

# Worldline Susceptibility as a Surrogate for Spectral Gaps in Quantum Annealing Schedule Design

Suraj Singh,<sup>1,2,3,\*</sup> Lakshya Nagpal,<sup>1,2,3,\*</sup> Vikas Chauhan,<sup>4,2</sup> and S.R. Hassan<sup>1,2,5</sup>

<sup>1</sup>*The Institute of Mathematical Sciences (IMSc),  
C.I.T Campus, Taramani, Chennai 600113, India*

<sup>2</sup>*QCAR Group, The Institute of Mathematical Sciences (IMSc), Chennai 600113, India*

<sup>3</sup>*Pecslab Research*

<sup>4</sup>*Department of Physics, Ramjas College, University of Delhi, Delhi 110007, India*

<sup>5</sup>*Homi Bhabha National Institute, Anushakti Nagar, Mumbai, Maharashtra 400094*

The performance of quantum annealing depends critically on the annealing schedule, which determines how computational effort is distributed along the evolution. Although the Roland–Cerf local-adiabatic schedule is theoretically optimal, it requires complete knowledge of the instantaneous spectral gap, whose computation is exponentially expensive and therefore impractical for large optimization problems. In this work, we propose a computationally inexpensive surrogate based on the worldline magnetization susceptibility measured during simulated quantum annealing (SQA). The susceptibility is obtained directly from the Suzuki–Trotter representation through classical Monte Carlo sampling and identifies the critical region of the anneal without requiring any knowledge of the quantum spectrum. Using exact diagonalization of small Sherrington–Kirkpatrick spin-glass instances as ground truth, we compare the surrogate schedule with both the conventional linear schedule and the exact Roland–Cerf schedule. The surrogate consistently outperforms the linear schedule and, for a substantial fraction of instances, also surpasses the exact local-adiabatic schedule. We show that this unexpected behaviour originates from two finite-time failure modes of the exact schedule: a boundary-gap trap and a coherent oscillatory instability arising from highly non-uniform time allocation. In contrast, the surrogate schedule remains robust because it identifies the critical region without overfitting to the detailed structure of the spectral gap. These results suggest that inexpensive classical observables can provide a practical and scalable foundation for annealing schedule design, and that robustness may be more important than exact spectral information for optimizing finite-time quantum annealing.

## I. INTRODUCTION

Quantum annealing has emerged as a promising approach for solving combinatorial optimization problems by encoding them into the ground state of an Ising Hamiltonian. A wide variety of practical problems, including scheduling, routing, portfolio optimization, machine learning, and spin-glass optimization, can be formulated in this framework and solved through an adiabatic evolution from a simple driver Hamiltonian to the target problem Hamiltonian [1–5]. The availability of dedicated quantum annealing hardware has further stimulated interest in understanding how annealing dynamics can be exploited to solve increasingly difficult optimization problems.

The success of quantum annealing, however, depends not only on the problem Hamiltonian but also on how the annealing parameter is varied in time. For a fixed total annealing time, traversing the difficult part of the evolution too rapidly can drive the system out of its instantaneous ground state, while slowing down uniformly over the entire evolution wastes valuable computational time. The annealing schedule therefore plays a central role

in determining the final ground-state success probability. Designing schedules that allocate computational effort where it is most needed offers one of the few ways to improve annealing performance without modifying either the optimization problem or the underlying hardware.

Theoretically, this problem has an elegant solution. Roland and Cerf showed that the annealing schedule should satisfy a local adiabaticity condition in which the evolution slows down in inverse proportion to the square of the instantaneous spectral gap [6]. This schedule minimizes the total annealing time while maintaining adiabatic evolution and has become the standard theoretical benchmark for optimal scheduling. Unfortunately, implementing this prescription requires complete knowledge of the instantaneous spectral gap throughout the annealing trajectory. Since evaluating the low-energy spectrum generally requires repeated diagonalization of an exponentially large Hamiltonian, the exact local-adiabatic schedule rapidly becomes computationally intractable for realistically sized optimization problems.

This practical limitation naturally raises an important question. Rather than attempting to compute the exact spectral gap, can one identify a computationally inexpensive observable that carries essentially the same scheduling information? Such a surrogate would eliminate the need

---

\* These authors contributed equally to this work.

for repeated diagonalization while retaining the ability to identify the difficult region of the anneal where additional computational time should be allocated.

## II. EXACT AND SURROGATE SCHEDULING

### A. Quantum Annealing and Local-Adiabatic Scheduling

Quantum annealing solves an optimization problem by continuously transforming the ground state of a simple driver Hamiltonian into the ground state of a problem Hamiltonian. The evolution is governed by

$$\hat{H}(s) = A(s)\hat{H}_D + B(s)\hat{H}_P, \quad 0 \leq s \leq 1, \quad (1)$$

where  $\hat{H}_D$  is the driver Hamiltonian,  $\hat{H}_P$  encodes the optimization problem, and  $A(s)$  and  $B(s)$  are annealing schedules satisfying

$$A(0) = 1, \quad B(0) = 0, \quad A(1) = 0, \quad B(1) = 1. \quad (2)$$

Initially the system is prepared in the easily accessible ground state of  $\hat{H}_D$ . During the anneal the transverse field is gradually reduced while the problem Hamiltonian is simultaneously switched on, with the goal of ending in the ground state of  $\hat{H}_P$ .

The success of this evolution is governed by the instantaneous spectral gap

$$\Delta(s) = E_1(s) - E_0(s), \quad (3)$$

where  $E_0(s)$  and  $E_1(s)$  denote the instantaneous ground- and first-excited-state energies of  $\hat{H}(s)$ . Small spectral gaps require slower evolution to suppress non-adiabatic transitions and therefore determine the computational difficulty of the annealing process.

Roland and Cerf showed that, under the local adiabatic approximation, the annealing speed should satisfy

$$\left| \frac{ds}{dt} \right| = \varepsilon \frac{\Delta(s)^2}{\left| \left\langle E_1(s) \left| \partial_s \hat{H} \right| E_0(s) \right\rangle \right|}, \quad (4)$$

where  $\varepsilon$  controls the allowable adiabatic error. This schedule allocates more annealing time to regions where the spectral gap becomes small and is optimal within the assumptions of local adiabatic evolution.

Although Eq. (4) provides an elegant theoretical prescription, its practical applicability is severely limited. Constructing the schedule requires knowledge of the complete instantaneous low-energy spectrum throughout the anneal, including both  $\Delta(s)$  and the transition matrix element in Eq. (4). For generic combinatorial optimization problems these quantities are exponentially expensive to compute, making the local-adiabatic schedule inaccessible for problem sizes of practical interest.

This limitation motivates the central question of the present work: can a computationally inexpensive quantity, directly measurable within simulated quantum annealing, identify the difficult regions of the annealing process well enough to construct an effective annealing schedule without explicit knowledge of the spectral gap?

### B. Worldline Magnetization Susceptibility in Simulated Quantum Annealing

To construct an annealing schedule without explicit knowledge of the spectral gap, we seek an observable that can be measured directly during a simulated quantum annealing (SQA) simulation. SQA replaces the quantum annealing dynamics of the transverse-field Ising model by an equivalent classical statistical-mechanical system obtained through the Suzuki–Trotter decomposition. This mapping allows quantum expectation values to be estimated using classical Monte Carlo sampling while avoiding direct diagonalization of the quantum Hamiltonian.

Applying the Suzuki–Trotter decomposition to Eq. (1) yields the effective classical action

$$S_{\text{eff}} = \frac{\beta}{M} \sum_{k=1}^M H_P(\mathbf{s}^{(k)}) - J_{\perp} \sum_{k=1}^M \sum_{i=1}^N s_i^{(k)} s_i^{(k+1)}, \quad (5)$$

where  $M$  is the number of Trotter slices,  $\beta$  is the inverse temperature, and

$$J_{\perp}(\beta, \Gamma) = \frac{1}{2} \ln \left[ \frac{1}{\tanh(\beta\Gamma/M)} \right] \quad (6)$$

is the effective coupling along the imaginary-time direction. The resulting classical system consists of  $M$  coupled replicas of the original Ising problem, with the additional Trotter coupling encoding the quantum fluctuations generated by the transverse field.

The natural order parameter for the worldline system is the imaginary-time-averaged magnetization

$$m = \frac{1}{NM} \sum_{i=1}^N \sum_{k=1}^M s_i^{(k)}, \quad (7)$$

which measures the mean polarization of spins across all  $N$  sites and  $M$  Trotter slices. This quantity encodes the Edwards–Anderson-type ordering that develops as the system crosses from the quantum-disordered to the classically ordered regime during the anneal.

The susceptibility of this worldline order parameter is

$$\chi_m = NM(\langle m^2 \rangle - \langle |m| \rangle^2), \quad (8)$$

where the angular brackets denote equilibrium averages over the Monte Carlo ensemble. The prefactor  $NM$  ensures that  $\chi_m$  is an extensive quantity that scales with

system and worldline volume. This susceptibility is evaluated directly from the Monte Carlo ensemble generated during the SQA simulation and is therefore computationally inexpensive.

The physical significance of  $\chi_m$  follows from its relation to the static structure factor at zero wavevector,

$$\chi_m \propto S(q=0, \omega=0) \propto \frac{1}{\Delta(s)^2} \quad \text{near the QCP}, \quad (9)$$

where the second proportionality is the quantum-critical scaling relation between the longitudinal susceptibility and the spectral gap. As the transverse field is varied during the annealing process,  $\chi_m$  changes continuously and exhibits a pronounced maximum in the crossover region between the quantum-disordered and classically ordered regimes. This peak is located near the quantum critical point where  $\Delta(s)$  is smallest, making  $\chi_m$  a direct probe of the same critical region as the spectral gap, without requiring any diagonalization of the quantum Hamiltonian.

In the following subsection we exploit this observable to construct an annealing schedule directly from  $\chi_m(s)$ .

### C. Construction of the Worldline-Susceptibility Schedule

The central idea of this work is to replace the inaccessible spectral gap in the local-adiabatic scheduling condition by the worldline magnetization susceptibility, which can be measured directly during an SQA simulation. From the scaling relation Eq. (9), the surrogate approximation

$$\chi_m(s) \approx \frac{\text{const}}{\Delta(s)^2} \quad \text{near the QCP} \quad (10)$$

motivates replacing  $\Delta(s)^{-2}$  in the Roland-Cerf time-budget integrand with  $\chi_m(s)$ .

To convert the measured susceptibility into an annealing schedule, we first define a positive time-allocation weight

$$w(s) = \chi_m(s) + \chi_0, \quad (11)$$

where  $\chi_0 > 0$  is a small regularization constant that prevents the annealing velocity from diverging in regions where the susceptibility becomes very small.

The cumulative fraction of annealing time assigned up to parameter value  $s$  is then defined as

$$\tau(s) = \frac{\int_0^s w(s') ds'}{\int_0^1 w(s') ds'}. \quad (12)$$

By construction,

$$\tau(0) = 0, \quad \tau(1) = 1, \quad (13)$$

and  $\tau(s)$  is a monotonic function of  $s$ . It therefore provides a normalized cumulative distribution describing how the total annealing time is allocated along the annealing path.

The annealing schedule is obtained by inverting this cumulative mapping,

$$s(t) = \tau^{-1}\left(\frac{t}{T}\right), \quad 0 \leq t \leq T, \quad (14)$$

where  $T$  denotes the total annealing time. Differentiating this relation gives the instantaneous annealing velocity,

$$\frac{ds}{dt} = \frac{1}{T w(s)} = \frac{1}{T [\chi_m(s) + \chi_0]}. \quad (15)$$

This construction naturally slows the evolution in regions where the worldline susceptibility is large and accelerates it where the susceptibility is small, while preserving the prescribed total annealing time. Unlike the Roland-Cerf schedule, the proposed procedure requires no knowledge of the instantaneous energy spectrum or minimum spectral gap. Every quantity entering the schedule is obtained directly from the classical SQA simulation. The resulting algorithm therefore provides a practical surrogate for local-adiabatic scheduling that can be constructed without diagonalizing the quantum Hamiltonian.

### D. Overall Computational Workflow

The complete workflow of the proposed surrogate scheduling method is summarized below.

1. Construct the problem Hamiltonian

$$H(s) = (1-s)H_D + sH_P. \quad (16)$$

2. Compute the exact instantaneous spectral gap  $\Delta(s)$  by exact diagonalization. This step is performed only for benchmarking and is not required by the proposed method.
3. Perform an independent SQA simulation of the same annealing process and measure the worldline magnetization susceptibility

$$\chi_m(s) = NM(\langle m^2 \rangle - \langle |m| \rangle^2) \quad (17)$$

as a function of the annealing parameter, where  $m = (NM)^{-1} \sum_{i,k} s_i^{(k)}$  is the worldline-averaged magnetization.

4. Construct the surrogate annealing schedule from the cumulative susceptibility according to Eqs. (12)–(14).
5. Compare the surrogate schedule with both the linear schedule and the exact Roland-Cerf schedule by solving the time-dependent Schrödinger equation for identical total annealing times.

6. Evaluate the performance using the final ground-state probability

$$P_{\text{GS}} = |\langle \psi_0(T) | \psi(T) \rangle|^2. \quad (18)$$

The essential feature of the proposed approach is that the surrogate schedule requires only the worldline susceptibility obtained from a classical SQA simulation. Exact spectral information is used solely as a reference for evaluating the quality of the surrogate and is never required for constructing the schedule itself.

### III. RESULTS

We now evaluate the proposed worldline-susceptibility schedule and compare it with the linear and Roland–Cerf schedules. The results are presented in stages. We first examine whether the worldline susceptibility obtained from simulated quantum annealing identifies the same critical region as the exact spectral gap. We then compare the resulting annealing schedules and finally investigate their effect on ground-state success probability, together with the physical mechanisms responsible for their differences.

#### A. Worldline Susceptibility as a Surrogate for the Spectral Gap

The proposed scheduling method is based on the hypothesis that the worldline magnetization susceptibility measured during an SQA simulation contains sufficient information to identify the critical region of the annealing process. Unlike the instantaneous spectral gap, which requires exact diagonalization of the quantum Hamiltonian, the susceptibility is obtained directly from the classical Monte Carlo simulation and therefore remains computationally inexpensive even when the spectrum is inaccessible.

To test this hypothesis, Figure 1 compares the normalized worldline susceptibility with the exact instantaneous spectral gap for a representative Sherrington–Kirkpatrick instance with  $n = 10$ . Although the two quantities are physically different, they identify nearly the same critical region. The minimum spectral gap occurs at

$$s_{\text{ED}}^* = 0.883,$$

whereas the susceptibility reaches its maximum at

$$s_{\text{SQ}}^* = 0.920.$$

The difference between these estimates is small compared with the full annealing interval, indicating that the susceptibility provides an accurate estimate of where the evolution should be slowed.

More importantly, the surrogate is not intended to reproduce the detailed functional form of the spectral gap.

Its purpose is to identify the region where non-adiabatic transitions are most likely to occur so that additional annealing time can be allocated there. Since this information is extracted directly from an SQA simulation, the construction completely avoids diagonalization of the quantum Hamiltonian.




FIG. 1. Comparison between the exact instantaneous spectral gap obtained by exact diagonalization and the worldline magnetization susceptibility  $\chi_m$  measured from simulated quantum annealing for a representative  $n = 10$  Sherrington–Kirkpatrick instance. The susceptibility peak closely follows the minimum-gap region, providing a computationally inexpensive estimate of the critical region without requiring spectral information.

Using the worldline susceptibility, the surrogate schedule is constructed according to the procedure described in Sec. II. Figure 2 compares the resulting schedule with the linear schedule and the exact Roland–Cerf schedule.

The three schedules exhibit markedly different strategies for allocating annealing time. The linear schedule evolves at constant speed throughout the anneal. The Roland–Cerf schedule concentrates nearly all of the additional annealing time into a very narrow interval surrounding the minimum spectral gap, where the annealing velocity decreases by almost four orders of magnitude. In contrast, the worldline-susceptibility schedule slows the evolution over a broader critical region while maintaining a considerably smoother time profile. As will be shown in the following subsections, this difference in time allocation plays a central role in the observed annealing performance.




FIG. 2. Comparison of the three annealing schedules. Left: evolution of the annealing parameter  $s(t)$  for the linear, Roland–Cerf, and worldline-susceptibility schedules. Right: corresponding instantaneous annealing velocity  $ds/dt$ . The Roland–Cerf schedule allocates an overwhelming fraction of the annealing time to a narrow region surrounding the minimum spectral gap, whereas the worldline-susceptibility schedule distributes the slowdown more smoothly across the critical region.

## B. Performance of the Surrogate Schedule

We first compare the annealing performance of the three schedules on representative Sherrington–Kirkpatrick instances. The performance is quantified by the final ground-state probability,

$$P_{\text{GS}} = |\langle \psi_0(T) | \psi(T) \rangle|^2, \quad (19)$$

as a function of the total annealing time  $T$ .

Figure 3 shows the results for a representative instance with  $n = 10$ . As expected, the linear schedule exhibits a monotonic increase in ground-state probability with increasing annealing time. The surrogate schedule also improves monotonically, but consistently achieves a higher success probability throughout the entire range of annealing times investigated.

In contrast, the Roland–Cerf schedule exhibits a strikingly different behaviour. Rather than improving monotonically, its performance oscillates as the annealing time increases. Around  $T \approx 20$ , the ground-state probability decreases substantially before recovering at larger annealing times. This behaviour is unexpected because the Roland–Cerf schedule is constructed from the exact instantaneous spectral gap and is therefore theoretically optimal within the local-adiabatic framework.




FIG. 3. Ground-state probability as a function of total annealing time for a representative  $n = 10$  Sherrington–Kirkpatrick instance. The worldline-susceptibility schedule consistently outperforms both the linear and Roland–Cerf schedules over the entire range of annealing times. Unlike the other two schedules, the Roland–Cerf schedule exhibits pronounced non-monotonic behaviour.

A qualitatively different situation is observed for the representative  $n = 12$  instance shown in Fig. 4. Here the Roland–Cerf schedule no longer oscillates. Instead, its performance remains significantly below both the linear and surrogate schedules over the entire annealing-time range. The surrogate schedule again provides the highest ground-state probability.

The two representative examples therefore reveal two distinct modes of behaviour. In the first instance, the

Roland–Cerf schedule suffers from a non-monotonic dependence on annealing time. In the second instance, it remains systematically suppressed even though the total annealing time is increased. In contrast, the worldline-susceptibility schedule behaves robustly in both cases and consistently provides the best overall performance.




FIG. 4. Ground-state probability as a function of total annealing time for a representative  $n = 12$  Sherrington–Kirkpatrick instance. The worldline-susceptibility schedule again achieves the highest success probability, while the Roland–Cerf schedule remains strongly suppressed over the entire annealing-time range.

Table I summarizes the critical-point estimates together with the final ground-state probabilities for the two representative instances. The surrogate schedule consistently produces the highest success probability despite relying only on the worldline susceptibility measured during the SQA simulation rather than the exact instantaneous spectral gap.

TABLE I. Critical-point estimates and final ground-state probabilities for the representative instances.

| System   | $s_{\text{ED}}^*$ | $s_{\text{SQ}}^*$ | Linear | Roland–Cerf | Surrogate    |
|----------|-------------------|-------------------|--------|-------------|--------------|
| $n = 10$ | 0.883             | 0.920             | 0.207  | 0.078       | <b>0.295</b> |
| $n = 12$ | 1.000             | 0.876             | 0.431  | 0.150       | <b>0.517</b> |

These results establish the principal empirical observation of this work: a schedule constructed from a classically measurable worldline susceptibility can outperform a schedule constructed from the exact instantaneous spectral gap. This conclusion is counterintuitive and immediately raises the question of why the theoretically optimal Roland–Cerf schedule performs poorly for these instances. We address this question in the following subsection.

## C. Origin of the Roland–Cerf Failure Modes

The results of the previous subsection present an apparent paradox. The Roland–Cerf schedule is constructed from the exact instantaneous spectral gap and is therefore



expected to provide the most efficient adiabatic evolution. Nevertheless, the surrogate schedule, which is based only on the worldline susceptibility measured during an SQA simulation, achieves a higher ground-state probability for the representative instances studied here.

To understand this behaviour, we analyse the structure of the Roland–Cerf schedule in detail. We find that the observed degradation does not arise from a failure of the local-adiabatic principle itself. Instead, it originates from two distinct mechanisms associated with the distribution of annealing time along the evolution. The first occurs when the minimum spectral gap is located at the end of the annealing path, while the second appears when the schedule concentrates an overwhelming fraction of the annealing time into an extremely narrow interval surrounding an interior minimum gap.

These two mechanisms are illustrated schematically in Fig. 5 through the cumulative time-allocation function

$$F(s) = \frac{\int_0^s \Delta(s')^{-2} ds'}{\int_0^1 \Delta(s')^{-2} ds'}. \quad (20)$$

The function  $F(s)$  measures the fraction of the total annealing time allocated before reaching the annealing parameter  $s$ . A rapid increase of  $F(s)$  therefore indicates that the schedule is spending a disproportionate amount of the available annealing time within a very small interval of the annealing path.




FIG. 5. Cumulative time allocation for the Roland–Cerf schedule. Left: boundary-gap trap for the representative  $n = 12$  instance, where most of the annealing time is allocated near the terminal point  $s = 1$ . Right: interior minimum-gap case for the representative  $n = 10$  instance, where the schedule concentrates the annealing time into a narrow interval surrounding the critical region.

The two mechanisms are discussed separately below.

### 1. A. Boundary-Gap Trap

The first failure mechanism occurs when the minimum spectral gap is located at, or extremely close to, the end of the annealing path. In this situation the Roland–Cerf schedule allocates an overwhelming fraction of the

available annealing time to the final stage of the evolution, where the transverse-field driver has already been almost completely switched off.

This behaviour is illustrated by the representative  $n = 12$  instance shown in the left panel of Fig. 5. The minimum gap occurs at

$$s_{\text{ED}}^* = 1.000,$$

causing the cumulative time-allocation function to rise extremely rapidly near the end of the anneal. Approximately half of the total annealing time is spent after  $s \approx 0.995$ , while nearly 90% of the total runtime is concentrated in the narrow interval immediately preceding the final Hamiltonian.

Although this allocation follows directly from the local-adiabatic condition, it is physically ineffective. Near

$$s = 1,$$

the system Hamiltonian is already dominated by the problem Hamiltonian,

$$H(s) \approx H_P,$$

while the transverse-field driver,

$$(1 - s)H_D,$$

has almost completely vanished. Consequently, the additional evolution time contributes very little to suppressing diabatic transitions because the mechanism responsible for quantum transitions has essentially disappeared.

The poor performance of the Roland–Cerf schedule for this instance therefore does not indicate a breakdown of the local-adiabatic principle itself. Rather, it arises because the schedule faithfully follows a spectral feature whose location is no longer useful for improving the dynamics. In other words, the schedule optimally allocates time according to the instantaneous gap, but the minimum gap occurs in a region where additional annealing time has little physical benefit.

The worldline-susceptibility schedule naturally avoids this pathology. Instead of following the exact minimum gap, the susceptibility reaches its maximum at

$$s_{\text{SQ}}^* = 0.876,$$

well inside the annealing interval. The resulting schedule distributes the slowdown over the broader critical region rather than concentrating almost the entire runtime at the endpoint. This leads to the significantly higher ground-state probability observed in Fig. 4.

### 2. B. Oscillatory Failure

The second failure mechanism is qualitatively different from the boundary-gap trap. Here the minimum spectral

gap lies well inside the annealing interval rather than at its endpoint. One would therefore expect the Roland–Cerf schedule to provide an excellent approximation to the optimal adiabatic evolution. Surprisingly, the opposite is observed.

The representative  $n = 10$  instance shown previously in Fig. 3 exhibits a pronounced non-monotonic dependence of the final ground-state probability on the total annealing time. Instead of approaching the adiabatic limit smoothly, the success probability repeatedly increases and decreases as the annealing time is increased. Table II summarizes several representative values illustrating this behaviour.

TABLE II. Ground-state probability for the representative  $n = 10$  instance. Unlike the linear and surrogate schedules, the Roland–Cerf schedule exhibits pronounced non-monotonic behaviour.

| $T$   | Linear | Roland–Cerf | Surrogate |
|-------|--------|-------------|-----------|
| 10.46 | 0.053  | 0.059       | 0.103     |
| 14.63 | 0.079  | 0.062       | 0.141     |
| 20.45 | 0.112  | 0.038       | 0.181     |
| 28.60 | 0.154  | 0.115       | 0.237     |
| 40.00 | 0.207  | 0.078       | 0.295     |

To investigate this behaviour further, we computed the ground-state probability over a dense grid of annealing times. The resulting evolution is shown in Fig. 6. Rather than a numerical fluctuation, the oscillation forms a well-defined pattern with approximately five extrema over the interval

$$3 \leq T \leq 60.$$

In contrast, both the linear schedule and the worldline-susceptibility schedule evolve smoothly throughout the same interval.




FIG. 6. Dense scan of the total annealing time for the representative  $n = 10$  instance. The Roland–Cerf schedule exhibits a reproducible oscillatory dependence of the ground-state probability, whereas the linear and worldline-susceptibility schedules remain monotonic. The observed oscillation period is far shorter than expected from a simple two-level Landau–Zener–Stückelberg description.

Since such oscillations could potentially arise from numerical inaccuracies, we repeated the calculation using multiple numerical solvers, integration tolerances ranging from  $10^{-8}$  to  $10^{-12}$ , and different discretizations of the annealing schedule. In every case the oscillatory behaviour

remained unchanged within numerical precision, ruling out solver-dependent artifacts.

A natural explanation is coherent Landau–Zener–Stückelberg interference within an effective two-level system. To test this hypothesis, we estimated the oscillation period predicted from the accumulated dynamical phase. The predicted period is approximately

$$T_{\text{LZS}} \approx 312,$$

whereas the observed oscillation period is only

$$T_{\text{obs}} \approx 20.$$

The discrepancy exceeds an order of magnitude, demonstrating that the oscillation cannot be explained by the conventional two-level Landau–Zener–Stückelberg mechanism.

Instead, we attribute the observed behaviour to multi-level quantum interference generated by the extreme time concentration of the Roland–Cerf schedule. As shown in the right panel of Fig. 5, nearly the entire annealing budget is confined to a narrow interval surrounding the minimum gap. Although the gap between the ground and first excited states remains well isolated, the strongly non-uniform evolution enhances interference among multiple excited states during the finite-time dynamics. The resulting interference produces the oscillatory dependence of the final ground-state probability observed in Fig. 6.

The worldline-susceptibility schedule avoids this behaviour because it distributes the slowdown over a wider region of the annealing path. Rather than concentrating almost the entire runtime into a narrow interval, it provides a smoother evolution through the critical region, thereby suppressing the strong multilevel interference responsible for the oscillatory failure.

#### D. Statistical Performance over Disorder Realizations

The representative examples discussed above establish the existence of two distinct failure mechanisms of the Roland–Cerf schedule. An important question, however, is whether these examples are exceptional or whether the same behaviour persists across independent disorder realizations.

To address this question, we average the ground-state probability over independent Sherrington–Kirkpatrick instances generated from the same disorder distribution. Figure 7 shows the disorder-averaged performance across all system sizes  $n \in \{10, 12, 14, 16, 18, 20\}$ . Each panel displays the mean  $P_{\text{GS}}(T) \pm \text{SEM}$  over 20–100 independent instances, with the failure-class composition annotated (B: boundary-gap trap, O: oscillatory, C: clean RC behaviour).

Several important observations emerge. First, the worldline-susceptibility schedule consistently achieves the highest average ground-state probability over the entire annealing-time range at every system size. Second, the Roland-Cerf schedule does not uniformly outperform the linear schedule despite having access to the exact spectral gap, and its relative suppression becomes more pronounced as  $n$  increases due to the growing prevalence of boundary-trap instances. Finally, the averaged curves remain smooth, indicating that the oscillatory behaviour identified in the single-instance analysis contributes generically to the ensemble statistics.




FIG. 7. Average ground-state probability over independent Sherrington-Kirkpatrick instances for all system sizes  $n \in \{10, 12, 14, 16, 18, 20\}$ . Shaded bands show  $\pm 1$  SEM. The worldline-susceptibility schedule (Surrogate) consistently achieves the highest average success probability. The annotations B/O/C count boundary-gap, oscillatory, and clean Roland-Cerf instances in each ensemble.




FIG. 8. Average ground-state probability over twenty independent  $n = 10$  Sherrington-Kirkpatrick instances. The worldline-susceptibility schedule consistently achieves the highest average success probability throughout the annealing-time range. The annotations indicate the number of instances exhibiting the boundary-gap trap, oscillatory failure, and conventional Roland-Cerf behaviour.




FIG. 9. Average ground-state probability over ten independent  $n = 12$  Sherrington-Kirkpatrick instances. The suppression of the Roland-Cerf schedule becomes more pronounced because terminal minimum-gap instances occur more frequently at this system size.

Taken together, Figs. 7–9 demonstrate that the improved performance of the worldline-susceptibility schedule is not limited to isolated examples. Instead, it persists after averaging over disorder realizations and across all system sizes considered in this work.

### E. Classification of Instance-Dependent Failure Modes

The disorder-averaged results establish that the worldline-susceptibility schedule performs better on average than the Roland-Cerf schedule. To understand the origin of this improvement, we now examine the behaviour of individual disorder realizations.

Figure 10 presents six representative Sherrington-Kirkpatrick instances spanning the full range of behaviours observed in this study. Three distinct classes emerge naturally.

The first class consists of instances for which the Roland-Cerf schedule performs as expected, consistently exceeding the performance of the linear schedule and remaining competitive with the surrogate schedule. These correspond to conventional local-adiabatic evolution in which the minimum spectral gap is well separated from the end of the annealing path and no anomalous finite-time dynamics are observed.

The second class exhibits the boundary-gap trap discussed in the previous subsection. In these instances the minimum spectral gap occurs at, or extremely close to, the endpoint of the anneal, causing the Roland-Cerf schedule to allocate an overwhelming fraction of the total annealing time where the transverse-field driver has already become negligible. As a result, the Roland-Cerf schedule performs significantly worse than both the linear and surrogate schedules.



The third class is characterized by oscillatory behaviour of the ground-state probability as the total annealing time is increased. These instances display the multilevel interference effects discussed previously and again show a clear advantage for the worldline-susceptibility schedule.




FIG. 10. Ground-state probability for six representative Sherrington–Kirkpatrick instances. The instances naturally separate into three categories: conventional Roland–Cerf behaviour, the boundary-gap trap, and oscillatory failure. The worldline-susceptibility schedule remains robust across all three classes.

To quantify this classification, Fig. 11 summarizes the performance of each instance as a function of its minimum spectral gap. The left panel shows the pairwise comparison between the three schedules, while the right panel plots the fraction of annealing times for which the surrogate schedule outperforms the Roland–Cerf schedule against the exact minimum gap.

A clear separation between the different classes is observed. Instances with relatively large minimum gaps are generally well described by the Roland–Cerf schedule, whereas instances with very small minimum gaps are much more susceptible to the two failure mechanisms identified above. The transition between these regimes occurs over approximately one order of magnitude in the minimum spectral gap.

Although the present dataset is too small to establish a universal threshold, the observed correlation strongly suggests that the minimum spectral gap provides a useful indicator of when precise local-adiabatic scheduling becomes vulnerable to finite-time pathologies. In contrast, the worldline-susceptibility schedule remains comparatively insensitive to this distinction, exhibiting stable performance across the entire range of instances considered here.

Taken together, Figs. 10 and 11 show that the superior performance of the worldline-susceptibility schedule is not associated with a particular disorder realization. Instead, it arises from its robustness against multiple classes of finite-time dynamical behaviour that affect the exact local-adiabatic schedule.

Figure 12 shows the failure-class fractions as a function of




FIG. 11. Classification of the representative instances. Left: pairwise comparison of the three schedules for each realization. Right: fraction of annealing times for which the worldline-susceptibility schedule outperforms the Roland–Cerf schedule as a function of the minimum spectral gap. The different failure mechanisms occupy distinct regions of the parameter space.

system size, together with Wilson 95% confidence intervals. The boundary-gap fraction increases monotonically with  $n$ , while the oscillatory fraction remains broadly constant, consistent with the interpretation that boundary trapping becomes the dominant RC failure mode at larger system sizes.




FIG. 12. Fraction of instances in each failure class as a function of system size  $n$ . Error bars are Wilson 95% confidence intervals. The boundary-gap trap (B) becomes increasingly prevalent with  $n$ , while the oscillatory failure (O) fraction is roughly constant.

## F. Robustness of the Oscillatory Failure under Decoherence

The oscillatory behaviour identified in the previous subsection was obtained from unitary Schrödinger evolution. Since physical quantum annealers inevitably operate in the presence of environmental decoherence, it is important to determine whether this oscillatory behaviour survives under open-system dynamics or is merely an artifact of the idealized closed-system description.

To address this question, we repeated the simulations

using a Lindblad master equation,

$$\dot{\rho} = -i[H(s(t)), \rho] + \sum_k \gamma \mathcal{D}[\sigma_z^{(k)}]\rho, \quad (21)$$

where  $\gamma$  denotes the dephasing rate and

$$\mathcal{D}[L]\rho = L\rho L^\dagger - \frac{1}{2}\{L^\dagger L, \rho\}$$

is the standard Lindblad dissipator. The calculations were performed for the same representative instance that exhibited oscillatory behaviour under unitary evolution.

Figure 13 shows the resulting ground-state probability for several dephasing strengths. The coherent oscillation survives essentially unchanged for weak and intermediate decoherence, remaining clearly visible up to

$$\gamma = 10^{-2}.$$

Only for stronger dephasing,

$$\gamma = 5 \times 10^{-2},$$

does the oscillatory structure become strongly suppressed, approaching the monotonic behaviour expected for incoherent evolution.




FIG. 13. Ground-state probability under Lindblad dephasing. Left: evolution for increasing dephasing strength. The oscillatory behaviour survives weak and intermediate decoherence and disappears only for sufficiently strong dephasing. Right: deviation from the unitary result as a function of annealing time.

The observed crossover is consistent with the analytical estimate obtained from the corresponding two-level dephasing model, which predicts suppression of coherent oscillations once the accumulated dephasing becomes comparable to the coherent phase accumulation. Although the simple two-level model fails to explain the oscillation period itself, it correctly estimates the decoherence scale at which the interference pattern disappears.

These results therefore support the interpretation that the oscillatory failure is a genuine coherent dynamical phenomenon rather than a numerical artifact. Moreover, they suggest that the mechanism should remain observable in quantum annealers operating in the weak-to-moderate decoherence regime.

## G. Robustness with Respect to the SQA Parameters

The proposed scheduling method is constructed from quantities measured during an SQA simulation. It is therefore important to determine whether the resulting surrogate depends sensitively on the Trotter discretization or on the simulation temperature. If the estimated critical point varied strongly with these numerical parameters, the practical usefulness of the surrogate would be significantly reduced.

To investigate this issue, we performed SQA simulations over a range of Trotter numbers,

$$M = \{8, 16, 32, 64\},$$

and inverse temperatures,

$$\beta = \{2, 5, 10, 20\},$$

and compared the resulting estimate of the critical point with the value obtained from exact diagonalization.

Figure 14 summarizes the average deviation

$$|s_{ED}^* - s_{SQ}^*|$$

over five independent Sherrington–Kirkpatrick instances. The error remains remarkably uniform across the entire  $(\beta, M)$  parameter space. No systematic improvement is observed upon increasing either the Trotter number or the inverse temperature, and the variation between different parameter choices remains considerably smaller than the fluctuations arising from different disorder realizations.




FIG. 14. Average difference between the exact critical point and the SQA estimate over the  $(\beta, M)$  parameter space. The surrogate shows little systematic dependence on either the inverse temperature or the number of Trotter slices.

The same conclusion is evident from the line cuts shown in Fig. 15. At fixed inverse temperature, increasing the number of Trotter slices produces only minor changes in the

estimated critical point. Likewise, at fixed Trotter number, varying the inverse temperature does not introduce any systematic trend. Instead, the dominant uncertainty originates from the intrinsic instance-to-instance variation of the Sherrington–Kirkpatrick ensemble.




FIG. 15. Dependence of the surrogate critical point on the Trotter number (left) and inverse temperature (right). The shaded regions denote the variation over different disorder realizations. The instance-to-instance fluctuations exceed the systematic dependence on the numerical SQA parameters.

These results demonstrate that the proposed surrogate schedule is numerically robust. Once the SQA simulation has reached a moderate Trotter resolution and inverse temperature, the location of the critical region becomes essentially independent of the simulation parameters. Consequently, the dominant source of uncertainty is the physical variation between problem instances rather than the numerical approximation used to construct the surrogate. This robustness is an important practical advantage, since it implies that reliable schedules can be obtained without fine tuning the SQA parameters.

#### IV. CONCLUSION

We have proposed a practical surrogate for local-adiabatic quantum annealing based on the worldline magnetization susceptibility  $\chi_m = NM(\langle m^2 \rangle - \langle |m| \rangle^2)$  measured during simulated quantum annealing. Unlike the Roland–Cerf schedule, which requires complete knowledge of the instantaneous spectral gap, the proposed approach relies only on quantities obtained directly from an SQA simulation and therefore avoids repeated diagonalization of the quantum Hamiltonian.

Using exact diagonalization as ground truth for Sherrington–Kirkpatrick spin-glass instances across system sizes  $n \in \{10, 12, 14, 16, 18, 20\}$ , we showed that the worldline susceptibility consistently identifies the critical region of the annealing process and produces schedules that outperform the conventional linear schedule. More importantly, the surrogate schedule also outperforms the exact local-adiabatic schedule on a substantial fraction of problem instances.

The origin of this unexpected behaviour was investigated in detail. We identified two distinct failure mechanisms of the exact Roland–Cerf schedule. The first is a boundary-gap trap, in which the minimum spectral gap occurs at the end of the annealing path, causing the algorithm to allocate most of the annealing time after the transverse-field driver has effectively vanished. The second is a coherent oscillatory failure associated with highly non-uniform time allocation. Numerical convergence tests, dense annealing-time scans, and Lindblad simulations demonstrate that this behaviour is a genuine physical effect rather than a numerical artifact.

We further showed that the surrogate schedule remains robust with respect to the numerical parameters of the SQA simulation, including the Trotter discretization and inverse temperature. The dominant uncertainty arises from the physical variation between problem instances rather than from the surrogate construction itself.

Taken together, these results suggest that the objective of schedule design is not necessarily to reproduce the exact spectral gap as accurately as possible. Instead, a computationally inexpensive surrogate that reliably identifies the critical region while avoiding over-optimization of the annealing trajectory can provide a more robust strategy for practical quantum annealing. This observation opens the possibility of designing scalable annealing schedules directly from inexpensive classical simulations, without requiring explicit knowledge of the quantum spectrum.

#### Acknowledgments

We thank the IMSc high-performance computing facility. [Funding and further acknowledgments to be added.]

[1] T. Kadowaki and H. Nishimori, Phys. Rev. E **58**, 5355 (1998).  
[2] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, Science **292**, 472 (2001).  
[3] G. E. Santoro, R. Martoňák, E. Tosatti, and R. Car, Science **295**, 2427 (2002).

[4] A. Lucas, Front. Phys. **2**, 5 (2014).  
[5] T. Albash and D. A. Lidar, Rev. Mod. Phys. **90**, 015002 (2018).  
[6] J. Roland and N. J. Cerf, Phys. Rev. A **65**, 042308 (2002).  
[7] S. Jansen, M.-B. Ruskai, and R. Seiler, J. Math. Phys. **48**, 102111 (2007).