Quantum-inspired dynamical models on quantum and classical annealers

Philipp Hanussek[1,2,*] Jakub Pawłowski[3,4,†] Zakaria Mzaouali[5,6,‡] and Bartłomiej Gardas[1,§]

[1] *Institute of Theoretical and Applied Informatics, Polish Academy of Sciences, Bałtycka 5, 44-100 Gliwice, Poland*

[2] *Faculté des Sciences et Ingénierie, Sorbonne Université, 75005 Paris, France*

[3] *Institute of Theoretical Physics, Faculty of Fundamental Problems of Technology, Wrocław University of Science and Technology, 50-370 Wrocław, Poland*

[4] *Quantumz.io Sp. z o.o., Puławska 12/3, 02-566 Warsaw, Poland*

[5] *Jülich Supercomputing Centre, Institute for Advanced Simulation, Forschungszentrum Jülich, Wilhelm-Johnen-Straße, 52428 Jülich, Germany*

[6] *Institut für Theoretische Physik, Universität Tübingen, Auf der Morgenstelle 14, 72076 Tübingen, Germany*

(Received 4 September 2025; revised 2 February 2026; accepted 9 March 2026; published 21 April 2026)

We propose a practical, physics-inspired benchmarking suite to challenge both quantum and classical computers by mapping real-time quantum dynamics to a common optimization format. Using a parallel-in-time encoding, we convert the real-time propagator of an n-qubit, possibly non-Hermitian, Hamiltonian into quadratic unconstrained binary optimization (QUBO) instances that are executable in a solver-agnostic manner on quantum annealers and classical optimizers alike. This enables direct, like-for-like performance comparisons across fundamentally different computational paradigms. To stress-test the framework, we consider eight representative dynamical models spanning single-qubit rotations, multiqubit entangling gates (Bell, Greenberger-Horne-Zeilinger, cluster), and $\mathcal{PT}$-symmetric and other non-Hermitian generators, and evaluate success probability and time to solution as standard benchmarking metrics. Applying this method to two generations of D-Wave quantum annealers and to state-of-the-art classical solvers (simulated annealing and the graphics-processing unit–accelerated VeloxQ), we find that Advantage2 consistently outperforms its predecessor, while VeloxQ retains the shortest absolute runtimes, reflecting the maturity of classical heuristics. We further extend the benchmarks to large-scale instances ($N \approx 10^5$), establishing a demanding classical baseline for future hardware. Together, these results position the parallel-in-time QUBO framework as a versatile and physically motivated test bed for quantitatively tracking progress toward quantum-competitive simulation of dynamical systems.

DOI: 10.1103/rrf3-jm5m

I. INTRODUCTION

Simulating the real-time evolution of many-body quantum systems is central to condensed matter physics [1–3], quantum chemistry [4–6], and emerging quantum technologies [7]. However, the exponential growth of the Hilbert space makes classical simulations infeasible beyond a few dozen qubits. This limitation motivated the idea that quantum hardware could outperform classical methods in simulating quantum dynamics [8–10]. Quantum annealers, such as those developed by D-Wave Systems, offer a platform for exploring such advantages. Recent experiments have suggested that they can simulate quantum dynamics beyond the capabilities of classical devices [11,12], but these results are still a subject of debate [13–15].

In the current noisy intermediate-scale quantum era, decisive demonstrations of advantage have proved elusive [17–22]. Gate-based processors suffer from limited qubit counts and decoherence, while analog platforms, such as ultracold atoms, trapped ions, and photonic devices, excel at modeling specific Hamiltonians but lack flexibility [23–25]. Quantum annealers, produced most prominently by D-Wave Systems, occupy a specific niche. Their main task is the minimization of quadratic unconstrained binary optimization (QUBO) cost functions on fabricated Ising graphs, offering thousands of physical qubits that operate

*Contact author: philipp.hanussek@etu.sorbonne-universite.fr

†Contact author: jakub.pawlowski@pwr.edu.pl

‡Contact author: z.mzaouali@extern.fz-juelich.de

§Contact author: bgardas@iitis.pl

continuously rather than in discrete gate cycles. Because the hardware is already available as a cloud service, annealers are attractive test beds for near-term advantage claims [11].

Yet two technical hurdles have so far blocked a rigorous assessment of annealers for dynamical simulation [26]. First, mapping Schrödinger's equation—fundamentally a sequence of unitary operators—onto a static QUBO cost function has lacked an efficient, scalable recipe. Second, without agreed-upon benchmark suites it is impossible to separate genuine quantum speedups from artifacts caused by, for example, a particularly easy instance or a fortuitous embedding on the hardware graph. Consequently, previous performance studies have either relied on synthetic optimization problems that are only indirectly related to real-time quantum mechanics or compared quantum and classical solvers on disjoint workloads.

Our work contributes to closing this gap via a rigorous benchmarking protocol built upon a parallel-in-time encoding of quantum dynamics (see Fig. 1). The same family of instances is executed (1) on two generations of D-Wave quantum annealers—Advantage (Pegasus topology) and the newer Advantage2 (Zephyr topology with 33% higher connectivity)—(2) on VeloxQ, a graphics-processing unit (GPU)-accelerated classical heuristic that solves the identical QUBOs but without hardware-embedding constraints [16], and (3) with use of both CPU and GPU implementations of simulated annealing, which serves as the well-known point of reference for classical solvers. This like-for-like design removes confounding factors and exposes the scaling behavior that ultimately determines whether and when quantum advantage can emerge. The empirical picture is mixed but instructive. D-Wave Advantage2 already delivers an order of magnitude higher ground-state success probability and a smaller time-to-solution (T_S) exponent compared with its predecessor, signaling rapid hardware progress. Nevertheless, VeloxQ remains the overall leader on the problem sizes accessible today by quantum hardware, underscoring the strength of optimized classical algorithms and the need

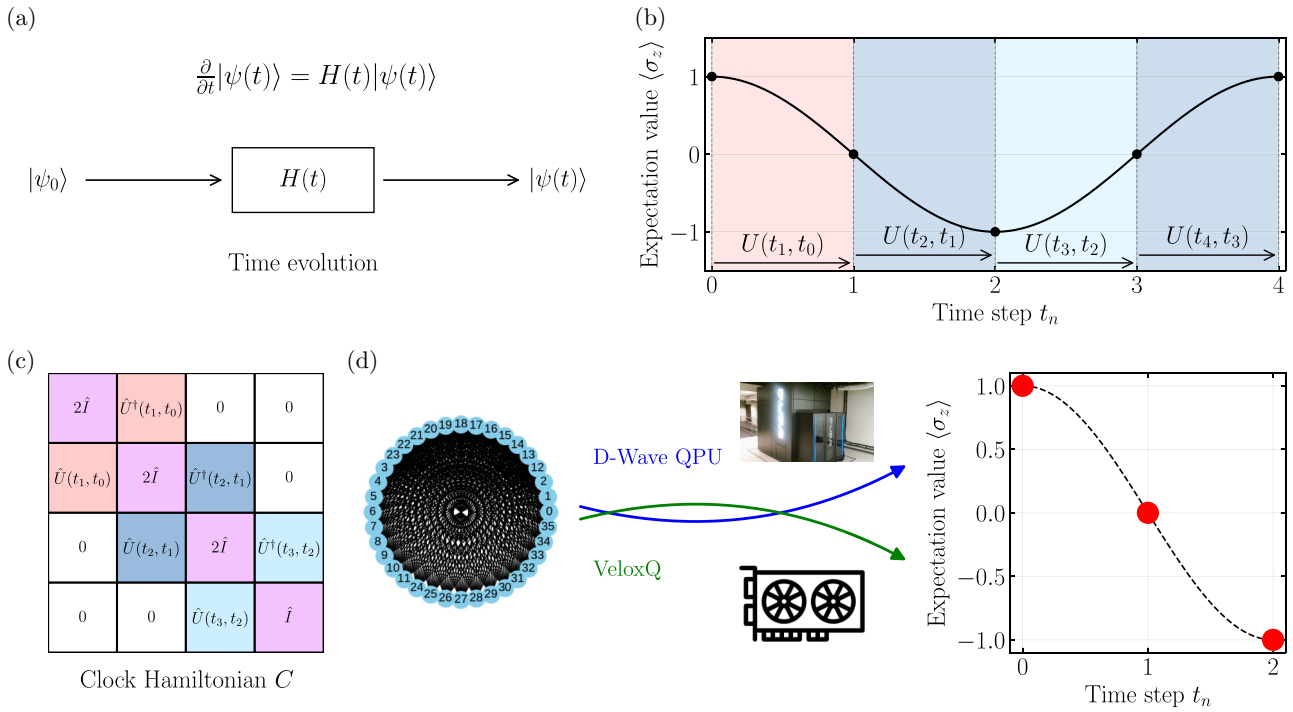

FIG. 1. Overview of the workflow. (a) A quantum state $|\psi_0\rangle$ of dimension N evolves under the (possibly non-Hermitian) generator $H(t)$ according to the Schrödinger equation $\partial_t|\psi(t)\rangle = H(t)|\psi(t)\rangle$. (b) The real-time interval $[t_0, t_f]$ is discretized into N equal Trotter slices of width Δt; each slice is propagated by the short-time operator $U(t_{n+1}, t_n) = \exp[H(t_n)\Delta t]$, so the sequence $\{U(t_{n+1}, t_n)\}_{n=0}^{N-2}$ fully specifies the history. (c) Discretized dynamics are encoded in the block-tridiagonal clock Hamiltonian C, Eq. (6); appending the initial-state projector yields the positive-definite matrix A. Writing the quadratic form $\frac{1}{2}\langle x|A|x\rangle - \langle x|\phi\rangle$ in fixed-point binary variables (R bits over the range $[-2^D, 2^D]$) produces a QUBO with $\mathcal{O}(LNR)$ spins, Eq. (10). (d) The QUBO instance is submitted to a D-Wave quantum annealer (Advantage and Advantage2) and to the GPU-accelerated classical heuristic VeloxQ [16]. Both solvers return low-energy spin configurations that are decoded unambiguously into the time-ordered amplitudes $\{|\psi(t_n)\rangle\}_{n=0}^{N-1}$; from these, physical observables—e.g., $\langle\sigma_z\rangle(t)$ shown on the right—are reconstructed without further fitting. Both solvers produce the same observable dynamics. This pipeline therefore maps continuous-time quantum dynamics to a solver-agnostic combinatorial problem, enabling a like-for-like performance comparison between quantum and classical optimization hardware.

for denser qubit connectivity or error-mitigation strategies before a categorical quantum advantage can be claimed. Through this work, we provide an open-source practical benchmark suite that spans single-qubit rotations, entangling gates, and $\mathcal{PT}$-symmetric non-Hermitian generators to supply the community with a transparent yardstick against which future hardware and algorithmic advances can be measured.

While our QUBO encoding builds on the parallel-in-time construction introduced in Ref. [27], the present work addresses a fundamentally different question. Rather than demonstrating feasibility on early-generation hardware, we establish a solver-agnostic, physics-derived benchmarking framework for quantum dynamics and use it to quantitatively assess the current performance gap between modern quantum annealers and state-of-the-art classical heuristics. In contrast to Ref. [27], we consider non-Hermitian and $\mathcal{PT}$-symmetric generators, systematically distinguish native and non-native hardware embeddings, and benchmark two generations of D-Wave processors against GPU-accelerated classical solvers. Unlike Ref. [16], which focuses on synthetic or optimization-driven QUBO instances, our benchmarks are derived directly from continuous quantum dynamics, providing a physically motivated and application-relevant test bed for tracking hardware and algorithmic progress.

II. METHOD

Before we formalize the benchmark problems and outline the protocol used to generate each instance, it is essential to introduce the computational solvers that underpin all subsequent experiments. Accordingly, the following two subsections are dedicated to an overview of every solver used throughout this study, establishing a common performance baseline and clarifying the methodological context in which our results should be interpreted.

A. Quantum annealing

Quantum annealing begins by initializing the system in the ground state of a simple driver Hamiltonian H_D whose ground state is the uniform superposition $|+\rangle^{\otimes N} = (1/\sqrt{2^N}) \sum_{x\in\{0,1\}^N} |x\rangle$, with N being the number of qubits. The goal is to find the ground state of a problem Hamiltonian

$$H_P = \sum_{i=1}^{N} h_i \hat{\sigma}_z^{(i)} + \sum_{i<j} J_{ij} \hat{\sigma}_z^{(i)} \hat{\sigma}_z^{(j)}, \tag{1}$$

where $\hat{\sigma}_z^{(i)}$ is the Pauli matrix for qubit q_i, h_i is the qubit bias, and $J_{i,j}$ is the strength of coupling between qubits. The ground state of H_P encodes the solution to an Ising optimization problem. During the annealing process, the system evolves under a time-dependent Hamiltonian

$$H_{\text{Ising}} = \frac{A(s)}{2} H_D + \frac{B(s)}{2} H_P \tag{2}$$

governed by quantum processing unit (QPU)-dependent annealing schedules $A(s)$ and $B(s)$ [see Fig. 2(a)]. If the evolution is sufficiently slow, according to the adiabatic theorem [28], the system remains in its instantaneous ground state, and the final state gives the solution to the original problem. In addition to standard forward annealing, we use cyclic annealing as a refinement strategy for

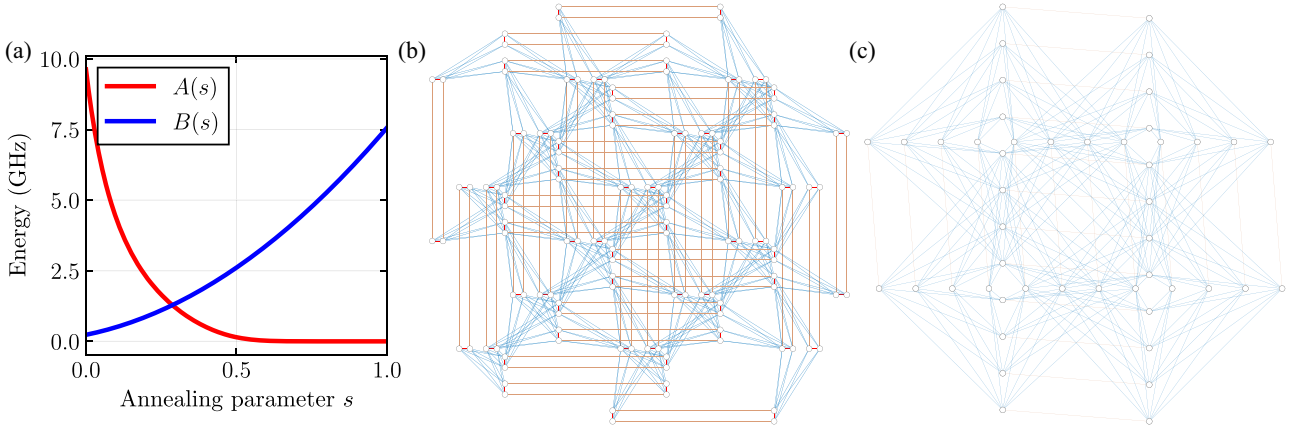

FIG. 2. (a) Time-dependent energy scales executed by the D-Wave quantum annealer. The coefficient $A(s)$ (red) multiplies the transverse-field term that enables quantum tunneling and therefore falls rapidly as the run progresses. The coefficient $B(s)$ (blue) multiplies the problem Hamiltonian and rises smoothly, so the device moves from a tunneling-dominated regime at $s = 0$ to a problem-dominated regime at $s = 1$. (b) Native connectivity map for the Pegasus architecture (D-Wave Advantage). Circles mark individual superconducting qubits, blue lines are programmable couplers that lie within the main module, and orange lines link neighboring modules. (c) Connectivity map for the newer Zephyr architecture (D-Wave Advantage2), which increases the number of couplers per qubit and introduces longer-range links.

selected QUBO instances [29–31]. The key idea is to repeatedly "soften" the optimization landscape around a candidate solution and then "refreeze" it. Operationally, the problem Hamiltonian is kept fixed, while the annealing control parameter s is driven through a short back-and-forth schedule: starting from a classical spin configuration at $s = 1$ (problem-dominated regime), the schedule is partially reversed to an intermediate point $s_p < 1$ to temporarily restore the transverse-field contribution, and is then returned to $s = 1$, where the state is measured. By repeating this cycle, the solver can leave shallow local minima, while still exploiting information contained in the current best configuration [32–34].

A single cycle can be described by a piecewise-linear control schedule $s(t)$,

$$s(t) = \begin{cases} 1 - (1 - s_p)\dfrac{t}{\tau_\downarrow}, & 0 \le t < \tau_\downarrow, \\ s_p, & \tau_\downarrow \le t < \tau_\downarrow + \tau_p, \\ s_p + (1 - s_p)\dfrac{t - \tau_\downarrow - \tau_p}{\tau_\uparrow}, & \tau_\downarrow + \tau_p \le t \\ & \le \tau_\downarrow + \tau_p + \tau_\uparrow, \end{cases} \tag{3}$$

where $\tau_\downarrow$ and $\tau_\uparrow$ set the reverse and forward ramp durations and τ_p is an optional pause. The turnaround point s_p controls the locality of the search: larger s_p keeps the dynamics close to the input configuration (local refinement), whereas smaller s_p induces broader exploration. We execute N_{cyc} cycles and keep the lowest objective value observed over the entire cycle sequence. In practice, cyclic annealing often yields rapid early improvement followed by diminishing returns (see the plateau behavior in Fig. 3), so the cycle budget is best treated as a tunable trade-off between solution quality and run-time.

D-Wave's superconducting quantum annealers are now applied across a spectrum of physics-driven tasks. Experiments on a 5000-qubit processor reproduced quantum critical dynamics in three-dimensional spin glasses, extracting critical exponents beyond the reach of classical simulation [35]. In molecular science, an annealer-based eigensolver has recovered excited-state spectra for small molecules, showing that electronic structure problems can be recast as Ising models amenable to quantum annealing [36]. Hybrid quantum-classical workflows have also tackled real-time traffic flow optimization, demonstrating measurable congestion reduction on realistic road networks [37–42]. In computer science, an adiabatic implementation of Simon's algorithm was recently proposed and solved on D-Wave quantum annealers [43]. Furthermore, the D-Wave machine has been used as a platform for thermodynamic experiments to validate quantum fluctuation theorems and simulate quantum thermal machines

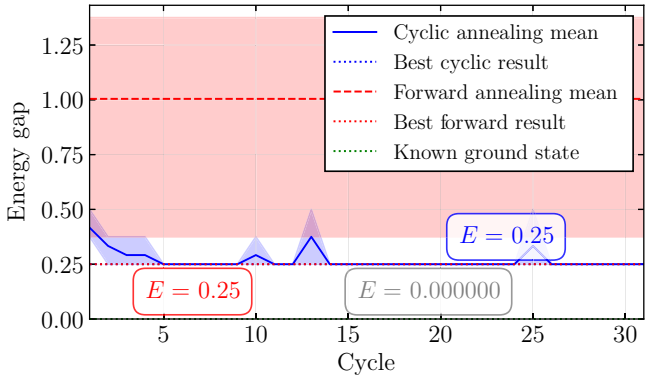

FIG. 3. Exemplary plateau behavior observed in cyclic annealing for instance H_3 with three time points. The blue line indicates the mean over the minimum energy found over each iteration, while the dashed red line shows the mean of the energies found by standard forward annealing. The shadowed box indicates the lower and upper quartiles for all samples for forward annealing, and for lowest cycle energy for cyclic annealing. For this plot, we performed three cyclic annealing runs over 30 cycles and 30 forward annealing runs on Advantage2 system 1.10.

[44–47]. Progress notwithstanding, widespread deployment is limited by sparse qubit connectivity, which imposes costly minor embeddings, and by analog control errors and thermal noise that dilute prospective speedups [48–51].

Because of the limited qubit connectivity in current D-Wave QPUs, problem structures must conform to the device's *working graph*. The Advantage and Advantage2 QPUs analyzed in this work implement distinct hardware topologies: Pegasus for Advantage and Zephyr for Advantage2, as shown in Figs. 2(b) and 2(c), respectively. The Zephyr graph increases native qubit connectivity, with 20 connections per qubit compared with 15 in Pegasus. When a problem instance does not natively fit the hardware graph, it must be embedded via a minor embedding procedure, which may introduce additional complexity. In this study, such embeddings are performed with D-Wave's heuristic *minorminer* tool when necessary.

B. VeloxQ

VeloxQ is a fast and efficient, physics-inspired solver based on the theory of dynamical systems, designed for solving QUBO problems, which are crucial for various real-world optimization tasks [52]. In recent work, it was extensively benchmarked against many state-of-the-art QUBO solvers, D-Wave quantum annealers among them [16]. Unlike them (or other physics-inspired optimization approaches such as gate-based quantum computers), VeloxQ does not rely on technological advancements in hardware to reach its full potential, as it leverages conventional computing technology (such as GPUs). While it is a classical solver, it emulates the probabilistic nature of

quantum annealers by producing a low-energy spectrum of solutions. By default, the solver operates with a carefully selected set of parameters that include instance-dependent tuning. However, it also offers flexibility, allowing users to adjust parameters to control the trade-off between solution quality and run-time on the basis of problem knowledge or prior runs. Crucially, it offers stable and repeatable results, which do not suffer from unstable hardware. VeloxQ distinguishes itself from D-Wave quantum annealers in several significant ways. A key difference lies in their handling of problem topology: D-Wave's QPUs have a fixed spin connectivity structure, and necessitate "minor embedding" for most real-world problems due to their specific topologies, which introduces substantial overhead. In contrast, VeloxQ natively handles all problem topologies without requiring such embedding. Furthermore, VeloxQ exhibits superior scalability, and is capable of tackling large-scale optimization problems with up to 2×10^8 sparsely connected variables, a feat currently intractable for D-Wave and its hybrid solutions, with estimates suggesting it could be decades before quantum annealers could reach similar capabilities [16].

C. Problem generation

We briefly describe a method introduced in Ref. [27] that converts the NP-hard problem of solving dynamical systems into finding the ground state of a QUBO formulation, and thus a problem formulation that is suitable for execution on quantum annealers. Consider an L-dimensional quantum system evolving under a (possibly non-Hermitian) generator $H(t) \in \mathbb{C}^{L\times L}$, governed by the first-order differential equation

$$\frac{d}{dt}\left|\psi(t)\right\rangle = H(t)\left|\psi(t)\right\rangle, \quad \left|\psi(t_0)\right\rangle = \left|\phi\right\rangle. \tag{4}$$

Upon discretizing the time interval $[t_0, t_f]$ into N steps of size Δt, one defines short-time propagators $U_n = \exp[H(t_n)\Delta t]$. The full evolution is encoded in a so-called history state,

$$\left|\Psi\right\rangle = \sum_{n=0}^{N-1} \left|t_n\right\rangle \otimes \left|\psi(t_n)\right\rangle. \tag{5}$$

This state is annihilated by the Hermitian clock operator

$$\mathcal{C} = \sum_{n=0}^{N-2} \left(\left|t_{n+1}\right\rangle\left\langle t_{n+1}\right| \otimes I - \left|t_{n+1}\right\rangle\left\langle t_n\right| \otimes U_n + \text{H.c.}\right), \tag{6}$$

i.e., $\mathcal{C}\left|\Psi\right\rangle = 0$. To enforce the initial condition $\left|\psi(t_0)\right\rangle = \left|\phi\right\rangle$, the following operator is introduced:

$$\mathcal{A} := \mathcal{C} + \left|t_0\right\rangle\left\langle t_0\right| \otimes I. \tag{7}$$

If the generator $H(t)$ is Hermitian, then $\mathcal{A}$ is positive definite, and the solution is given by the minimization of the quadratic function

$$f(\mathbf{x}) = \frac{1}{2}\left\langle \mathbf{x}\right|\mathcal{A}\left|\mathbf{x}\right\rangle - \left\langle \mathbf{x}\right|\left.\phi\right\rangle. \tag{8}$$

The real variables x_i must be converted into binary form so the optimization can be executed on a quantum annealer (or any other QUBO solver). Each component is written in fixed-point notation

$$x_i = 2^D\left(\sum_{\alpha=0}^{R-1} 2^{-\alpha} q_i^\alpha - 1\right), \quad q_i^\alpha \in \{0,1\}, \tag{9}$$

where R bits of precision cover the interval $[-2^D, 2^D]$. Substitution of this expression into the quadratic objective $f(\mathbf{x})$, yields a QUBO form,

$$f(\mathbf{q}) = \sum_{i,\alpha} a_i^\alpha q_i^\alpha + \sum_{i,j,\alpha,\beta} b_{ij}^{\alpha\beta} q_i^\alpha q_j^\beta + f_0, \tag{10}$$

with explicitly computable coefficients,

$$\begin{aligned} b_{ij}^{\alpha\beta} &= A_{ij} 2^{1-\alpha-\beta+2D}, \\ a_i^\alpha &= 2^{1-\alpha+D}\left[A_{ii} - 2^D\left(\sum_j A_{ij} - \varphi_i\right)\right], \\ f_0 &= 2^D\left(2^{D-1}\sum_{ij} A_{ij} + \sum_i \varphi_i\right), \end{aligned} \tag{11}$$

where A_{ij} are elements of the symmetric coefficient matrix $\mathcal{A}$ and quantify the quadratic coupling between variables x_i and x_j, and φ_i is a component of the linear-term vector and acts as an external bias on the variable x_i. The constant shift f_0 does not influence the minimizer and can be dropped. The edge set mirrors the nonzero pattern of the matrix A derived from the discretized dynamics. This binary quadratic form is the instance that is finally programmed into the quantum annealer and VeloxQ.

D. Problem instances

We examine entangled and nonentangled quantum systems, as well as non-Hermitian Hamiltonians that model quantum dynamics such as $\mathcal{PT}$-symmetric qubits. To perform an accurate performance comparison between classical and quantum solvers, we select dynamical systems that can be precisely modeled in QUBO form, as defined by eq. (10) even with limited precision. Some of these dynamical systems produce QUBO instances that are native to the Pegasus graph, others require an embedding.

1. System 1: Single-qubit Y rotation

The Hamiltonian is

$$H_1 = \frac{\pi}{2}\sigma_y, \tag{12}$$

so $U(t) = e^{-iH_1 t}$ enacts a rigid rotation of the Bloch vector about the y axis by an angle $\theta(t) = (\pi/2)t$. Eigenvalues $\pm\frac{\pi}{2}$ give a fundamental period $T = 2$ (in the dimensionless unit of this paper), making this an analytically solvable sanity-check instance for the annealer, which is also native to the D-Wave Advantage (Pegasus graph) architecture.

2. System 2: Oblique single-qubit drive

The Hamiltonian is

$$H_2 = \frac{\pi}{\sqrt{2}}(\sigma_x + \sigma_z). \tag{13}$$

The generator points along the $(1, 0, 1)$ direction of the Bloch sphere. Real eigenvalues $\pm\pi/\sqrt{2}$ yield unitary dynamics that mix computational basis states with Rabi frequency $\Omega = \pi/\sqrt{2}$, testing the algorithm's ability to resolve nonorthogonal control axes. Similarly to Eq. (12), this generator is native to the D-Wave Advantage (Pegasus graph) architecture.

3. System 3: Single-qubit Rabi oscillations (Y-driven)

The Hamiltonian is

$$H_3 = \pi\left(\frac{3}{4}\mathbb{1} - \frac{1}{4}\sigma_y\right). \tag{14}$$

This Hamiltonian can be interpreted as a combination of a global phase rotation plus a spin rotation around the y axis. The off-diagonal terms generate coherent oscillations (Rabi oscillations) between the qubit basis states. The resulting QUBO is not native to the D-Wave Advantage and Advantage2 (Pegasus and Zephyr graphs, respectively) architectures.

4. System 4: Native two-qubit entanglement

The Hamiltonian is

$$H_4 = \frac{\pi}{4}(\sigma_x \otimes \sigma_x + \sigma_y \otimes \sigma_y). \tag{15}$$

This Hamiltonian generates maximally entangled Bell states under time evolution. Crucially, its QUBO encoding maps natively to the Pegasus topology of D-Wave Advantage processors, eliminating embedding overhead and providing a direct test of entanglement dynamics on quantum annealers.

5. System 5: Two-qubit pair flip

The Hamiltonian is

$$H_5 = \frac{\pi}{4}(\sigma_x \otimes \sigma_x - \sigma_y \otimes \sigma_y). \tag{16}$$

Restricted to the $\{|00\rangle, |11\rangle\}$ subspace, the spectrum is $\{\pm 1\}$ and the dynamics perform Rabi oscillations between these product states. At half-period the system passes through the Bell state, providing a clean entanglement-generation benchmark and Pegasus-native problem for the D-Wave Advantage solver.

6. System 6: Greenberger-Horne-Zeilinger flip on three qubits

The Hamiltonian is

$$H_6 = \frac{\pi}{8}(\sigma_x \otimes \sigma_x \otimes \sigma_x - \sigma_x \otimes \sigma_y \otimes \sigma_y - \sigma_y \otimes \sigma_x \otimes \sigma_y - \sigma_y \otimes \sigma_y \otimes \sigma_x). \tag{17}$$

This is the natural three-qubit analog of H_5, Eq. (16). Evolution swaps $|000\rangle$ and $|111\rangle$ with period $T = 2$ while creating genuine tripartite (Greenberger-Horne-Zeilinger) entanglement that is also Pegasus-native for the D-Wave Advantage device.

7. System 7: $\mathcal{PT}$-symmetric qubits

Besides Hermitian operators, which are usually used as observables in quantum theory, quasi-Hermitian Hamiltonians allow one to model unitary evolution as well. We simulate the behavior of a specific type of quasi-Hermitian system, the so-called $\mathcal{PT}$-symmetric Hamiltonians. $\mathcal{PT}$-symmetric systems remain invariant under the combined parity ($\mathcal{P}$) and time-reversal ($\mathcal{T}$) operators. Both these transformations are Hermitian and they are independent of each other, with $[\mathcal{P}, \mathcal{T}] = 0$ [53]. Notably, systems with *unbroken* $\mathcal{PT}$ symmetry have a real eigenspectrum and exhibit quasi-Hermitian behavior. A system is unbroken if the eigenstates of H correspond to the eigenstates of the $\mathcal{PT}$ operator. We evaluate the dynamics of the non-Hermitian Hamiltonian

$$H_7 = \omega\left(b\sigma_x + \sin(\alpha)\sigma_y\right), \quad \omega, \alpha, b \in \mathbb{R}, \tag{18}$$

with $|\alpha| < \pi/2$ to ensure that the $\mathcal{PT}$ symmetry of H is unbroken [54]. H is $\mathcal{PT}$-symmetric with regard to $\mathcal{P} = \sigma_x = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$ and $\mathcal{T}$ being complex conjugation. With $b = \sqrt{1 + (\sin\alpha)^2}$, the system's eigenvalues are ± 1, allowing the simulation of its exact time evolution at given time points even with the sparse quantum processing unit graph.

8. System 8: Two-qubit system

The Hamiltonian is

$$H_8 = \frac{\pi}{4}\left(\mathbb{1}\otimes\sigma_y - \mathbb{1}\otimes\sigma_z + \sigma_x\otimes\sigma_x - \sigma_y\otimes\sigma_x\right). \quad (19)$$

This Hamiltonian is also entanglement-producing, and the resulting QUBO is a connected graph that does not fit natively on the Pegasus and Zephyr graphs, i.e., D-Wave Advantage and Advantage2 architectures, respectively.

III. PERFORMANCE EVALUATION

A. Sampling efficiency

The quality of a quantum annealer is most clearly revealed by the success probability, defined as the fraction of returned bit strings that coincide with the exact ground state of the Ising cost function introduced in Sec. II C. For every data point we collect 5000 annealing runs and repeat the procedure for four schedule lengths $T_A \in \{10, 100, 200, 500\}$ μs. We distinguish two classes of instances: (1) Native problems that fit directly onto the Pegasus connectivity of the chip, requiring no auxiliary qubits as shown in Fig. 4(a) for system 1, Eq. (12). (2) Embedded problems that violate Pegasus constraints and therefore need minor-embedding chains, for example, system 3, Eq. (14), as shown in Fig. 4(d). Figures 4(b) and 4(e) compare the raw success probabilities obtained on two generations of D-Wave processors and illustrate how the corresponding low-energy samples reconstruct the target wave function. Across the entire dataset the Advantage2

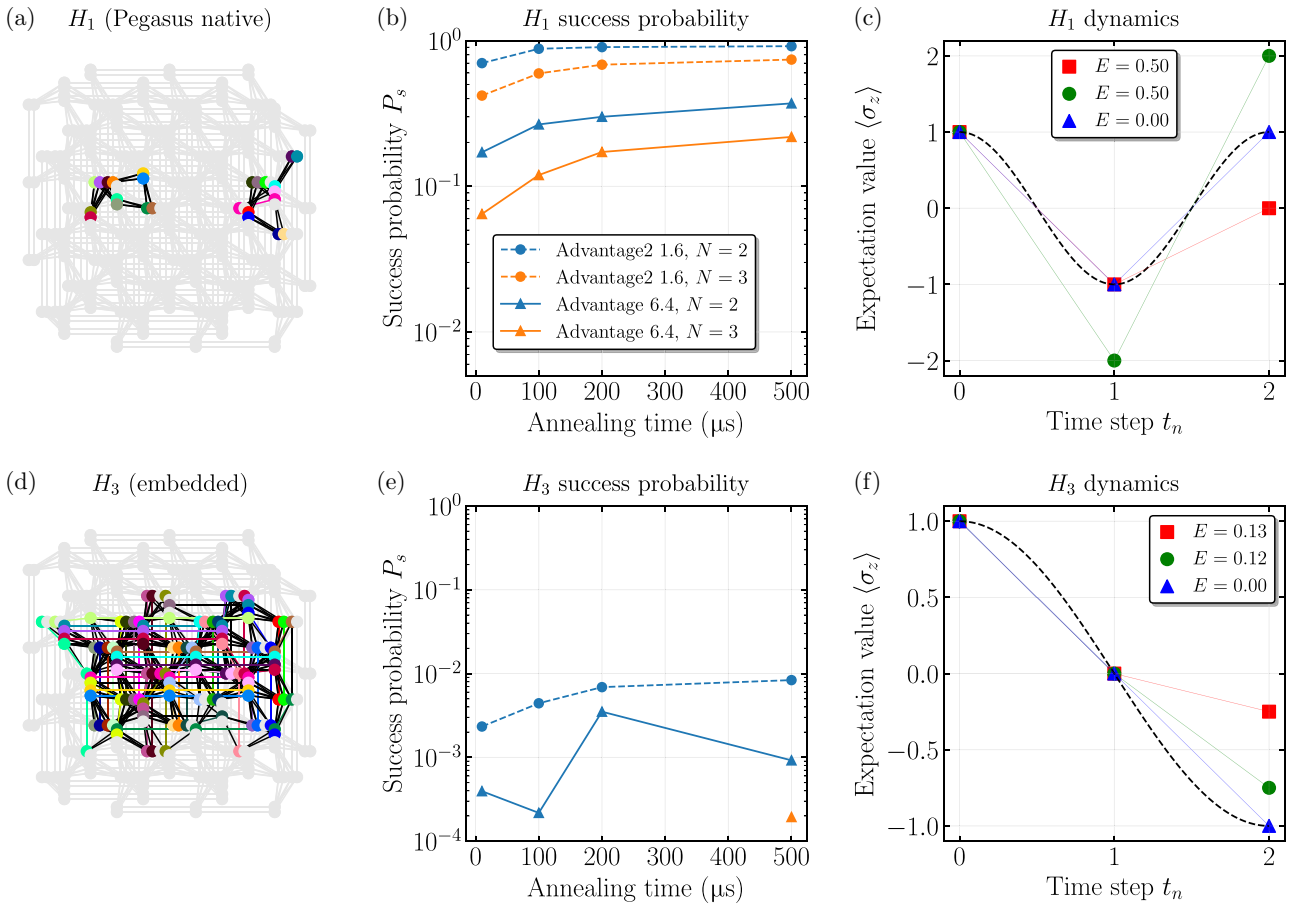

FIG. 4. (a),(d) Colored chains illustrate how the logical qubits of two representative model Hamiltonians are mapped onto the physical Pegasus P16 topology of the D-Wave hardware. (a) The native single-qubit problem H_1, Eq. (12), which can be directly solved on the quantum annealer without an additional embedding step. (d) The non-native single-qubit problem H_3, Eq. (14), where minor embedding produces chains that span several unit cells. (b),(e) Ground-state success probability obtained on the two quantum annealers Advantage2 (solid lines) and the earlier Advantage system (dashed lines) as a function of the annealing time. Circles correspond to a temporal discretization with two time points, and triangles correspond to a temporal discretization with three rimes points. Each marker averages 5000 annealing runs; if no ground-state sample is observed, the marker is omitted. (c),(f) Expectation value $\langle\sigma_z\rangle$ reconstructed from the annealer output of the three lowest-energy states of the QUBO model of the target Hamiltonian. Blue triangles denote the ground state, and green squares and red circles denote the first and second excited states, respectively. The dashed black curve is the exact unitary evolution computed with QuTiP [55–59] and serves as a baseline.

TABLE I. Ratio of success probabilities between Advantage2 and Advantage systems for Pegasus-native and non-native problems. Success probabilities are averaged over 5000 samples and system type (Pegasus native and non-native).

	Advantage2:Advantage success probability ratio	
T_A	Native	Non-native
10	5.280	20.500
100	4.778	14.000
200	4.185	4.000
500	3.768	26.500

processor leads the benchmark test. For native problems the median success probability increases by a factor of between 4 and 5; for embedded problems the gain grows to 1 order of magnitude. Table I summarizes these trends, which also confirm observations made for other problem types, such as maximum clique and unweighted maximum cut problems [60]. At the shortest schedule ($T_A = 10$ μs) higher success rates—in particular the more than 20-fold gain on non-native problems—point to tighter control of the time-dependent Hamiltonian and reduced calibration error. At intermediate schedules ($T_A = 100$–200 μs) the advantage remains, with ratios ranging from about 4 to 14, showing that improvements are not confined to the shortest anneals. At the longest schedule ($T_A = 500$ μs) the advantage grows again, especially for non-native problems. Minor embedding possibly introduces chain breaks and therefore lowers the absolute success probability on *all* processors. Even so, as a result of its greater connectivity, Advantage2 secures ground states roughly 4 times to 25 times more often than its direct predecessor, pushing several previously intractable parameter regimes into the range of routine experimentation.

B. Performance scaling

We evaluate the samplers' scaling using the T_S metric, which is the computing time required to solve the optimization problem with a given certainty P_{target} or higher. T_S is defined as

$$T_S = \frac{\ln(1 - P_{\text{target}})}{\ln(1 - P_{\text{success}})} T_r, \tag{20}$$

where P_{success} is the probability of finding the ground state in a single run and T_r is the time required for one run. In the following, we define P_{success} to be 99%. To allow comparison between the classical solver VeloxQ and the D-Wave QPUs, we measure the success probability obtained in one run, which we set to be 1000 samples both on the D-Wave QPUs and on VeloxQ. The time used to calculate T_S is the time required for one run, averaged over five runs in the case of D-Wave QPUs, and 20 runs in the case of VeloxQ and simulated annealing. Figure 5 shows scaling of T_S for

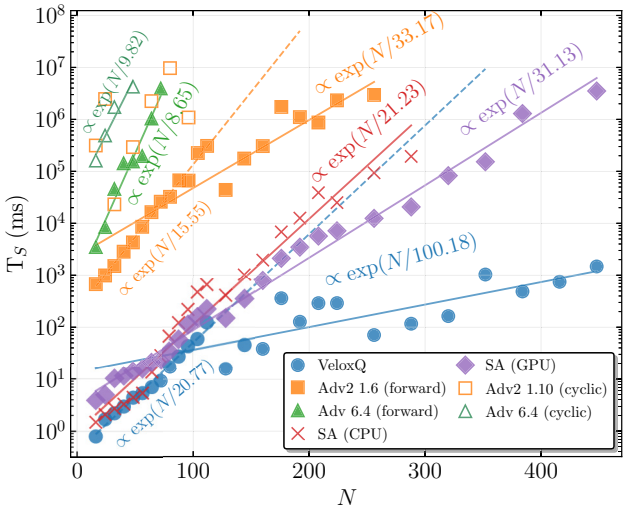

FIG. 5. Scaling properties of T_S for Pegasus-native systems [H_1, H_2, H_5, H_6, and H_7; Eqs. (12), (13) and (16)–(18), respectively] on the Advantage (Adv) and Advantage2 (Adv2) QPUs, as well as the classical solvers: simulated annealing (SA), both CPU and GPU based, and the GPU-based VeloxQ. N denotes the number of problem variables. Data points indicate T_S values averaged over system size for all systems considered. Missing data points indicate failure to find the ground state. Five (20) runs with 1000 samples each were performed on quantum (classical) solvers for each system at time points $t \in \{2, 3, \ldots, 14\}$.

different instances on quantum (using forward and cyclic annealing) and classical solvers. Since T_S grows exponentially with the size of the system, we use a fitting function of the form

$$T_S \sim \exp\left(\frac{N}{\beta}\right), \tag{21}$$

where and N is the system size [61]. The most important fitting parameter is β, which is related to the average residual energy density, and subsequently is loosely identified with an effective temperature of the generated ensemble of states [62]. In other words, it is a measure of computational complexity associated with a given problem type, and can be used as a quantitative way to compare different solvers, with a larger value of β indicating better performance. Of course, one should not expect this quantity to be universal across different QUBO problems or spin glass models, with evidence pointing to its dependence on the interaction range, and with values ranging from approximately 10^3 in case of the Edwards-Anderson model on a cubic lattice [62] to as low as approximately 4.5 for an all-to-all Sherrington-Kirkpatrick model [63]. In particular, it is not yet fully clear whether quantum solvers (e.g., nonadiabatic quantum annealing) exhibit any inherent advantage over the fully classical solver.

Figure 5 shows that on the Advantage processor the fitted exponential scale parameter β increases from approximately 8.65 (forward annealing) to approximately 9.82 (cyclic annealing), indicating that cycling can increase the likelihood of reaching the lowest-energy basin, but does not change the overall exponential character of the workload. This is consistent with the mechanism of cyclic annealing as a guided refinement step: it can help escape nearby suboptimal traps, yet it remains constrained by the same hardware noise floor and the same energy landscape structure [29].

Within the overlap region, the cyclic points show substantial scatter and do not exhibit a systematic reduction of T_S relative to forward annealing. A natural interpretation is that the extra control cycles add run-time per attempt, and—once the best reachable energy has been found—additional cycles mainly accumulate overhead without proportionally increasing the probability of hitting the true ground state. This interpretation is directly supported by Fig. 3, where the best-found objective value drops quickly in the first few cycles and then plateaus.

To allow comparison between quantum and classical solvers we measure the system size as the number of variables before embedding on the QPU's working graph. Notably, the Advantage2 QPU not only has T_S that is smaller than T_S of its predecessor, but it also exhibits a better scaling exponent. It exhibits two scaling regimes, $N \lesssim 100$ ($\beta \approx 15$) and $N \gtrsim 100$ ($\beta \approx 33$), compared with the uniform Advantage QPU's scaling exponent β of approximately 9, for the instances that are native to the QPU's working graph (see Fig. 5). In the latter regimes, its scaling behavior even matches that of an efficient, GPU-based implementation of simulated annealing.

In the case of classical solvers, one should again distinguish two regimes: $N \lesssim 100$ and $N \gtrsim 100$. In the former, all classical solvers behave similarly in terms of scaling, as well as the individual T_S values. For the latter, the advantage of GPU over CPU becomes clearly visible, with GPU-based simulated annealing outperforming CPU-based simulated annealing. Nevertheless, VeloxQ's highly optimized GPU parallelization allows it to take the lead, with scaling exponent $\beta \approx 100$. For instances that are neither Pegasus nor Zephyr native, we are also not able to evaluate the scaling performance, because problem instances quickly become intractable for the QPU (see the system-resolved T_S plots in Fig. 7 in the Appendix).

Modern quantum annealers are limited by the sparse links that connect their qubits. This limitation is clear in our tests on the single-qubit model H_3, Eq. (14), and the two-qubit model H_8, Eq. (19). The hardware reached the true ground state only twice for H_3 and never for H_8. The classical algorithm, VeloxQ, faces no such connectivity constraints and solved the same dynamical problems at every time step. Our results reinforce earlier reports showing that VeloxQ outperforms Pegasus-based and Zephyr-based quantum processors on problems that do not match their native coupling graphs [16].

Finally, we extend our benchmarks beyond instances that fit (even with embedding) on current generation annealers. This allows us to establish a classical baseline, within the framework introduced in this work, against which the performance of future quantum devices might be measured. We select three Pegasus-native systems (H_1, H_2, and H_6), and construct corresponding QUBO instances with 100, 1000, 2000, 5000, and 10 000 discretization time points for H_1 and H_2 and 100, 1000, and 2000 discretization time points for H_6. This results in QUBO instances with sizes N ranging from 800 to 80 000 variables. Since it is infeasible for a heuristic solver to obtain the true ground state in reasonable time, we relax our performance gauge and use the so-called time to ε (T_ε) [12,15], which measures the expected time needed to reach a solution with optimality gap at most ε:

$$T_\varepsilon = \frac{\ln(1-0.99)}{\ln(1-P_{E\le E_0+\varepsilon|E_0|})} T_r. \tag{22}$$

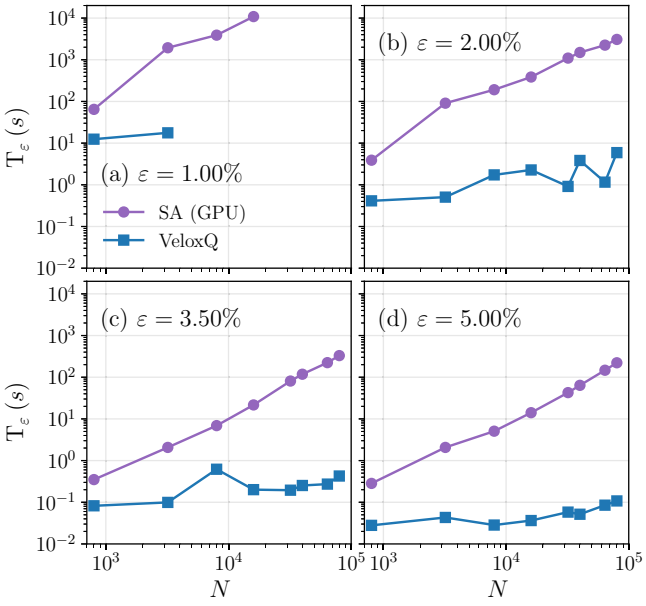

FIG. 6. Scaling of T_ε for GPU-based classical heuristics, VeloxQ and simulated annealing (SA), on large-scale QUBO instances (N up to 10^5) generated from systems H_1, H_2, and H_6. Both algorithms display characteristic power-law scaling, where the steepness of the slope is inversely related to the target optimality gap ε. This behavior indicates a crossover toward exponential complexity as the target error vanishes ($\varepsilon \to 0$). For almost all approximate targets shown ($\varepsilon \in \{1.0, 2.0, 3.5, 5.0\}\%$), VeloxQ demonstrates a consistent advantage over simulated annealing. The only exception is for $\varepsilon = 1.0\%$, where VeloxQ fails to reach the target earlier for increasing system size than simulated annealing.

For the comparison, we consider only the GPU-based VeloxQ solver and the GPU version of simulated annealing. To estimate the success probability, we performed 100 independent runs in the former case, and 25 in the latter. The obtained values of T_ε were then optimized over the relevant control parameters of the solvers, their being the number of steps and simultaneously evolved trajectories for VeloxQ, and the number of Monte Carlo sweeps and simultaneously evolved replicas for simulated annealing. The results are shown in Fig. 6. Both solvers exhibit power-law-like scaling, with the slope increasing with decreasing ε, and eventually diverging to infinity as ε tends to 0. This indicates the transition to the exponential scaling of the time to solution and is consistent with the results presented in Fig. 6. We see that for all considered optimality targets $\varepsilon \in \{1.0, 2.0, 3.5, 5.0\}\%$, VeloxQ produces solutions of the desired quality significantly faster (on average) than simulated annealing. The only exception is the case of the strictest target of $\varepsilon = 1\%$, where both solvers cannot reach it for the largest instances, but VeloxQ fails earlier for increasing system size than simulated annealing. Nevertheless, VeloxQ proves itself as a capable solver, especially in the realm of approximate optimization of very large problems.

IV. CONCLUSION

We introduced a workflow that recasts continuous quantum dynamics into a sequence of binary optimization tasks. Those tasks can be processed without modification on two very different kinds of hardware: superconducting quantum annealers and classical quadratic unconstrained binary optimization solvers. The common encoding allowed us to perform a direct, quantitative comparison between the D-Wave Advantage2 processor and the GPU-accelerated classical solver VeloxQ.

For the set of dynamical problems tested, D-Wave Advantage2 produced the correct ground-state answer roughly 10 times more often than its predecessor, and the increase held across all annealing schedules that we explored. In practical terms the probability of success rises, while the time required to reach a fixed 99% confidence level—often used as an application benchmark—increases far more slowly with problem size than on the earlier chip. These gains mirror the ongoing improvements in qubit coherence, control precision, and on-chip connectivity delivered by the latest fabrication cycle. We additionally tested cyclic annealing as a refinement step for selected instances, and found that it typically improves the best objective rapidly in the first few cycles before saturation occurs. At the aggregate level, this translates into only modest changes in time-to-solution scaling. These observations indicate that cyclic annealing is most useful as a targeted local-improvement heuristic, while the dominant performance gains in our benchmark are driven by the hardware generation upgrade to Advantage2.

At the same time the classical reference solver, VeloxQ, retains the shortest absolute run-time. Optimized data structures, problem-specific heuristics, and the massive parallelism of modern graphics cards still outweigh the quantum speedups that current devices can offer. This balance of power illustrates a central theme in contemporary computing: hardware breakthroughs and algorithmic ingenuity progress on separate tracks, and leadership can pass from one to the other multiple times before a mature technology emerges.

The prospects for quantum hardware are encouraging. Our mapping requires additional qubits only in proportion to the number of simulated time steps, rather than the square or cubic growth familiar from many other encodings. As denser qubit lattices and wider coupler graphs move from prototype to production, medium-scale dynamical simulations covering tens to hundreds of effective degrees of freedom should become accessible within a single annealing run. Whether the next generation of processors merely matches or overtakes the best classical heuristics will depend on continued advances in both camps.

Our large-scale benchmarks (N up to approximately 10^5) confirm that VeloxQ significantly outperforms simulated annealing, exhibiting favorable power-law scaling. This establishes a high-performance classical baseline that helps to identify the threshold for future quantum advantage. The open-source benchmark suite released with this work provides a transparent reference against which that progress can be measured.

ACKNOWLEDGMENTS

The authors acknowledge the Jülich Supercomputing Centre for providing computing time on the D-Wave Advantage system JUPSI through the Jülich UNified Infrastructure for Quantum computing (JUNIQ). P.H. and B.G. acknowledge funding from the National Science Centre (NCN), Poland, under the project Sonata Bis 10 (Grant No. 2020/38/E/ST3/00269). Z.M. acknowledges funding from the Ministry of Economic Affairs, Labour and Tourism Baden-Württemberg in the frame of the Competence Center Quantum Computing Baden-Württemberg (project KQCBW25). GPU mining icons were created by Ahmad Roaayala [64].

DATA AVAILABILITY

The data that support the findings of this article are openly available [65].

APPENDIX: SYSTEM-RESOLVED TIME-TO-SOLUTION RESULTS

In the appendix we show system-resolved time-to-solution results, which show the impact of different quantum Hamiltonians on the difficulty of resulting QUBO problems.

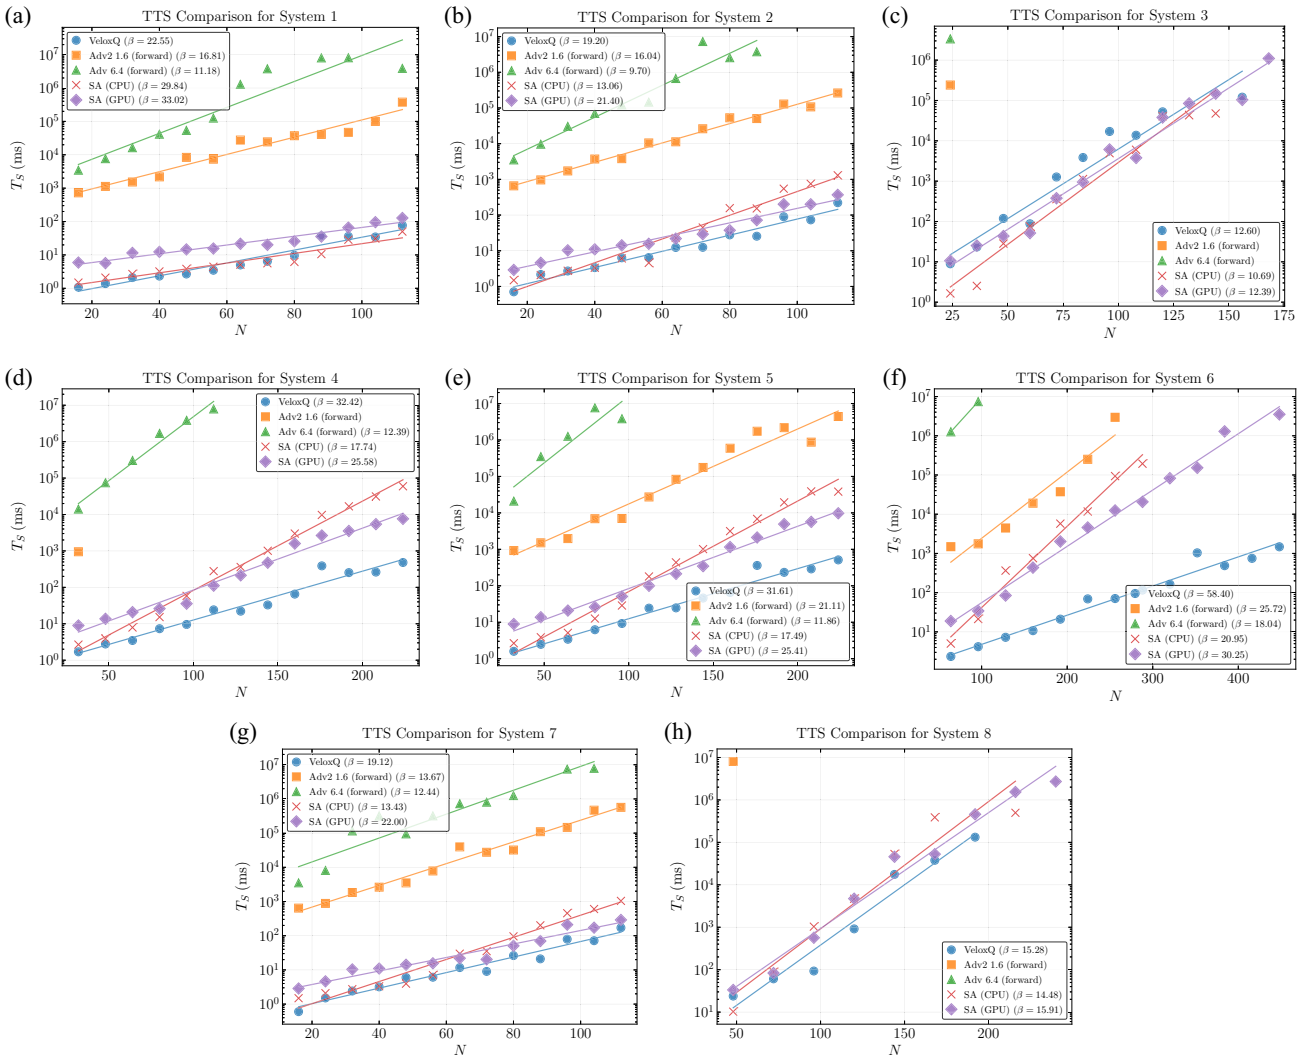

FIG. 7. System-resolved T_S as a function of system size, demonstrating clear exponential scaling proportional to $\exp(N/\beta)$, with β indicated in the legends. Missing data points indicate zero success probability, which is the case, for example, for the quantum solver in systems H_3 and H_8. In the $N \gtrsim 100$ regime the advantage of GPU solvers becomes clearly visible. Nevertheless, the Advantage2 1.6 QPU demonstrates the developing potential of quantum annealers. (a) System H_1. (b) System H_2. (c) System H_3. (d) System H_4. (e) System H_5. (f) System H_6. (g) System H_7. (h) System H_8. SA, simulated annealing.

[1] B. P. Lanyon, C. Hempel, D. Nigg, M. Müller, R. Gerritsma, F. Zähringer, P. Schindler, J. T. Barreiro, M. Rambach, G. Kirchmair, M. Hennrich, P. Zoller, R. Blatt, and C. F. Roos, Universal digital quantum simulation with trapped ions, Science **334**, 57 (2011).

[2] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, Probing many-body dynamics on a 51-atom quantum simulator, Nature **551**, 579 (2017).

[3] T. I. Andersen, N. Astrakhantsev, A. H. Karamlou, J. Berndtsson, J. Motruk, A. Szasz, J. A. Gross, A. Schuckert, T. Westerhout, Y. Zhang *et al.*, Thermalization and criticality on an analogue–digital quantum simulator, Nature **638**, 79 (2025).

[4] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, Simulated quantum computation of molecular energies, Science **309**, 1704 (2005).

[5] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, Nat. Commun. **5**, 4213 (2014).

[6] T. N. Georges, M. Bothe, C. Sünderhauf, B. K. Berntson, R. Izsák, and A. V. Ivanov, Quantum simulations of chemistry in first quantization with any basis set, npj Quantum Inf. **11**, 55 (2025).

[7] S. Euchner and I. Lesanovsky, Rydberg atom arrays as quantum simulators for molecular dynamics, Phys. Rev. Res. **7**, L042009 (2025).

[8] R. P. Feynman, Simulating physics with computers, Int. J. Theor. Phys. **21**, 467 (1982).

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

[10] O. Lanes, M. Beji, A. D. Corcoles, C. Dalyac, J. M. Gambetta, L. Henriet, A. Javadi-Abhari, A. Kandala, A. Mezzacapo, C. Porter, S. Sheldon, J. Watrous, C. Zoufal, A. Dauphin, and B. Peropadre, A framework for quantum advantage, arXiv:2506.20658 [quant-ph].

[11] A. D. King, A. Nocera, M. M. Rams, J. Dziarmaga, R. Wiersema, W. Bernoudy, J. Raymond, N. Kaushal, N. Heinsdorf, R. Harris *et al.*, Beyond-classical computation in quantum simulation, Science **388**, 199 (2025).

[12] H. Munoz-Bauza and D. Lidar, Scaling advantage in approximate optimization with quantum annealing, Phys. Rev. Lett. **134**, 160601 (2025).

[13] J. Tindall, A. Mello, M. Fishman, M. Stoudenmire, and D. Sels, Dynamics of disordered quantum systems with two- and three-dimensional tensor networks, arXiv:2503.05693 [quant-ph].

[14] L. Mauron and G. Carleo, Challenging the quantum advantage frontier with large-scale classical simulations of annealing dynamics, arXiv:2503.08247 [quant-ph].

[15] J. Pawłowski, P. Tarasiuk, J. Tuziemski, L. Pawela, and B. Gardas, Closing the quantum-classical scaling gap in approximate optimization, arXiv:2505.22514 [quant-ph].

[16] J. Pawłowski, J. Tuziemski, P. Tarasiuk, A. Przybysz, R. Adamski, K. Hendzel, L. Pawela, and B. Gardas, VeloxQ: A fast and efficient QUBO solver, arXiv:2501.19221 [quant-ph].

[17] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell *et al.*, Quantum supremacy using a programmable superconducting processor, Nature **574**, 505 (2019).

[18] Y. Wu, W.-S. Bao, S. Cao, F. Chen, M.-C. Chen, X. Chen, T.-H. Chung, H. Deng, Y. Du, D. Fan *et al.*, Strong quantum computational advantage using a superconducting quantum processor, Phys. Rev. Lett. **127**, 180501 (2021).

[19] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. van den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, and A. Kandala, Evidence for the utility of quantum computing before fault tolerance, Nature **618**, 500 (2023).

[20] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu *et al.*, Quantum computational advantage using photons, Science **370**, 1460 (2020).

[21] L. S. Madsen, F. Laudenbach, M. F. Askarani, F. Rortais, T. Vincent, J. F. F. Bulmer, F. M. Miatto, L. Neuhaus, L. G. Helt, M. J. Collins *et al.*, Quantum computational advantage with a programmable photonic processor, Nature **606**, 75 (2022).

[22] D. Hangleiter and J. Eisert, Computational advantage of quantum random sampling, Rev. Mod. Phys. **95**, 035001 (2023).

[23] K. Michielsen, M. Nocon, D. Willsch, F. Jin, T. Lippert, and H. De Raedt, Benchmarking gate-based quantum computers, Comput. Phys. Commun. **220**, 44 (2017).

[24] B. Fauseweh, Quantum many-body simulations on digital quantum computers: State-of-the-art and future challenges, Nat. Commun. **15**, 2123 (2024).

[25] I. M. Georgescu, S. Ashhab, and F. Nori, Quantum simulation, Rev. Mod. Phys. **86**, 153 (2014).

[26] V. Mehta, H. De Raedt, K. Michielsen, and F. Jin, Understanding the physics of D-wave annealers: From Schrödinger to Lindblad to Markovian dynamics, Phys. Rev. A **112**, 032616 (2025).

[27] K. Jałowiecki, A. Więckowski, P. Gawron, and B. Gardas, Parallel in time dynamics with quantum annealers, Sci. Rep. **10**, 13534 (2020).

[28] M. Born and V. Fock, Beweis des Adiabatensatzes, Z. Phys. **51**, 165 (1928).

[29] H. Zhang, K. Boothby, and A. Kamenev, Cyclic quantum annealing: Searching for deep low-energy states in 5000-qubit spin glass, Sci. Rep. **14**, 30784 (2024).

[30] H. Zhang and A. Kamenev, Computational complexity of three-dimensional Ising spin glass: Lessons from D-wave annealer, Phys. Rev. Res. **7**, 033098 (2025).

[31] H. Wang, H.-C. Yeh, and A. Kamenev, Many-body localization enables iterative quantum optimization, Nat. Commun. **13**, 5503 (2022).

[32] M. H. Muñoz-Arias and P. M. Poggi, Quantum ergodicity and scrambling in quantum annealers, Quantum Sci. Technol. **10**, 045027 (2025).

[33] H. Zhang and A. Kamenev, Quantum sequel of neural network training, Commun. Phys. **8**, 478 (2025).

[34] S. Schulz, D. Willsch, and K. Michielsen, Learning-driven annealing with adaptive Hamiltonian modification for solving large-scale problems on quantum devices, Quantum **9**, 1898 (2025).

[35] A. D. King, J. Raymond, T. Lanting, R. Harris, A. Zucca, F. Altomare, A. J. Berkley, K. Boothby, S. Ejtemaee, C. Enderud *et al.*, Quantum critical dynamics in a 5,000-qubit programmable spin glass, Nature **617**, 61 (2023).

[36] A. Teplukhin, B. K. Kendrick, S. M. Mniszewski, Y. Zhang, A. Kumar, C. F. A. Negre, P. M. Anisimov, S. Tretiak, and P. A. Dub, Computing molecular excited states on a D-Wave quantum annealer, Sci. Rep. **11**, 18796 (2021).

[37] F. Neukart, G. Compostella, C. Seidel, D. von Dollen, S. Yarkoni, and B. Parney, Traffic flow optimization using a quantum annealer, Front. ICT **4**, 29 (2017).

[38] M. Koniorczyk, K. Krawiec, L. Botelho, N. Bešinović, and K. Domino, Solving rescheduling problems in heterogeneous urban railway networks using hybrid quantum–classical approach, J. Rail Transp. Plan. Manage. **34**, 100521 (2025).

[39] K. Domino, E. Doucet, R. Robertson, B. Gardas, and S. Deffner, On the Baltimore Light RailLink into the quantum future, Sci. Rep. **15**, 29576 (2025).

[40] T. Śmierzchalski, J. Pawłowski, A. Przybysz, Ł. Pawela, Z. Puchała, M. Koniorczyk, B. Gardas, S. Deffner, and K. Domino, Hybrid quantum-classical computation for automatic guided vehicles scheduling, Sci. Rep. **14**, 21809 (2024).

[41] K. Domino, M. Koniorczyk, K. Krawiec, K. Jałowiecki, S. Deffner, and B. Gardas, Quantum annealing in the NISQ era: Railway conflict management, Entropy **25**, 191 (2023).

[42] E. Kędziera, W. Gamon, M. Koniorczyk, Z. Mzaouali, A. Galadíková, and K. Domino, EMU circulation planning for Silesian railways: Case study and a quantum approach, arXiv:2512.19340 [quant-ph].

[43] R. Robertson, E. Doucet, Z. Mzaouali, K. Domino, B. Gardas, and S. Deffner, in *2025 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, Albuquerque, NM, USA, 2025), Vol. 01, pp. 190–196.

[44] L. Buffoni and M. Campisi, Thermodynamics of a quantum annealer, Quantum Sci. Technol. **5**, 035013 (2020).

[45] M. Campisi and L. Buffoni, Improved bound on entropy production in a quantum annealer, Phys. Rev. E **104**, L022102 (2021).

[46] T. Śmierzchalski, Z. Mzaouali, S. Deffner, and B. Gardas, Efficiency optimization in quantum computing: balancing thermodynamics and computational performance, Sci. Rep. **14**, 4555 (2024).

[47] E. Doucet, Z. Mzaouali, R. Robertson, B. Gardas, S. Deffner, and K. Domino, Thermodynamic significance of QUBO encoding on quantum annealers, arXiv:2601.04402 [quant-ph].

[48] M. S. Könz, W. Lechner, H. G. Katzgraber, and M. Troyer, Embedding overhead scaling of optimization problems in quantum annealing, PRX Quantum **2**, 040322 (2021).

[49] V. Mehta, H. De Raedt, K. Michielsen, and F. Jin, Unraveling reverse annealing: A study of D-Wave quantum annealers, Phys. Rev. A **112**, 012414 (2025).

[50] V. Mehta, H. De Raedt, K. Michielsen, and F. Jin, Performance of quantum annealing for 2-satisfiability problems with multiple satisfying assignments, Phys. Rev. A **112**, 012405 (2025).

[51] J. Vodeb, V. Eržen, T. Hrga, and J. Povh, Accuracy and performance evaluation of quantum, classical and hybrid solvers for the max-cut problem, arXiv:2412.07460 [math.OC].

[52] A. Lucas, Ising formulations of many NP problems, Front. Phys. **2**, 74887 (2014).

[53] B. Gardas, S. Deffner, and A. Saxena, $\mathcal{PT}$-symmetric slowing down of decoherence, Phys. Rev. A **94**, 040101(R) (2016).

[54] S. Croke, $\mathcal{PT}$-symmetric Hamiltonians and their application in quantum information, Phys. Rev. A **91**, 052113 (2015).

[55] N. Lambert, E. Giguère, P. Menczel, B. Li, P. Hopf, G. Suárez, M. Gali, J. Lishman, R. Gadhvi, R. Agarwal, A. Galicia, N. Shammah, P. Nation, J. R. Johansson, S. Ahmed, S. Cross, A. Pitchford, and F. Nori, QuTiP 5: The Quantum Toolbox in Python, arXiv:2412.04705 [quant-ph].

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

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

[58] B. Li, S. Ahmed, S. Saraogi, N. Lambert, F. Nori, A. Pitchford, and N. Shammah, Pulse-level noisy quantum circuits with QuTiP, Quantum **6**, 630 (2022).

[59] N. Lambert, T. Raheja, S. Cross, P. Menczel, S. Ahmed, A. Pitchford, D. Burgarth, and F. Nori, QuTiP-BoFiN: A bosonic and fermionic numerical hierarchical-equations-of-motion library with applications in light-harvesting, quantum control, and single-molecule electronics, Phys. Rev. Res. **5**, 013181 (2023).

[60] E. Pelofske, Comparing three generations of D-Wave quantum annealers for minor embedded combinatorial optimization problems, Quantum Sci. Technol. **10**, 025025 (2025).

[61] V. Mehta, F. Jin, H. De Raedt, and K. Michielsen, Quantum annealing for hard 2-satisfiability problems: Distribution and scaling of minimum energy gap and success probability, Phys. Rev. A **105**, 062406 (2022).

[62] H. Zhang and A. Kamenev, Computational complexity of three-dimensional Ising spin glass: Lessons from D-Wave annealer, Phys. Rev. Res. **7**, 033098 (2025).

[63] A. Montanaro, Quantum speedup of branch-and-bound algorithms, Phys. Rev. Res. **2**, 013056 (2020).

[64] https://www.flaticon.com/free-icons/gpu-mining, 433 Gpu mining icons, Ahmad Roaayala, Flaticon.

[65] P. Hanussek, J. Pawłowski, Z. Mzaouali, and B. Gardas, Quantum dynamics benchmarking framework, 2025, https://github.com/atg205/dynamics-benchmark.