. Jpn. Acad., Ser. B* **69**, 161 (1993).
- [52] H. De Raedt and K. Michielsen, Computational methods for simulating quantum computers, in *Quantum and Molecular Computing, Quantum Simulations, Handbook of Theoretical and Computational Nanotechnology*, Vol. 3, edited by M. Rieth and W. Schommers (American Scientific, Los Angeles, 2006) pp. 1–48.
- [53] A. P. Young, S. Knysh, and V. N. Smelyanskiy, Size dependence of the minimum excitation gap in the quantum adiabatic algorithm, *Phys. Rev. Lett.* **101**, 170503 (2008).
- [54] A. P. Young, S. Knysh, and V. N. Smelyanskiy, First-order phase transition in the quantum adiabatic algorithm, *Phys. Rev. Lett.* **104**, 020502 (2010).
- [55] S. Schulz, GPU-accelerated simulation of guided quantum walks, Master's thesis, RWTH Aachen, 2022.