

PAPER • OPEN ACCESS

Diabatic quantum annealing for the frustrated ring model

To cite this article: Jeremy Côté *et al* 2023 *Quantum Sci. Technol.* **8** 045033

View the [article online](#) for updates and enhancements.

You may also like

- [Investigating the potential for a limited quantum speedup on protein lattice problems](#)

Carlos Outeiral, Garrett M Morris, Jiye Shi et al.

- [Quantum annealing for industry applications: introduction and review](#)

Sheir Yarkoni, Elena Raponi, Thomas Bäck et al.

- [Quantum annealing with capacitive-shunted flux qubits](#)

Yuichiro Matsuzaki, Hideaki Hakoshima, Yuya Seki et al.



Easy-to-use and Helium-3 free cryogenics solutions

[LEARN MORE](#)

Quantum Science and Technology



CrossMark

OPEN ACCESS**RECEIVED**
9 February 2023**REVISED**
31 July 2023**ACCEPTED FOR PUBLICATION**
20 September 2023**PUBLISHED**
6 October 2023

Original Content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.

**PAPER**

Diabatic quantum annealing for the frustrated ring model

Jeremy Côté^{1,2,*} Frédéric Sauvage¹ Martín Larocca^{1,3} Matías Jonsson^{1,4} Lukasz Cincio¹ and Tameem Albash^{5,6}

¹ Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545, United States of America

² Institut Quantique & Département de Physique, Université de Sherbrooke, Sherbrooke, Québec J1K 2R1, Canada

³ Center for Nonlinear Studies, Los Alamos National Laboratory, Los Alamos, NM 87545, United States of America

⁴ Department of Physics, Carnegie Mellon University, Pittsburgh, PA 15217, United States of America

⁵ Department of Electrical and Computer Engineering, University of New Mexico, Albuquerque, NM 87131, United States of America

⁶ Department of Physics and Astronomy and Center for Quantum Information and Control, University of New Mexico, Albuquerque, NM 87131, United States of America

* Author to whom any correspondence should be addressed.

E-mail: jeremy.cote.science@gmail.com

Keywords: diabatic, quantum, annealing, optimization

Abstract

Quantum annealing (QA) is a continuous-time heuristic quantum algorithm for solving or approximately solving classical optimization problems. The algorithm uses a schedule to interpolate between a driver Hamiltonian with an easy-to-prepare ground state and a problem Hamiltonian whose ground state encodes solutions to an optimization problem. The standard implementation relies on the evolution being adiabatic: keeping the system in the instantaneous ground state with high probability and requiring a time scale inversely related to the minimum energy gap between the instantaneous ground and excited states. However, adiabatic evolution can lead to evolution times that scale exponentially with the system size, even for computationally simple problems. Here, we study whether non-adiabatic evolutions with optimized annealing schedules can bypass this exponential slowdown for one such class of problems called the frustrated ring model. For sufficiently optimized annealing schedules and system sizes of up to 39 qubits, we provide numerical evidence that we can avoid the exponential slowdown. Our work highlights the potential of highly-controllable QA to circumvent bottlenecks associated with the standard implementation of QA.

1. Introduction

The ubiquity of optimization problems across disciplines and the incredible computational resources they consume continues to motivate researchers to explore new and more efficient approaches to tackle them. Though quantum computing has so far offered limited provable computational advantages for combinatorial optimization problems [1–14], we expect that the library of quantum approaches for tackling different classes of optimization problems will expand as quantum information processors mature and become more readily available.

Quantum annealing (QA) is a generic quantum approach to tackling combinatorial optimization problems. Starting from an efficiently prepared ground state of an initial Hamiltonian \hat{H}_d , the system is evolved according to an interpolating Hamiltonian between \hat{H}_d and a problem Hamiltonian \hat{H}_p , which ground state encodes the solution to the optimization problem. At the end of the evolution, the algorithm is deemed successful if the quantum state has high overlap with the ground state of \hat{H}_p .

A sufficient but not necessary condition for QA to succeed is for the evolution to satisfy the adiabatic condition, which requires the total evolution time T , also called the annealing time, to scale as an inverse power of the minimum gap encountered along the interpolations [15–17]. When it does, QA is sometimes referred to as quantum adiabatic optimization (QAO) [18, 19], and it belongs to the adiabatic paradigm of quantum computing [20, 21]. The adiabatic theorem of quantum mechanics [15–17] provides a guarantee

that if the evolution is sufficiently slow, then the state at the end of the evolution will have high overlap with the ground state of \hat{H}_p . The efficiency of the algorithm is thus given by the scaling of the minimum gap encountered along the interpolating schedule with the system size.

However, enforcing adiabaticity for all realizations of QA can result in exponentially slow evolutions even for simple problems [22, 23]. Furthermore, the only provable exponential speedups currently known for QA involves an evolution that is *not* adiabatic [8]. This suggests that it might be meaningful to consider heuristic continuous-time quantum optimization algorithms that are not necessarily adiabatic, and for which we can design the interpolation to take advantage of non-adiabatic transitions. This approach is sometimes called ‘diabatic quantum annealing’ (DQA) [24]. The added flexibility of DQA comes at the cost of identifying a suitable interpolation schedule, which is a non-trivial optimization problem in itself. This is in contrast to other approaches to shortcuts to adiabaticity [25], such as counterdiabatic driving [26] whereby new Hamiltonian terms are added to the Hamiltonian. If our goal is to search for interpolating schedules that improve on the evolution times of the adiabatic approach, we must ensure that the computational resources to identify these schedules also scale favorably with the system size.

In this work, we consider the performance of DQA on a simple class of optimization problems, known to be exponentially hard for QAO [27]. The problem is a one-dimensional frustrated ring model whose ground state can be found analytically, but poses a challenge for standard implementations of QAO because it exhibits a perturbative crossing [28], leading to exponentially slow evolution times. We will show that we can identify interpolating schedules that solve the optimization problem more efficiently than QAO.

The paper is organized as follows. In section 2, we review the frustrated ring model and why it exhibits an exponentially closing gap with the system size. In section 3, we present our algorithm for optimizing the annealing schedules and annealing times. In section 4, we present our numerical results for our optimized schedules and the performance of the DQA algorithm. We conclude and highlight some future directions in section 5.

2. Model

For our problem Hamiltonian, we focus on the frustrated ring model of [27] defined on N Ising spins arranged in a ring. The problem Hamiltonian is given by:

$$\hat{H}_p = - \sum_{j=1}^N J_j \hat{Z}_j \hat{Z}_{j+1}, \quad (1)$$

where J_j denotes the coupling between spins j and $j+1$, \hat{Z}_j is the Pauli operator \hat{Z} acting on spin j , and $\hat{Z}_{N+1} \equiv \hat{Z}_1$. We denote the $+1$ eigenstate of \hat{Z} by $|0\rangle$, and the -1 eigenstate by $|1\rangle$. For simplicity, we restrict ourselves to the case where N is odd. There are three types of couplings, and using the indexing convention in figure 1, they are:

$$J_j = \begin{cases} -J_R & \text{if } j = N, \\ J_L & \text{if } j = \frac{1}{2}(N \pm 1), \\ J & \text{otherwise.} \end{cases} \quad (2)$$

For the rest of the paper, we take the couplings to satisfy $0 < J_R < J_L < J = 1$. With these choices, the couplings associated with J_L and J are ferromagnetic, and the couplings associated with J_R are antiferromagnetic.

The doubly degenerate ground states of this Hamiltonian are given by the states $|0\rangle^{\otimes N}$ and $|1\rangle^{\otimes N}$, which satisfy all the ferromagnetic couplings but violate the one antiferromagnetic coupling of the Hamiltonian. The ground state energy is thus given by:

$$E_0 = - \sum_{j=1}^N J_j = -(N-3)J + J_R - 2J_L. \quad (3)$$

On the other hand, the first excited states satisfy the antiferromagnetic coupling but violate one of the two J_L ferromagnetic couplings, giving rise to a four-fold degenerate first excited eigenspace with energy:

$$E_1 = -(N-3)J - J_R = E_0 + 2(J_L - J_R). \quad (4)$$

Thus, the difference $J_L - J_R$ determines the ground state energy gap of the problem Hamiltonian.

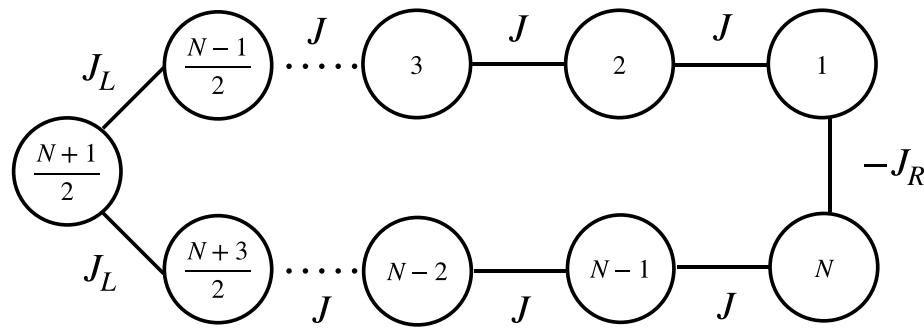


Figure 1. A visualization of the frustrated ring model with a total of N (odd) spins with our indexing convention for the sites.

We now consider a continuous-time algorithm with a time-dependent Hamiltonian $\hat{H}(t)$ that interpolates between a uniform transverse field Hamiltonian $\hat{H}_d = -\sum_{j=1}^N \hat{X}_j$, where \hat{X}_j is the Pauli operator \hat{X} acting on site j , and the problem Hamiltonian. That is,

$$\hat{H}(t) = (1 - A(t)) \hat{H}_d + A(t) \hat{H}_p = \hat{H}_d + A(t) (\hat{H}_p - \hat{H}_d) , \quad (5)$$

where $A(t)$ is the annealing schedule of our algorithm. For the remainder of the paper, we require $A(0) = 0$ and $A(T) = 1$, where T is the total annealing time for the protocol.

An important feature of this combination of problem Hamiltonian and mixer Hamiltonian is that it gives rise to a ‘perturbative crossing’ [28] in the energy spectrum. To see this, we first express the ground states and first excited states as symmetric and anti-symmetric combinations (for simplicity we write this for the case of $N = 7$):

$$|\Phi_0^{(\pm)}\rangle = \frac{1}{\sqrt{2}} (|0000000\rangle \pm |1111111\rangle) , \quad (6a)$$

$$|\Phi_1^{(\pm)}\rangle = \frac{1}{\sqrt{2}} (|0001111\rangle \pm |1110000\rangle) , \quad (6b)$$

$$|\Phi_{1'}^{(\pm)}\rangle = \frac{1}{\sqrt{2}} (|0000111\rangle \pm |1111000\rangle) . \quad (6c)$$

Because the Hamiltonian $\hat{H}(t)$ commutes with the operator $\hat{X}^{\otimes N}$, the energy eigenstates can be labeled by the two eigenvalues ± 1 of this operator. We only need to concern ourselves with the $+1$ eigenvalue states (the reason for this will become apparent when we discuss the initial state of our continuous-time algorithm), which in turn means we only need to consider the symmetric combinations of states in (6). The two first excited states $\Phi_1^{(+)}$ and $\Phi_{1'}^{(+)}$ are coupled by the transverse field, so the degeneracy of the first excited state of the problem Hamiltonian is broken when we turn on the parameter $1 - A$. At first order in perturbation theory, the unique instantaneous first excited state has energy

$$E_1(A) = E_1 A - (1 - A) . \quad (7)$$

The ground state energy is corrected to $E_0(A) = A E_0$ at the same order. As A decreases from 1, first order perturbation theory predicts an energy level crossing at:

$$A_* = \frac{1}{1 + E_1 - E_0} , \quad (8)$$

which will be corrected to an exponentially small (in N) avoided level crossing in the full quantum theory [27], with (8) being approximately the position of the minimum gap. For example, with $J_R = 0.45, J_L = 0.5$ we find $A_* \approx 0.9091$.

3. Optimizing the annealing schedules

Starting from the state $|\psi(0)\rangle = |+\rangle^{\otimes N}$, which is the ground state of $\hat{H}(0)$, our objective is to prepare a state $|\psi(T)\rangle = \text{Texp}(\int_0^T \hat{H}(t) dt) |\psi(0)\rangle$, with Texp being the time ordered exponential, that has a large overlap with the (relevant) ground state of \hat{H}_p . We provide details of how we simulate the dynamics in appendix A. In order to avoid unfairly steering the optimization of the annealing schedule towards the known solution, we

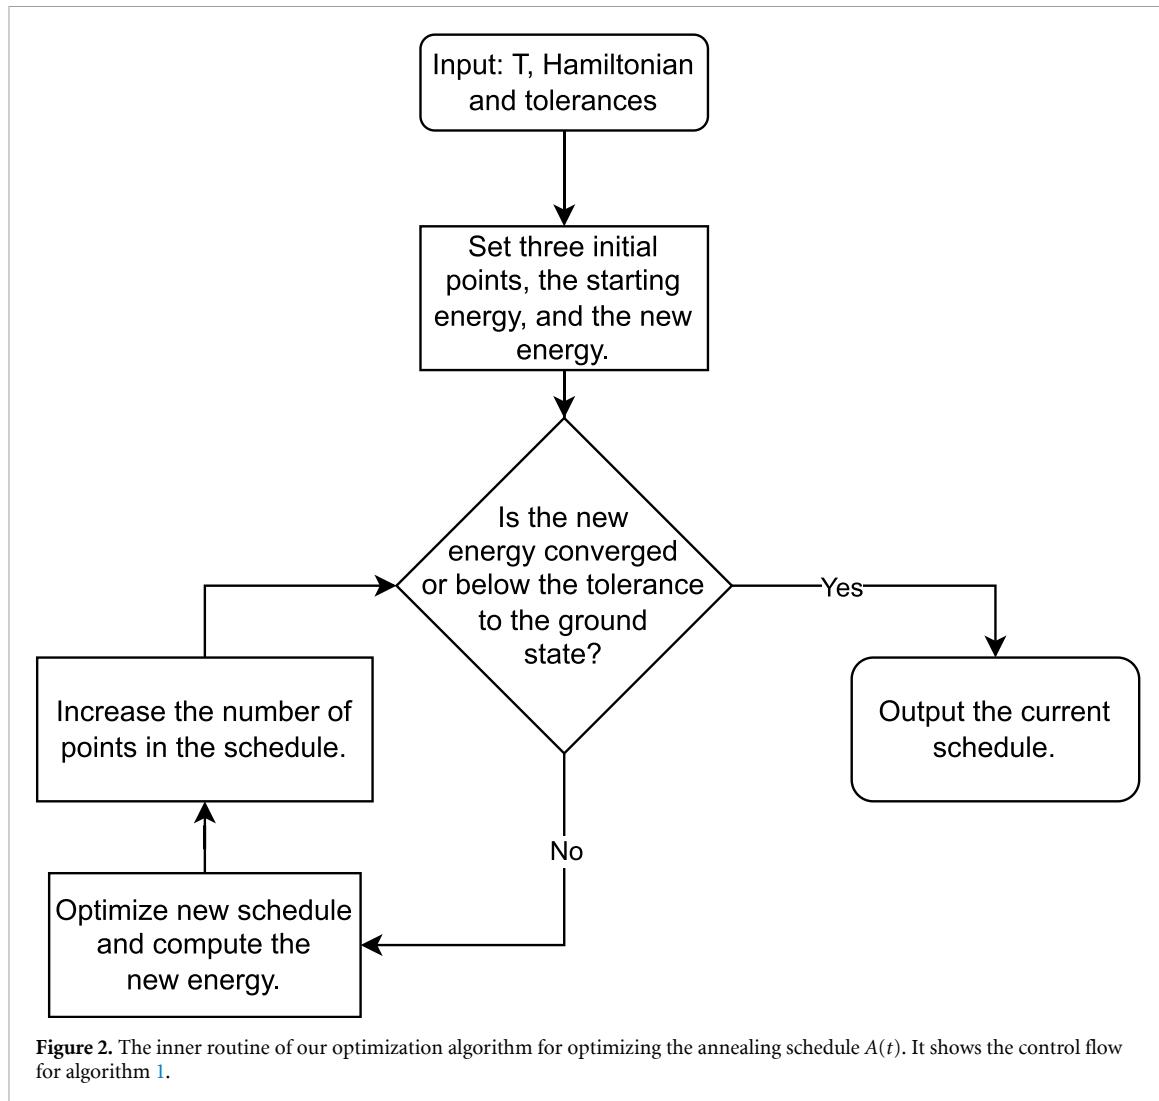


Figure 2. The inner routine of our optimization algorithm for optimizing the annealing schedule $A(t)$. It shows the control flow for algorithm 1.

choose the cost function to minimize to be the energy $E(T) = \langle \psi(T) | \hat{H}_p | \psi(T) \rangle$. We note that this choice is general and can be applied to any problem Hamiltonian. In order to determine the success of our optimization, we choose a threshold value Δ_E and consider any state $|\psi(T)\rangle$ with $E(T) \leq E_0 + \Delta_E$ to be a success. This choice does use our knowledge of the true ground state, and we use it to provide no ambiguity in the scaling performance of our optimization method. In more realistic situations, where the ground state energy of the problem Hamiltonian is not known, one can use more sophisticated stopping criteria such as the optimal stopping discussed in [29], where a further cost associated with every call to the algorithm is introduced.

Our goal is to identify the lowest time T and corresponding schedule which, given a target energy threshold value, ensures the success of the optimization. Other objectives besides the lowest time could be used, but we do not pursue them in this work. For example, if our objective is to minimize the total time to find a state with or below a target energy $E_0 + \Delta_E$, a more suitable metric might be the time-to-solution [30], which balances spending more time per independent run of the algorithm for a higher success probability against performing multiple fast runs of the algorithm to boost a lower success probability.

As a baseline, we consider the linear annealing schedule $A(t) = t/T$, with T large enough to satisfy the adiabatic condition. As we discussed in section 2, Hamiltonian $\hat{H}(t)$ exhibits along its interpolation a minimum energy gap between its first excited state and ground state that closes exponentially with the system size. This implies that an adiabatic protocol would require an evolution time that scales exponentially with the system size, which means that the protocol is exponentially slow in finding the ground state (with high probability) of the one-dimensional frustrated ring Ising problem, despite its computational simplicity.

Our objective is to optimize the annealing time and schedule of DQA to avoid this exponential slowdown. To do so, we divide our optimization procedure into an inner and outer routine. Within the inner routine, which we visualize in figure 2, the annealing time T is fixed and we optimize the schedule $A(t)$ to minimize the energy of the prepared state. We discuss the details of this in section 3.1.

Algorithm 1. FindSchedule.

```

Input:  $J_R, J_L, J, E_0, N, T, \delta t, \delta E, c, k_0, \text{trials}$ 
Output:  $A(t)$ 
 $k \leftarrow k_0$  // Starting number of points;
 $E_{\text{best}} \leftarrow \infty$  // We set at  $10^5$ ;
for  $i < \text{trials}$  do
   $a \leftarrow \text{InitialPoints}(k)$  // Get  $\{a_j\}$ ;
   $a(T) \leftarrow \text{Optimize}(a, T)$ ;
   $E(T) \leftarrow \text{ComputeEnergy}(a, T)$ ;
  if  $E(T) < E_{\text{best}}$  then
     $E_{\text{best}} \leftarrow E(T)$ ;
     $a_{\text{best}}(T) \leftarrow a(T)$ ;
  end if
   $i \leftarrow i + 1$ ;
end for
 $k \leftarrow 2k + 1$ ;
 $a(T) \leftarrow a_{\text{best}}(T)$ ;
 $E_{\text{previous}} \leftarrow E_{\text{best}}$ ;
repeat
   $a \leftarrow \text{EmbedSchedule}(a(T), k)$  // Get  $\{a_j\}$ ;
   $a(T) \leftarrow \text{Optimize}(a, T)$ ;
   $E(T) \leftarrow \text{ComputeEnergy}(a, T)$ ;
  if  $E(T) < E_{\text{previous}}$  then
     $E_{\text{previous}} \leftarrow E(T)$ ;
  end if
   $k \leftarrow 2k + 1$ ;
until  $E(T) - E_0 < 2c(J_L - J_R)$  or  $|E(T) - E_{\text{previous}}| < \delta E$ ;
 $A(t) \leftarrow \text{Schedule}(a, T)$  // Schedule from  $\{a_j\}$ ;
return  $A(t)$ ;

```

In addition, the outer routine (figure 4) aims at minimizing the annealing time T required to achieve a target threshold energy. We do this by iteratively decreasing (or increasing) T depending on the success (or failure) of the inner routine. The algorithm repeats these steps until we identify a sufficiently narrow range of the values of T for which the algorithm has found a successful schedule (section 3.2). We now describe the inner and outer routines in further detail.

3.1. Optimizing the schedule

The inner routine of our algorithm involves optimizing the annealing schedule $A(t)$ such that we minimize $E(T) - E_0$ (algorithm 1), for a fixed time T . In particular, we want to ensure our output state $|\psi(T)\rangle$ has enough overlap with the true ground state $|\Phi_0^+\rangle$ so that we could in principle implement this schedule on a quantum processor and have a finite probability of measuring the ground states.

We parameterize the schedule $A(t)$ in terms of a finite number of continuous parameters corresponding to weights over a predefined set of functions. For instance, these functions could be a basis of trigonometric functions [31], a basis of polynomials [32, 33], ‘bang-bang’ or on-off pulses (standard in the quantum approximate optimization algorithm (QAOA) [34, 35]), or others [36, 37]. In our case, we opt for a linear piecewise decomposition [38], which we iteratively refine and now describe.

To parameterize $A(t)$, we begin by choosing k_0 random values a_1, a_2, \dots, a_{k_0} in the interval $[0, 1]$ and split our annealing interval $[0, T]$ into $k_0 + 1$ equal intervals separated at $t_j = jT/(k_0 + 1)$ for $j \in [1, k_0]$. We then set $A(t_j) = a_j$. This defines the initial annealing schedule $A(t)$, with $k_0 + 1$ linear functions connecting the k_0 points (as well as the start and end points, $A(0) = 0$ and $A(T) = 1$).

After initializing a random schedule, we optimize the parameters a_j such that we minimize the energy of the output state $|\psi(T)\rangle$. We perform the optimization using a gradient-free method from SciPy, which we detail further in appendix B. The result is a locally optimal schedule $A(t)$ using a piece-wise linear function with $k_0 + 1$ segments. In our experiments, we use ten different random instances of the initial tuple $(a_1, a_2, \dots, a_{k_0})$ and select the best set of parameters (in terms of the energy achieved) when proceeding to the following step. We also experiment with $k_0 \in \{3, 5, 7, 9\}$ so that the averaged results we present in section 4 do not depend on a specific number of initial parameters k_0 .

Next, we refine the schedule by adding new points (i.e. additional free parameters) bisecting the line segments connecting our existing points. For example, suppose we start with $k_0 = 3$. Then, we will initially

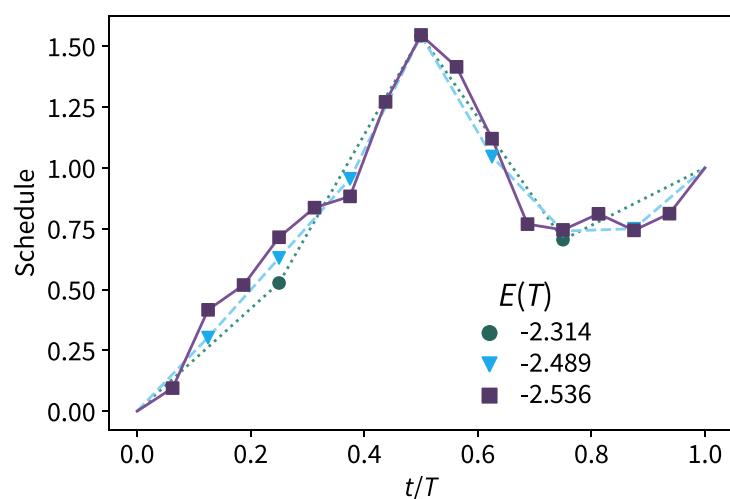


Figure 3. An example annealing schedule $A(t)$ for $N = 5$ during three stages of the optimization process, for $T = 12.5$. The green dotted line with circular markers is the initial guess with $k_0 = 3$ points, the blue dashed line with triangular markers is the first optimized schedule with 7 points, and the purple solid line with square markers is the final optimized schedule with 15 points. The ground state energy is $E_0 = -2.55$, and only the purple schedule reaches it within a tolerance of $c = 0.5$. The tolerance c is defined in (9).

have the points $(0, 0)$, (t_1, a_1) , (t_2, a_2) , (t_3, a_3) and $(T, 1)$. In this case, we would go from three free parameters to seven (including the original three). In general, when we split the interval $[0, T]$ with k points (not including the fixed endpoints), the points are located at $t_j = jT/(k+1)$, with $j \in \{1, 2, \dots, k\}$. We set the new values of a_j so that the points (t_j, a_j) bisect the line segments connecting the original points. This ensures that the introduction of new points does not initially alter the schedule. Now, the optimization performed on the refined schedule should either maintain or improve the quality of the result (measured in energy $E(T)$) as the interpolated schedule is initially identical, and further optimization is expected to improve the energy. We note that running the optimization algorithm on a schedule with more points does not guarantee a lower energy will be found.

For each new set of values $\{a_j\}$ defining $A(t)$, we run our optimization routine and compare the lowest energy found to the previous lowest energy value. If the lowest energy found does not decrease by more than a threshold value δE , we declare the energy ‘converged’ and output $A(t)$ as our optimal annealing schedule (figure 2). Otherwise, we add more points in between the current ones and optimize again. In figure 3, we show these intermediate steps of schedule refinement, from 3 parameters to 15 parameters that are required to converge close enough to the desired ground state.

In our optimization approach, we place no restrictions on the range of $A(t)$, but the schedules we find have a bounded range. This is important since in any physical implementation of the time-dependent Hamiltonian, the coupling strengths made available by the hardware will be limited.

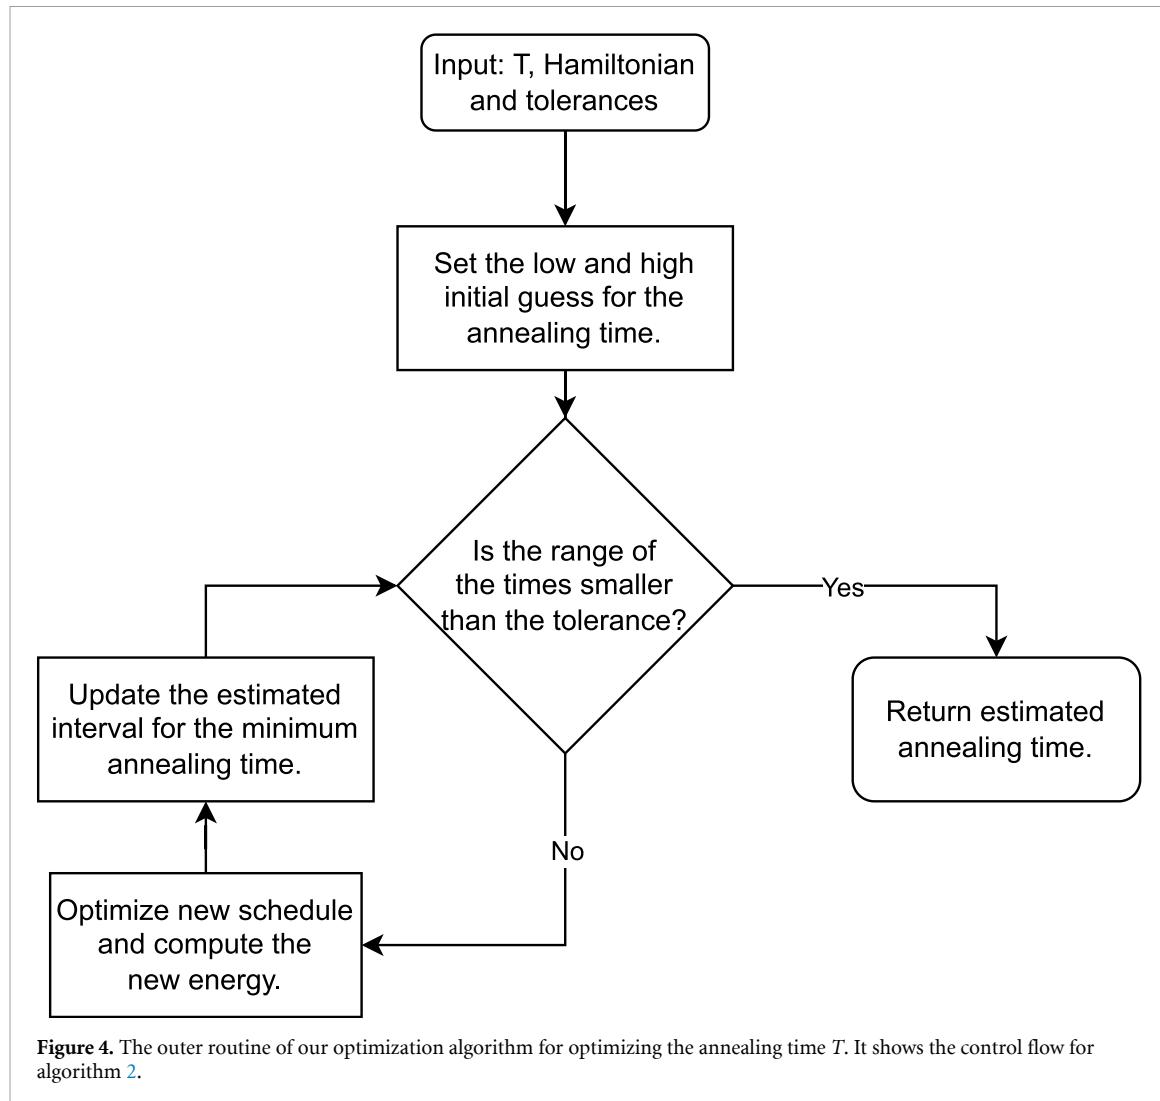
3.2. Optimizing the annealing time

The outer loop of the algorithm aims at determining the minimum annealing time T_{\min} that keeps us within our threshold $E(T) - E_0 \leq \Delta_E$ (see figure 4 and algorithm 2). We choose Δ_E to be a fraction c of the problem Hamiltonian gap (4):

$$\Delta_E(c) = 2c(J_L - J_R), \quad c > 0. \quad (9)$$

When $c = 1$, the threshold corresponds to the gap between the ground state and the first excited state of the problem Hamiltonian. In practice, we vary c to determine the performance of our algorithm.

To estimate T_{\min} , we build an estimation interval $[T_{\text{low}}, T_{\text{high}}]$ and take the estimate to be the upper limit of this interval. We start by setting $T_{\text{low}} = 0$. To set an initial T_{high} , we find a time T such that the corresponding energy $E(T)$ after optimizing the schedule satisfies $E(T) - E_0 \leq \Delta_E$. We declare such an optimization successful and set our current upper bound for the minimal annealing time as $T_{\text{high}} = T$.



If we are not successful, we set $T_{\text{low}} = T$ (see algorithm 2 for the edge case where $T_{\text{high}} = T_{\text{low}}$). Then, we update T :

$$T \leftarrow \frac{T_{\text{high}} + T_{\text{low}}}{2}. \quad (10)$$

As we continue refining our best estimate for T_{min} , our interval of uncertainty $[T_{\text{low}}, T_{\text{high}}]$ will shrink. We stop our algorithm when

$$\frac{T_{\text{high}} - T_{\text{low}}}{T_{\text{high}} + T_{\text{low}}} \leq \delta T = 0.05, \quad (11)$$

which is just the ratio between the uncertainty of the interval (half the range) and its midpoint $(T_{\text{high}} + T_{\text{low}})/2$.

By calculating T_{min} for multiple system sizes, we can study how the minimal annealing time scales with the system size. This is only an upper bound for the minimal annealing time $T_{\text{min}}(N)$ satisfying $E(T) \leq E_0 + \Delta_E$, because our optimization algorithm could always fail to find a schedule with a smaller T_{min} .

Algorithm 2. MinimizeAnnealingTime.

Input: $J_R, J_L, J, E_0, N, T, \delta T, \delta E, c, k_0$
Output: T_{\min}

$T_{\text{low}} \leftarrow 0;$
 $T_{\text{high}} \leftarrow T;$
while $\frac{T_{\text{high}} - T_{\text{low}}}{T_{\text{high}} + T_{\text{low}}} \geq \delta T$ **do**

$A(t) \leftarrow \text{FindSchedule} / / \text{Use algorithm 1};$
 $E(T) \leftarrow \text{ComputeEnergy}(A(t));$
if $E(T) - E_0 \leq 2c(J_L - J_R)$ **then**

$T_{\text{high}} \leftarrow T;$
 $T \leftarrow \frac{T_{\text{high}} + T_{\text{low}}}{2};$

else

$T_{\text{low}} \leftarrow T;$
if $T_{\text{high}} = T_{\text{low}}$ **then**

$T_{\text{high}} \leftarrow 2T_{\text{high}} / / \text{Increase guess};$
 $T \leftarrow T_{\text{high}};$

else

$T \leftarrow \frac{T_{\text{high}} + T_{\text{low}}}{2};$

end if

end if

end while

$T_{\min} \leftarrow T_{\text{high}};$
return $T_{\min};$

4. Results

In this section, we first compare the minimum annealing times obtained by means of our optimization (as described in section 3) to the times for a linear schedule to reach the same energy $E(T) \leq E_0 + \Delta_E$. We show that the optimized schedules can avoid the exponential times required by a linear schedule, resulting in annealing times scaling only polynomially with the system size. Second, we investigate what distinguishes the optimized schedules from the linear ones by studying the overlap of the instantaneous state with the ground state and first excited state of the instantaneous Hamiltonian $\hat{H}(t)$. In all the results we present, we set the system parameters to $(J_R, J_L, J) = (0.45, 0.5, 1)$.

4.1. Scaling of annealing time

In figure 5, we report the annealing times corresponding to both the linear schedules and the optimized ones, over different values of the threshold factor $c = 0.10, 0.25$ and 0.50 appearing in (9) (colors in legend). As expected from the discussion in section 2, we find that the linear schedules result in annealing times that grow exponentially fast with the size N of the system. Already, for $N > 9$ this would entail times $T > 10^6$ that are too computationally demanding to simulate.

In stark contrast, figure 5 demonstrates that our optimization protocol is capable of finding schedules significantly shorter than the linear ones (up to 5 orders of magnitude shorter for $N = 9$). Here, each data point corresponding to the optimized schedules is obtained as the minimum time found over 100 optimizations. Overall, we see that times only scale polynomially with N for all the values of $c \in \{0.10, 0.25, 0.50\}$ we study.

While this already provides evidence that optimized schedules can be exponentially shorter than linear ones, it is also of great interest to understand the scaling of annealing times entailed by a *typical* optimized schedule. For that purpose, we report in the inset of figure 5, the median values of the times found over the 100 repetitions. As shown, these median times can be up to one order of magnitude larger than the minimum times found. Despite this increase, these median times still scale polynomially in N . Fits of the form $\text{median}(T_{\min}) \sim \alpha(c)N^2$ (as a leading term) are displayed as dashed lines in the inset. As c decreases, the slopes $\alpha(c)$ increases.

4.2. Population levels

To understand the underlying dynamics associated with our optimized annealing schedules and low annealing times, we first inspect the schedules themselves. In figure 6, we report a selection of representative optimized schedules, identified for the threshold value $c = 0.10$. These schedules share a common structure across various system sizes, in particular a similar ‘up-down-up’ trajectory.

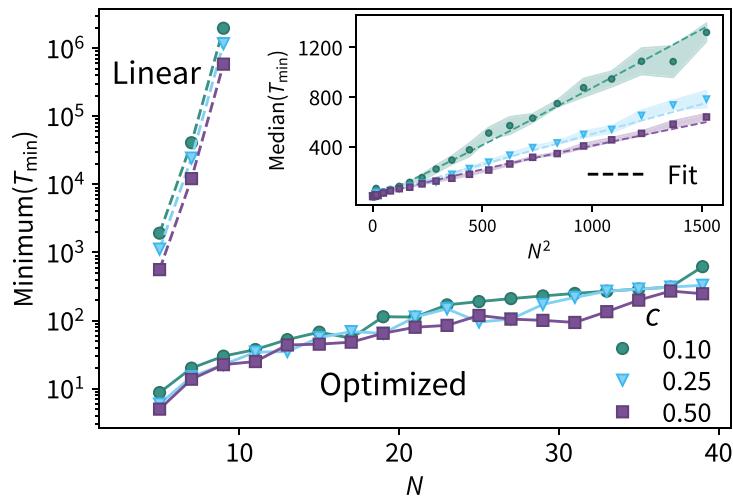


Figure 5. Scaling of the minimum annealing times T_{\min} found for the optimized schedules (solid lines) and the linear ones (dashed lines). These are obtained for odd system sizes $N \in [5, 39]$ and for different values of the energy threshold c (colors in legend). Each data point in the ‘Optimized’ part of the main figure represents the minimum annealing time found over 100 runs of our algorithm. Furthermore, in the inset we report the median values of these times over the 100 repetitions, along with a 95% confidence interval in these estimates (shaded region). In both cases, we see evidence of the polynomial scaling of the optimized T_{\min} with respect to N . In particular, the median values are found to scale quadratically with N , with a fit of the form αN^2 as a leading term, depicted in the inset (dashed lines).

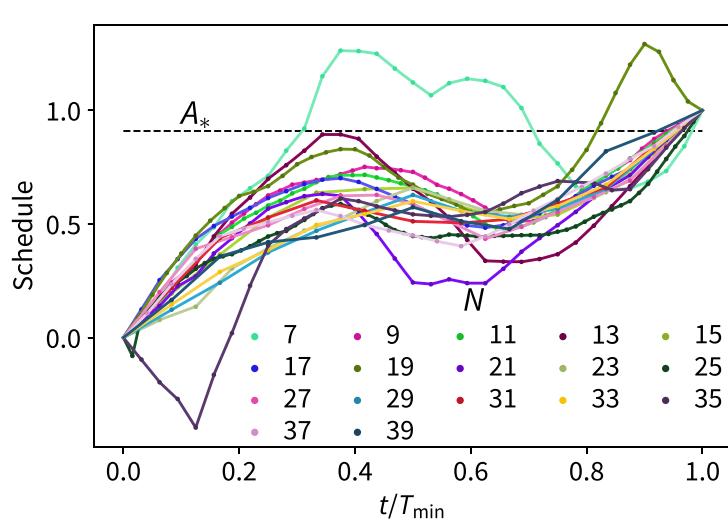


Figure 6. A selection of optimized schedules $A(t)$ across many system sizes that have been found for the case $c = 0.10$ in (9). We plot each schedule along the time t/T_{\min} that has been scaled by the overall time T_{\min} of the schedule. Dots in the figure depict the optimized points for the schedule, with the initial and final ones fixed to $(0, 0)$ and $(1, 1)$. The dashed black horizontal line represents the value of schedule A_* where there is a perturbative gap.

To understand the role of the non-monotonic schedule, we look at the populations in the energy levels of the instantaneous Hamiltonian $\hat{H}'(t) = \hat{P}_N \hat{H}(t) \hat{P}_N$, where $\hat{P}_N = (\hat{1}_N + \hat{X}^{\otimes N})/2$, and $\hat{1}_N$ is the identity operator of dimension 2^N . Because the initial state satisfies $\hat{P}_N |\psi(0)\rangle = |\psi(0)\rangle$ and our evolution preserves the symmetry ($[\hat{H}(t), \hat{X}^{\otimes N}] = 0$), we can restrict ourselves to states in that symmetry subspace. This means that the relevant eigenvalues and eigenvectors are those of the Hamiltonian $\hat{H}'(t)$.

By tracking how the populations in the energy levels of $\hat{H}'(t)$ change over time, we can observe how $A(t)$ affects the quantum state. Our simulations show that the majority of the energy population is restricted to the ground state and first excited state of $\hat{H}'(t)$. We show two examples in figure 7. The strategy followed by the optimized schedule is to first transfer population from the ground state to the first excited state and then cross the perturbative gap at $A \approx A_*$ one last time to diabatically transfer the population back into the ground state [39]. What is of particular interest is that the transfer of population to the first excited state is done by using a non-monotonic schedule. In some cases, the function $A(t)$ crosses the value $A_* \approx 0.9091$ a total of three times (twice increasing above A_* and once decreasing below it). In others, the function $A(t)$ approaches the

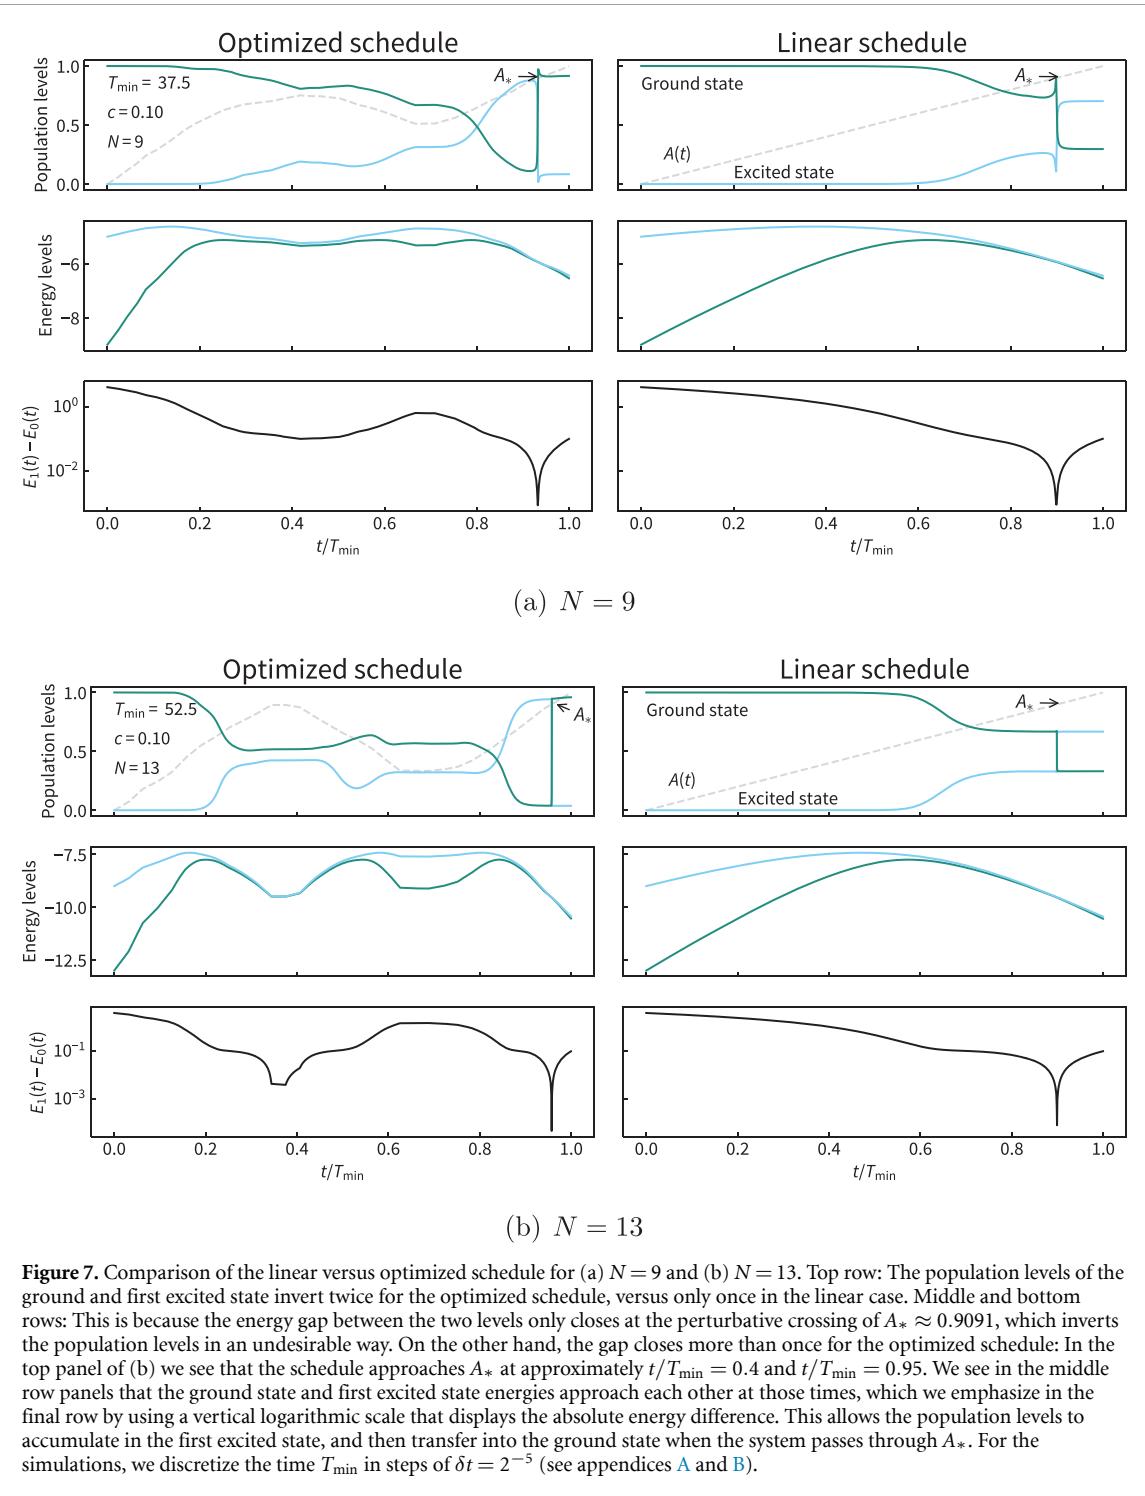


Figure 7. Comparison of the linear versus optimized schedule for (a) $N = 9$ and (b) $N = 13$. Top row: The population levels of the ground and first excited state invert twice for the optimized schedule, versus only once in the linear case. Middle and bottom rows: This is because the energy gap between the two levels only closes at the perturbative crossing of $A_* \approx 0.9091$, which inverts the population levels in an undesirable way. On the other hand, the gap closes more than once for the optimized schedule: In the top panel of (b) we see that the schedule approaches A_* at approximately $t/T_{\min} = 0.4$ and $t/T_{\min} = 0.95$. We see in the middle row panels that the ground state and first excited state energies approach each other at those times, which we emphasize in the final row by using a vertical logarithmic scale that displays the absolute energy difference. This allows the population levels to accumulate in the first excited state, and then transfer into the ground state when the system passes through A_* . For the simulations, we discretize the time T_{\min} in steps of $\delta t = 2^{-5}$ (see appendices A and B).

perturbative gap at A_* but does not cross it, moves away from it, and only crosses it near the end of the evolution. We show this in figure 6, where the horizontal dashed line indicates the schedule value A_* which leads to the perturbative gap.

To put our results into context, we compare our optimized schedules to linear schedules for the same optimized time T_{\min} in figure 7. Doing so highlights the advantages of a non-adiabatic evolution. We look at the instantaneous population levels, energy levels, and gap to the ground state along both schedules. We observe that the linear schedule only creates a small gap near the exponential closing A_* . If we anneal too fast, this exponentially small gap will cause the population levels of the ground and first excited state to invert. The only option is to anneal more slowly, which lowers the probability that the population levels invert. By contrast, some of our optimized schedules first make the gap to the ground state small *without* crossing A_* ,

allowing most of the population to migrate from the ground state to the first excited state. Then, the quantum state anneals through the point A_* , inverting the populations so the ground population is the majority.

A common feature of our optimized schedules is the ‘plateau’ in the populations at intermediate times. Naively, this might suggest that the annealing time could be further reduced since the system’s energy appears to be ‘idle’ during the plateau period, but we believe that the relative phase accumulation during this period is crucial to get a high first excited population near the end of the anneal that then transitions to a high ground state population.

In addition to gaining a better understanding of the physics at play, identifying common features of the optimized schedules as shown in figure 6 could serve to develop more informed initialization strategies, or to potentially refine the way the schedules are parameterized. For instance, one would expect that using the patterns identified so far as an initial guess for the schedules to be optimized over (rather than the random initialization adopted in this study) would significantly improve the convergence of our algorithm. We leave this study for future work.

5. Conclusion

Our work shows that we can trade the exponential annealing time required using an adiabatic approach for a polynomial annealing time using an optimized schedule with a non-adiabatic evolution on certain problems that are efficient to solve classically. By examining the population levels of a few schedules, we showed that the non-adiabatic nature of the schedules helps the system anneal more quickly into the ground state. Our Ising Hamiltonian (1) is one dimensional, so from a classical optimization perspective it is easy to find the ground state using heuristic algorithms since growing the size of domain walls along the spin chain does not cost additional energy. The problem’s exponential hardness for adiabatic QA stems from a perturbative crossing that arises from how the transverse field breaks the degeneracy of the low-lying energy eigenstates. Our work suggests that there are relatively simple annealing schedules that can allow for a polynomial time annealing algorithm to solve this problem. This highlights the importance of sufficient flexibility in the annealing schedule to achieve such an advantage.

However, it should be noted that the saving in annealing time could be offset by the effort spent in optimizing the schedules in the first place, rendering such an approach less attractive. We provide several points of indirect evidence that the cost of optimization remains under control. Most importantly, for the fixed amount of computational resources allocated for each run of the optimization of the algorithm, we were able to optimize schedules for up to $N = 39$, which we would not expect if the optimization effort was growing exponentially with the system size. Furthermore, we fixed the amount of computational resources when running each schedule optimization and did not see any particular increase in the number of steps of schedules refinement required as the system size increases. This suggests that for this problem Hamiltonian, the cost of the schedule optimization is unlikely to be exponential in the system size.

We can consider following a similar strategy for solving more complex problems. For spin glasses, we can expect our schedule optimization to help us reach the low energy states of the rugged landscape more quickly, but it is unrealistic to imagine that an efficient schedule optimization would allow us to identify the global optimum efficiently. Nevertheless, it may be sufficient to find high quality low-energy solutions quickly depending on the optimization context.

We do not claim that the schedules we have identified are optimal in the sense that they achieve the highest ground state probability for the lowest annealing time. We have observed that different choices of initial schedules can affect the best schedules found by our heuristic algorithm, so the optimization landscape is non-monotonic and includes schedules with different characteristics and possibly some with higher quality than the ones we present. A thorough investigation of this possibility using quantum optimal control for QA [35] would however be limited to small system sizes due to the high computational cost.

It is worth noting that the performance of DQA is likely to be related to the performance of another quantum algorithm for tackling optimization problems, the QAOA [40], in the quantum circuit paradigm. Recent works [31, 41] have highlighted that in the heuristic setting where the QAOA parameters are trained, the algorithm can be understood as implementing a diabatic evolution as in DQA. In fact, it has been shown that a hybrid of the two approaches, whereby the unitary evolution is allowed to be defined by both a continuous interpolating Hamiltonian as well as Hamiltonians that are turned on and off, is in fact optimal for solving optimization problems in this heuristic way [35, 42]. Our results for DQA may therefore indicate that similar behavior might be observed for QAOA.

As identified and discussed in section 4.2, the existence of a shared structure of the optimized schedules across different system sizes could serve as a basis of improved initialization of the algorithm. Such feature

has been observed recently in the context of QAOA [31, 41, 43, 44]. Given the difficulties that would be encountered when performing optimization directly based on measurement data—with in particular the flattening of the cost function values [45, 46]—this may prove crucial for realistic optimization of annealing schedules at scale.

There are several aspects of our optimization protocol that could be modified. These include the choice of driver Hamiltonian or intermediate Hamiltonian [38, 47], the initial quantum state, and the way we parameterize the schedule. Instead of using the average energy $E(T)$ for the optimization, we could for example try to minimize different quantiles of the energy distribution generated by the final state. An optimization procedure that includes the cost of identifying the optimal schedule using optimal stopping techniques [29] would verify the viability of such strategies for real world optimization.

Our hope is that such optimization protocols can be used to create targeted schedules for highly-controllable quantum annealers, offering non-trivial operating benchmarks for such devices as well as removing the bottlenecks of the standard adiabatic paradigm when possible.

Data availability statements

The data that support the findings of this study are openly available at the following URL/DOI: <https://doi.org/10.5281/zenodo.7613047>.

The repository [48] includes: data for the optimized schedules in figure 6, the scaling of the annealing time, the population levels, and a Qiskit implementation of our algorithm, as well as code for calculating the population levels.

Acknowledgments

This work was supported by the U.S. Department of Energy (DOE) through a quantum computing program sponsored by the Los Alamos National Laboratory (LANL) Information Science & Technology Institute. L C and M L were supported by the U.S. DOE, Office of Science, Office of Advanced Scientific Computing Research, under the Accelerated Research in Quantum Computing (ARQC) program. F S acknowledges Laboratory Directed Research and Development (LDRD) program of LANL under Project Number 20220745ER. J C was supported by a B2X scholarship from the Fonds de recherche—Nature et technologies and a scholarship from the Natural Sciences and Engineering Research Council of Canada (Funding Reference Number: 456431992). J C also acknowledges the physics PhD program at the Université de Sherbrooke and funding from the Canada First Research Excellence Fund. This material is based upon work supported by the National Science Foundation under Grant No. 2037755. T A thanks Jarrett Smalley for useful discussions.

Appendix A. Simulating the schedule

For the optimizations in section 4.1, we study system sizes of up to $N = 39$ qubits. Simulations with such a large number of qubits are possible by numerical routines first proposed in [49] and implemented in [50]. The simulation works at the level of the system's dynamical Lie algebra [51, 52], the Lie closure of the set $\{i\hat{H}_d, i\hat{H}_p\}$, which in our case scales efficiently (polynomially) in the system size. We perform Heisenberg evolution of the observable \hat{H}_p . Crucially, irrespective of the time T considered and irrespective of the schedule employed, the evolved operator remains contained in the dynamical Lie algebra, and therefore admits an efficient description. That is, by keeping track of the instantaneous support of the observable over a basis of the dynamical Lie algebra, one can assess the energy of the evolved system at scale.

However, the study of the instantaneous population in section 4.2 requires full knowledge of the system over time, and therefore we resort to conventional full-state vector simulations, thus limiting the number of qubits we can simulate. For the example schedule of figure 3 and the population levels in section 4.2, we use Qiskit [53] to simulate the time evolution of the schedule. We write the time-evolved state $|\psi(T)\rangle$ at the end of the annealing procedure as

$$|\psi(T)\rangle = \hat{U}(T)|\psi(0)\rangle, \quad (\text{A.1})$$

where $|\psi(0)\rangle = |+\rangle^{\otimes N}$, and the unitary $\hat{U}(T)$ is the time-ordered propagator

$$\hat{U}(T) \equiv \text{T exp} \left(-i \int_0^T \hat{H}(t) dt \right). \quad (\text{A.2})$$

We approximate $\hat{U}(T)$ with small time slices δt and implement the evolution as a quantum circuit. This means the unitary in (A.2) decomposes as:

$$\hat{U}(T) \approx U(T, T - \delta t) \dots \hat{U}(2\delta t, \delta t) \hat{U}(\delta t, 0), \quad (\text{A.3})$$

where $\hat{U}(t + \delta t, t)$ is the unitary between times t and $t + \delta t$. For small enough δt , we assume the Hamiltonian $\hat{H}(t)$ is constant in a given interval.

Let us focus our attention on one particular time slice at time t^* , with the associated unitary operator

$$\begin{aligned} \hat{U}(t^* + \delta t, t^*) &\equiv \exp[-i\hat{H}(t^*)\delta t] \\ &= \exp[-i\delta t(A(t^*)\hat{H}_p + (1 - A(t^*))\hat{H}_d)]. \end{aligned} \quad (\text{A.4})$$

Because the terms in \hat{H}_d and \hat{H}_p do not commute (one contains \hat{Z} terms and the other contains \hat{X} terms), we decompose the unitary slice with a first-order Trotter decomposition [54, 55]:

$$\hat{U}(t^* + \delta t, t^*) \approx \exp \left[-i\delta t \left(-A(t^*) \sum_{j=1}^N J_j \hat{Z}_j \hat{Z}_{j+1} \right) \right] \times \exp \left[-i\delta t \left(-(1 - A(t^*)) \sum_{j=1}^N \hat{X}_j \right) \right]. \quad (\text{A.5})$$

Within each exponential in (A.5), the operators commute with each other for all i , which allows us to rewrite the sum as a product. We then find the final form for our unitary:

$$\hat{U}(t^* + \delta t, t^*) \approx \prod_{j=1}^N \exp \left[-i\delta t \left(-A(t^*) J_j \hat{Z}_j \hat{Z}_{j+1} \right) \right] \times \prod_{j=1}^N \exp \left[-i\delta t \left(-(1 - A(t^*)) \hat{X}_j \right) \right]. \quad (\text{A.6})$$

Equation (A.6) tells us that each time slice of size δt requires $2N$ gates (there are N terms in each product), and the two types of quantum gates are:

$$\begin{aligned} \text{RZZ}(j, j+1, \theta_Z) &\equiv \exp \left[-i \frac{\theta_Z}{2} \hat{Z}_j \hat{Z}_{j+1} \right], \\ \text{RX}(j, \theta_X) &\equiv \exp \left[-i \frac{\theta_X}{2} \hat{X}_j \right], \end{aligned} \quad (\text{A.7})$$

where in our case the angles are $\theta_Z = -2\delta t A(t^*) J_j$ and $\theta_X = -2\delta t (1 - A(t^*))$.

This collection of gates in (A.6) defines a quantum circuit for each time slice. The full evolution involves the same layout of gates over and over again, but with varying θ_Z and θ_X , according to the schedule $A(t)$, and the couplings J_j . In figure A1, we show the circuit according to this layout.

In order to calculate $E(T) \equiv \langle \psi(T) | \hat{H}_p | \psi(T) \rangle$, we choose to discretize the unitary evolution operator in terms of time steps δt , and the output state is then given by $|\psi(T, \delta t)\rangle$. Calculating the energy then requires sending the discretization to zero:

$$E(T) = \lim_{\delta t \rightarrow 0} \langle \psi(T, \delta t) | \hat{H}_p | \psi(T, \delta t) \rangle. \quad (\text{A.8})$$

We do this by building the circuit for multiple values of δt , and checking how the energy $E(T)$ changes as we decrease the size of δt . We use the same threshold δE as in section 3.1 to stop this procedure, giving us our ‘converged’ estimate of $E(T)$ (see appendix B).

To calculate the quantities from section 4.2, we track the instantaneous quantum state $|\psi(t)\rangle$ during the annealing schedule. We also calculate the instantaneous Hamiltonian in the relevant symmetry subspace $\hat{H}'(t) = \hat{P}_N \hat{H}(t) \hat{P}_N$, where $\hat{P}_N = (\hat{1}_N + \hat{X}^{\otimes N})/2$, and $\hat{1}_N$ is the identity operator of dimension 2^N .

Diagonalizing $\hat{H}'(t)$ provides us with $E_0(t)$ and $E_1(t)$ in the relevant symmetry subspace, so we can calculate the relevant energy gap for our dynamics.

In order to calculate the instantaneous populations of a given energy eigenstate, we find the eigenvectors $|E_k(t)\rangle$ of $\hat{H}'(t)$, with k indexing the energy levels. We then define the population level as:

$$P_k(t) \equiv |\langle \psi(t) | E_k(t) \rangle|^2, \quad k \in \{0, 1, 2, \dots, 2^N - 1\}, \quad (\text{A.9})$$

where $|E_k(t)\rangle$ is the eigenstate of the Hamiltonian $\hat{H}'(t)$ with energy $E_k(t)$. The population levels are fractions, so they must sum to one:

$$\sum_{k=0}^{2^N-1} P_k(t) = 1. \quad (\text{A.10})$$

In figure 7, we show $P_0(t)$ and $P_1(t)$, since they contribute the most to the annealing dynamics.

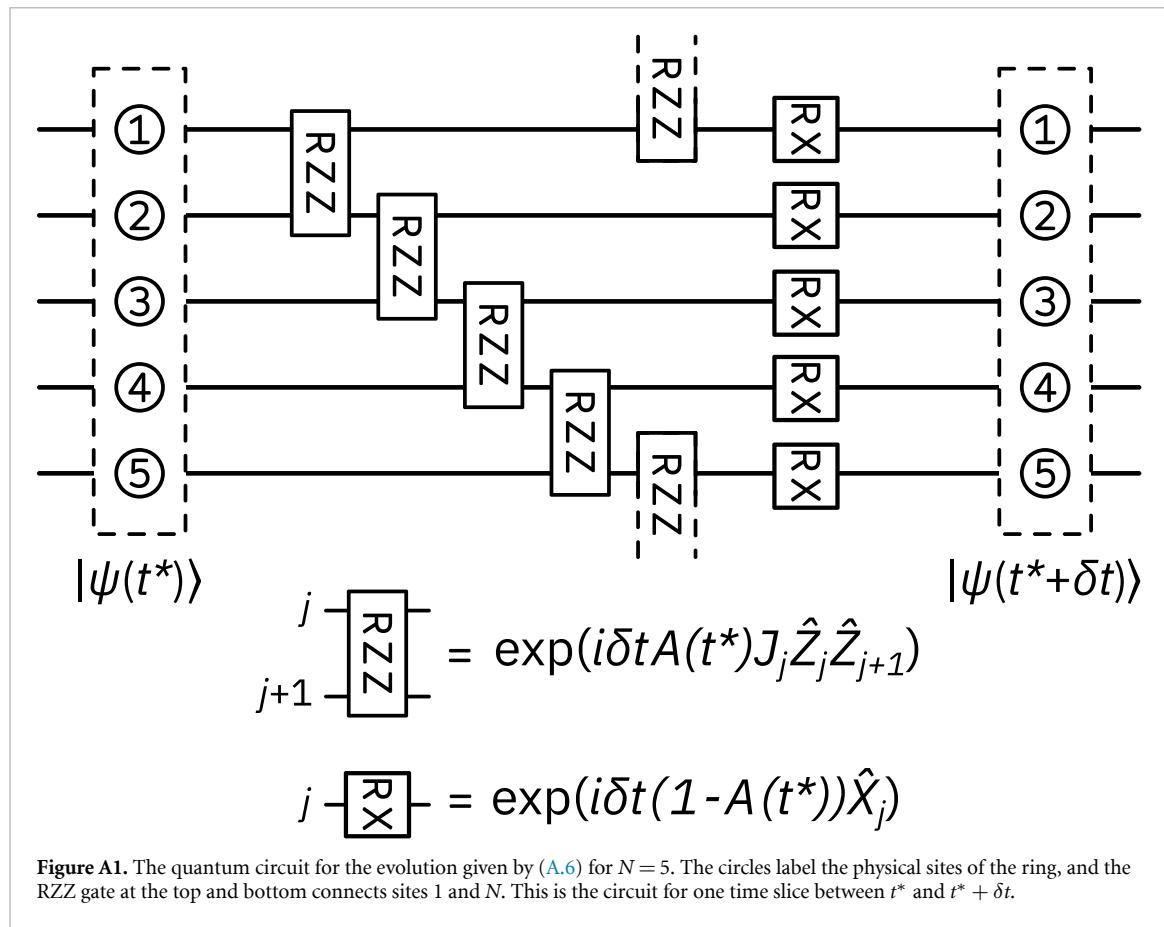


Figure A1. The quantum circuit for the evolution given by (A.6) for $N = 5$. The circles label the physical sites of the ring, and the RZZ gate at the top and bottom connects sites 1 and N . This is the circuit for one time slice between t^* and $t^* + \delta t$.

Appendix B. Implementation details

For our numerical work, we use the minimization procedure implemented in SciPy with the COBYLA method [56]. We set the maximum number of iterations ‘maxiter’ of the function to be 800, and the tolerance for the gradient ‘tol’ to be 0.001.

For every numerical experiment, we set $J_R = 0.45$, $J_L = 0.50$, and $J = 1$. We discretize the total annealing time T with an initial step size of $\delta t = 1$, except for the population levels, where we set the initial δt to be 2^{-4} (which is small enough such that the final δt we have when the algorithm converges is 2^{-5}).

The thresholds we use to determine when to stop our algorithm are:

- (i) $\delta T = 0.05$.
- (ii) $c \in \{0.10, 0.25, 0.50\}$.
- (iii) $\delta E = 0.001$.

ORCID iDs

Jeremy Côté <https://orcid.org/0000-0001-9594-2145>
 Frédéric Sauvage <https://orcid.org/0003-3363-5929>
 Tameem Albash <https://orcid.org/0003-3916-3985>

References

- [1] Grover L K 1997 *Phys. Rev. Lett.* **79** 325–8
- [2] Brassard G, Høyer P, Mosca M and Tapp A 2002 Quantum amplitude amplification and estimation *Quantum Computation and Information (Contemporary Mathematics)* vol 305 (American Mathematical Society) pp 53–74
- [3] Cerf N J, Grover L K and Williams C P 2000 *Phys. Rev. A* **61** 032303
- [4] Roland J and Cerf N J 2002 *Phys. Rev. A* **65** 042308
- [5] Childs A M, Cleve R, Deotto E, Farhi E, Gutmann S and Spielman D A 2003 Exponential algorithmic speedup by a quantum walk *Proc. 35th Annual ACM Symp. on Theory of Computing (STOC '03)* (Association for Computing Machinery) pp 59–68
- [6] Ambainis A 2004 *SIGACT News* **35** 22–35
- [7] Wiebe N, Braun D and Lloyd S 2012 *Phys. Rev. Lett.* **109** 050505
- [8] Somma R D, Nagaj D and Kieferová M 2012 *Phys. Rev. Lett.* **109** 050501

- [9] Montanaro A 2018 *Theory Comput.* **14** 1–24
- [10] Brandao F G and Svore K M 2017 Quantum speed-ups for solving semidefinite programs 2017 IEEE 58th Annual Symp. on Foundations of Computer Science (FOCS) pp 415–26
- [11] Mandrà S, Guerreschi G G and Aspuru-Guzik A 2016 *New J. Phys.* **18** 073003
- [12] Ambainis A and Kokainis M 2017 Quantum algorithm for tree size estimation, with applications to backtracking and 2-player games *Proc. 49th Annual ACM SIGACT Symp. on Theory of Computing* pp 989–1002
- [13] Campbell E, Khurana A and Montanaro A 2019 *Quantum* **3** 167
- [14] Brandão F G S L, Kueng R and Stilck França D 2022 *Quantum* **6** 625
- [15] Born M and Fock V 1928 *Z. Phys.* **51** 165–80
- [16] Kato T 1950 *J. Phys. Soc. Japan* **5** 435–9
- [17] Jansen S, Ruskai M B and Seiler R 2007 *J. Math. Phys.* **48** 102111
- [18] Kadewaki T and Nishimori H 1998 *Phys. Rev. E* **58** 5355–63
- [19] Farhi E, Goldstone J, Gutmann S, Lapan J, Lundgren A and Preduce D 2001 *Science* **292** 472–5
- [20] Farhi E, Goldstone J, Gutmann S and Sipser M 2000 arXiv:quant-ph/0001106
- [21] Aharonov D, van Dam W, Kempe J, Landau Z, Lloyd S and Regev O 2007 *SIAM J. Comput.* **37** 166–94
- [22] van Dam W, Mosca M and Vazirani U 2001 How powerful is adiabatic quantum computation? *Proc. 42nd IEEE Symp. on Foundations of Computer Science* pp 279–87
- [23] Reichardt B W 2004 The quantum adiabatic optimization algorithm and local minima *Proc. 36th Annual ACM Symp. on Theory of Computing* pp 502–10
- [24] Crosson E J and Lidar D A 2021 *Nat. Rev. Phys.* **3** 466–89
- [25] Guéry-Odelin D, Ruschhaupt A, Kiely A, Torrontegui E, Martínez-Garaot S and Muga J G 2019 *Rev. Mod. Phys.* **91** 045001
- [26] Berry M V 2009 *J. Phys. A: Math. Theor.* **42** 365303
- [27] Roberts D, Cincio L, Saxena A, Petukhov A and Knysh S 2020 *Phys. Rev. A* **101** 042317
- [28] Amin M H S and Choi V 2009 *Phys. Rev. A* **80** 062326
- [29] Vinci W and Lidar D A 2016 *Phys. Rev. Appl.* **6** 054016
- [30] Rønnow T F, Wang Z, Job J, Boixo S, Isakov S V, Wecker D, Martinis J M, Lidar D A and Troyer M 2014 *Science* **345** 420–4
- [31] Zhou L, Wang S T, Choi S, Pichler H and Lukin M D 2020 *Phys. Rev. X* **10** 021067
- [32] Quiroz G 2019 *Phys. Rev. A* **99** 062306
- [33] Hegde P R, Passarelli G, Scocco A and Lucignano P 2022 *Phys. Rev. A* **105** 012612
- [34] Yang Z C, Rahmani A, Shabani A, Neven H and Chamon C 2017 *Phys. Rev. X* **7** 021027
- [35] Brady L T, Baldwin C L, Bapat A, Kharkov Y and Gorshkov A V 2021 *Phys. Rev. Lett.* **126** 070505
- [36] Chen Y Q, Chen Y, Lee C K, Zhang S and Hsieh C Y 2022 *Nat. Mach. Intell.* **4** 269–78
- [37] Hegde P R, Passarelli G, Cantele G and Lucignano P 2023 Deep learning optimal quantum annealing schedules for random Ising models *New J. Phys.* **25** 073013
- [38] Matsuura S, Buck S, Senicourt V and Zaribafian A 2021 *Phys. Rev. A* **103** 052435
- [39] Crosson E, Farhi E, Yen-Yu Lin C, Lin H H and Shor P 2014 arXiv:1401.7320
- [40] Farhi E, Goldstone J and Gutmann S 2014 arXiv:1411.4028
- [41] Pagano G et al 2020 *Proc. Natl Acad. Sci.* **117** 25396–401
- [42] Brady L T, Kocia L, Bienias P, Bapat A, Kharkov Y and Gorshkov A V 2021 arXiv:2107.01218
- [43] Sauvage F, Sim S, Kunitsa A A, Simon W A, Mauri M and Perdomo-Ortiz A 2021 arXiv:2103.08572
- [44] Mele A A, Mbeng G B, Santoro G E, Collura M and Torta P 2022 Avoiding barren plateaus via transferability of smooth solutions in a Hamiltonian variational ansatz *Phys. Rev. A* **106** L060401
- [45] McClean J R, Boixo S, Smelyanskiy V N, Babbush R and Neven H 2018 *Nat. Commun.* **9** 1–6
- [46] Cerezo M, Sone A, Volkoff T, Cincio L and Coles P J 2021 *Nat. Commun.* **12** 1–12
- [47] Zeng L, Zhang J and Sarovar M 2016 *J. Phys. A: Math. Theor.* **49** 165305
- [48] Côté J, Sauvage F, Larocca M, Jonsson M, Cincio L and Albash T 2023 Diabatic quantum annealing for the frustrated ring model *Zenodo* (<https://doi.org/10.5281/zenodo.7613047>)
- [49] Somma R D 2005 arXiv:quant-ph/0512209
- [50] Goh M, Larocca M, Sauvage F, Cincio L and Cerezo M 2023 arXiv:2308.01432
- [51] Larocca M, Czarnik P, Sharma K, Muraleedharan G, Coles P J and Cerezo M 2022 *Quantum* **6** 824
- [52] Larocca M, Ju N, García-Martín D, Coles P J and Cerezo M 2023 Theory of overparametrization in quantum neural networks *Nat. Comput. Sci.* **3** 542–51
- [53] Anis M S 2021 Qiskit (Zenodo) (<https://doi.org/10.5281/zenodo.2573505>)
- [54] Trotter H F 1959 *Proc. Am. Math. Soc.* **10** 545–51
- [55] Suzuki M 1976 *Commun. Math. Phys.* **51** 183–90
- [56] Powell M J D 1994 A direct search optimization method that models the objective and constraint functions by linear interpolation *Advances in Optimization and Numerical Analysis* (Springer) pp 51–67