

Benchmark study of simulated quantum annealing and asynchronous parallel simulated annealing for quadratic unconstrained binary optimization-based semiprime factorization

Gianbiagio Curato ^{*}, Manuel Ponzi , Edoardo Chiatello , and Roberto Giorgetti 
Intesa Sanpaolo, Turin, Italy



(Received 3 March 2026; published 21 July 2026)

Simulated quantum annealing (SQA) has been widely studied as a heuristic for solving hard optimization problems, including integer factorization formulated as a quadratic unconstrained binary optimization (QUBO) problem. In this work, we present a replication-oriented benchmark of annealing-based solvers applied to biprime factorization using the direct QUBO mapping. We perform a controlled comparison of SQA, implemented via path-integral Monte Carlo, simulated annealing (SA), and asynchronous parallel simulated annealing (A-PSA), evaluating their performance in terms of probability of success and time to solution. Here, A-PSA denotes a multistart strategy consisting of multiple independent SA trajectories with no inter-replica interactions. Solver and encoding parameters are systematically optimized using Bayesian optimization to ensure a fair and reproducible comparison. Our results reproduce previously reported performance trends: SQA outperforms conventional SA on the considered QUBO landscapes, while SQA and A-PSA can exhibit closely comparable performance in specific regimes, highlighting the importance of multistart baselines in annealing-based benchmarking. All simulations are performed on commodity multithreaded CPU hardware using open-source software, enabling straightforward reproducibility without reliance on specialized experimental platforms. As a reference benchmark, the largest semiprime successfully factorized in our study is $8\,219\,999 = 32\,749 \times 251$, matching results previously obtained using quantum annealing hardware.

DOI: [10.1103/hj46-qw5t](https://doi.org/10.1103/hj46-qw5t)

I. INTRODUCTION

Integer factorization is a canonical computational problem with central relevance to public-key cryptography, most notably the Rivest—Shamir—Adleman (RSA) cryptosystem [1–8]. The problem has also become a standard benchmark for assessing emerging computational paradigms, including quantum and quantum-inspired optimization methods. While Shor's algorithm demonstrates that fault-tolerant quantum computers could factor large integers efficiently [9], current quantum hardware remains far from the scale required for cryptographically relevant instances [10–16]. As a result, alternative approaches based on heuristic optimization, quantum annealing, and their classical simulations continue to attract attention. Quantum annealing (QA) and related techniques aim to solve optimization problems by exploiting quantum fluctuations to navigate

^{*}Contact author: gianbiagio.curato@intesanpaolo.com

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

complex energy landscapes. In practice, access to physical quantum annealing hardware is limited, and device performance is constrained by noise, connectivity, and control imperfections. For this reason, *simulated quantum annealing* (SQA), typically implemented via path-integral Monte Carlo, has become a widely used classical proxy for studying QA dynamics and benchmarking annealing-based heuristics [17–26]. Classical *simulated annealing* (SA) [27] provides a natural reference point for such comparisons. In this work, we also consider a parallel classical baseline based on running multiple SA trajectories concurrently with no interprocess communication, which we denote as *asynchronous parallel simulated annealing* (A-PSA). This corresponds to the “asynchronous” class of parallel simulated annealing (PSA) methods in which parallel runs evolve independently and no information is exchanged among replicas [28,29]. Such “parallel multistart” or “independent trajectory” strategies are widely used as a scalable parallelization model for trajectory-based metaheuristics [30,31]. Importantly, A-PSA should not be confused with *Replicated Simulated Annealing* schemes based on coupled replicas and robust-ensemble objectives, as introduced by Baldassi *et al.* [32,33] and further analyzed by Gripon *et al.* [34]. In those replicated/robust-ensemble methods, replicas are *explicitly coupled* (e.g., via a reference/centroid term), whereas our A-PSA baseline uses *no inter-replica interactions*. A substantial body of prior work has investigated the relative performance of QA, SQA, and classical annealing-based heuristics on a variety of optimization problems [20,35–37]. Beyond comparisons with SA, SQA has also been evaluated against stronger classical sampling/optimization baselines that exploit parallelism, such as parallel tempering, in both algorithmic studies and hardware-oriented probabilistic Ising machines [38–40]. These works reinforce the need for careful, resource-aware baselines when assessing SQA performance and motivate including an explicit A-PSA control that captures the benefit of multistart exploration under comparable computational budgets. For stoquastic Hamiltonians, it is well established that SQA can outperform conventional SA on some landscapes, and that the precise relationship between SQA dynamics and quantum adiabatic optimization algorithm can be instance dependent [41,42]. Within this context, the contribution of the present work is *not* to introduce a factorization algorithm or claim quantum advantage, but rather to provide a careful and reproducible replication study focused on integer factorization formulated as a quadratic unconstrained binary optimization (QUBO) problem. We adopt the direct method for mapping biprime factorization into QUBO form [43] and benchmark SQA, SA, and A-PSA under controlled and explicitly documented conditions. To the best of our knowledge, controlled benchmarks that directly compare SQA against an explicitly asynchronous PSA (multistart SA) baseline on QUBO-based semiprime factorization remain limited; accordingly, our study aims to help fill this gap with a replication-oriented methodology. Recent experimental work by Grimaldi *et al.* [44] implemented SQA using stochastic magnetic tunnel junctions on field-programmable gate array (FPGA)-based probabilistic computing hardware, demonstrating factorization of integers up to 26 bits using thousands of p-bits. A key conclusion of that study was that, when normalized for computational resources, SQA does not significantly outperform an optimized classical approach. In this work, we treat such results as a *replication target* and investigate whether analogous performance trends can be observed using only commodity CPU hardware and open-source software implementations. Our study differs from prior hardware-focused investigations in three deliberate aspects: *hardware*, *software*, and *methodological transparency*. All simulations are performed on standard multithreaded CPU platforms using widely adopted open-source libraries [45–49], with the explicit goal of enabling straightforward reproducibility and independent verification. We systematically document solver configurations, stopping criteria, and performance metrics to support meaningful comparison across studies. To ensure a fair comparison between SQA, SA, and A-PSA, we employ Bayesian optimization to tune solver and encoding parameters under realistic computational constraints. This approach allows us to reproduce known performance regimes while controlling for confounding effects arising from suboptimal parameter choices or annealing schedule instabilities. Similar optimization strategies based on time to solution (TTS) have been employed in related benchmarking studies [50]. The primary objective of this work is therefore to provide a transparent, reproducible benchmark that clarifies under which conditions SQA offers an advantage over SA, and when its behavior converges

to that of A-PSA, in the specific setting of QUBO-based semiprime factorization. By focusing on replication rather than novelty, we aim to supply the community with a reliable reference point for future algorithmic, hardware, or hybrid comparisons. The paper is structured as follows. In Sec. II, we introduce the QUBO formulation of the factorization problem and the annealing-based solvers. Section III describes the Bayesian optimization pipeline used to tune solver parameters. In Sec. IV, we present the benchmarking results for SQA, SA, and A-PSA. Section V summarizes the findings and discusses their implications. Four Appendixes provide additional supporting material: Appendix A estimates the search-space cardinality; Appendix B discusses an alternative polynomial constrained Boolean optimization (PCBO) formulation; Appendix C presents a structural analysis of the derived Ising models; and Appendix D reports the optimized parameters used in the simulations.

II. METHODS

In 2001, Burges [43] demonstrated that factoring a biprime number N into its prime numbers p and q can be mapped to an optimization problem. Consider an integer N that could be represented in binary form using L_N bits:

$$\lceil \log_2(N) \rceil = L_N, \quad (1)$$

where $\lceil \cdot \rceil$ is the ceiling function. The factorization problem asks to find the two prime factors p and q such that

$$N = p \cdot q. \quad (2)$$

We assume that N is odd as otherwise the factorization problem is trivial. Expressing the above numbers N , p , and q in binary, the following relations are obtained:

$$N = \sum_{i=0}^{L_N-1} 2^i n_i, \quad p = \sum_{j=0}^{L_p-1} 2^j p_j, \quad q = \sum_{k=0}^{L_q-1} 2^k q_k, \quad (3)$$

where $n_i \in \{0, 1\}$ are known bits of N and $p_j, q_k \in \{0, 1\}$ are unknown binary variables. Since N is assumed to be an odd number, p and q are odd primes, which implies that the least significant bits, denoted as p_0 and q_0 , are always 1.

Our analysis mainly focuses on balanced factorizations, with L_p and L_q of comparable size, as the digital signature standard for RSA [51] recommends an even bit length for the modulus N and prime factors of approximately half that length. However, since unbalanced cases have also been explored in the literature, we consider an iterative encoding [52] capable of handling both balanced and unbalanced factorizations:

$$\begin{cases} L_p = \frac{L_N}{2} - a, \\ L_q = \frac{L_N}{2} + 1 + a, \end{cases} \quad L_N \text{ even}, \quad \begin{cases} L_p = \lfloor \frac{L_N}{2} \rfloor + 1 - a, \\ L_q = \lfloor \frac{L_N}{2} \rfloor + 1 + a, \end{cases} \quad L_N \text{ odd}, \quad (4)$$

where a is an integer initially set to zero and iteratively incremented, as described in Sec. III. Such encoding allows to handle asymmetries in bit length between factors while enforcing the condition

$$L_p + L_q = L_N + 1, \quad (5)$$

thereby ensuring that a sufficient number of bits is allocated to represent both factors.

Several approaches can be followed in order to convert this problem into an optimization problem [53]. Moreover, it often happens that the objective function $O(p, q)$ involves terms with order higher than 2. This means that $O(p, q)$ is not the desired QUBO form [53,54] but rather it is expressed as a higher-order unconstrained binary optimization (HUBO) problem [55]. The trick used to transform HUBOs into QUBOs is to add ancillary variables and substitute higher-order terms with an expression having the same minimum and containing only quadratic terms [56]. We delegate this task to the open-source library QUBOVERT [45]. A straightforward optimization model penalizes

the deviation of the product $p \cdot q$ from N :

$$\mathcal{O}(p, q) = (N - p \cdot q)^2, \quad (6)$$

where the square at the exponent ensures that the function has its global minimum only when p and q are factors of N , i.e., where the objective function has value 0. Knowing the value of the objective function at the global minimum is not trivial information, because it allows the computation of the probability of success metric [57,58]. This is the basis of our approach to the search of optimal parameters for simulations in Sec. III.

Expanding p and q in their binary forms yields

$$\mathcal{O}(p, q) = \left(\sum_{i=0}^{L_N-1} 2^i n_i - \left(\sum_{j=0}^{L_p-1} 2^j p_j \right) \left(\sum_{k=0}^{L_q-1} 2^k q_k \right) \right)^2. \quad (7)$$

This objective is polynomial in the binary variables but contains many cubic and higher-order terms, which can be burdensome when converting to a QUBO (quadratic) form. Various approaches have been proposed to find this minimum by using a Hamiltonian to find the ground state. Most methods use arithmetic simplifications of equations resulting from the binary multiplication table when multiplying p and q as binary numbers. For a detailed explanation of those approaches, we refer to Ref. [59].

The QUBOVERT library allows the conversion of the objective function to a QUBO formulation. At this point, the objective function is amenable to be mapped to an Ising model [60], in order to exploit energy minimization methods coming from statistical physics. The classical Ising model describes a system of N spins $s_i \in \{-1, +1\}$ with Hamiltonian

$$H_{\text{Ising}} = - \sum_{i>j} J_{ij} s_i s_j - \sum_{i=1}^d h_i s_i + b, \quad (8)$$

where J_{ij} are coupling constants, h_i are local fields, and d is the dimension of the system. The bias term b comes from the mapping process of the QUBO problem and allows to offset the energy, preserving the minimum global value to 0. The bias term is not relevant for the search of the optimal solution and in the following we do not report this term in the Hamiltonian representation in order to have more clarity in the exposition. A visualization of the Ising model corresponding to factorization problem for $N = 143$ is shown in Fig. 1.

The entries of matrix J and vector h are determined by the QUBO encoding, so their magnitudes and patterns reflect the arithmetic structure of the problem and not statistical disorder. A detailed structural analysis is reported in Appendix C for the values of $N = 59\,989$. The list of semiprime integers considered for factorization and the corresponding Ising-model statistics are summarized in Table I, where $|J|$ is defined by Eq. (30). The list is defined from lists already analyzed in different previous studies [52,59,61,62].

A. Simulated quantum annealing

In the quantum, i.e., transverse field, Ising model (TFIM), the classical spin variables s_i are promoted to quantum operators [63–65]. The Hamiltonian becomes

$$H_{\text{TFIM}} = - \sum_{i>j} J_{ij} \hat{\sigma}_i^z \hat{\sigma}_j^z - \sum_{i=1}^d h_i \hat{\sigma}_i^z - \Gamma(t) \sum_i \hat{\sigma}_i^x \quad (9)$$

$$= H_{\text{Ising}} + H_{\text{TF}}(t), \quad (10)$$

where $\hat{\sigma}_i^z$ and $\hat{\sigma}_i^x$ are Pauli matrices acting on spin i and $\Gamma(t)$ is the strength of the transverse field. The eigenvalues of $\hat{\sigma}_i^z$ are ± 1 , which correspond to the classical Ising spin states. The transverse-field term introduces quantum fluctuations, enabling tunneling between classical configurations and




FIG. 1. Visualization of the Ising model for the eight-bit number $N = 143$, consisting of the upper-triangular interaction matrix J , the local fields h , and the global energy offset, all plotted with a shared color scale centered at zero. Asymmetric encoding is used.

thus allowing a quantum search of the phase space. Initially, the strength of the transverse field $\Gamma(t)$ is chosen to be very large, and the total Hamiltonian is dominated by the second kinetic term. The coefficient is then gradually and monotonically decreased toward 0, leaving eventually only the potential term H_{Ising} . Accordingly, the state vector is expected to evolve from the trivial initial ground state of the transverse-field term to the nontrivial ground state of H_{Ising} , which is the solution of the optimization problem.

An important issue is how slowly $\Gamma(t)$ should be reduced to keep the state vector arbitrarily close to the instantaneous ground state of the total Hamiltonian [64]. The power decrease of the transverse field guarantees the adiabatic evolution. This annealing schedule, however, does not provide practically useful algorithms because infinitely long time is necessary to reach the exact solution. An approximate algorithm for finite annealing time τ should be used in practice. Since such a finite-time algorithm does not satisfy the generic convergence condition, the answer includes a certain amount of errors. Annealing schedules showing errors proportional to τ^{2m} ($m > 1$) were developed in Ref. [64]. The ordinary Hamiltonian of quantum annealing [66–69] with finite annealing time τ is composed of the potential term and the kinetic energy term:

$$H(s) = f(s)H_{\text{Ising}} + [1 - f(s)]H_{\text{TF}}. \quad (11)$$

Here, $s = t/\tau$ is a normalized time and $\Gamma(t) = \Gamma$. The function $f(s)$, representing the annealing schedule, satisfies $f(0) = 0$ and $f(1) = 1$. Thus, $H(0) = H_{\text{TF}}$ and $H(1) = H_{\text{Ising}}$. The above

TABLE I. Semiprime integers considered for factorization and corresponding Ising-model statistics.

N	p	q	n_N	n_p	n_q	a	n_{spins}	$J_{\min}$	$J_{\max}$	$ J $	$h_{\min}$	$h_{\max}$
15	3	5	4	2	3	0	4	-6.45×10^1	3.62×10^1	3.41×10^1	-8.82×10^1	9.65×10^1
21	3	7	5	2	3	0	8	-3.86×10^2	1.97×10^2	6.82×10^1	-4.17×10^2	5.78×10^2
35	5	7	6	3	3	0	11	-9.62×10^2	3.89×10^2	3.08×10^2	-1.75×10^3	1.44×10^3
57	3	19	6	2	5	1	8	-1.02×10^3	5.12×10^2	1.66×10^2	-2.14×10^3	1.54×10^3
77	7	11	7	3	4	0	15	-4.61×10^3	1.44×10^3	5.72×10^2	-6.02×10^3	6.91×10^3
123	3	41	7	2	6	2	10	-4.10×10^3	2.05×10^3	2.56×10^2	-9.85×10^3	6.14×10^3
143	11	13	8	4	4	0	19	-2.15×10^4	7.36×10^3	1.26×10^3	-2.61×10^4	3.23×10^4
253	11	23	8	4	5	0	19	-2.15×10^4	5.12×10^3	9.29×10^2	-3.94×10^4	3.23×10^4
451	11	41	9	4	6	1	24	-9.22×10^4	2.05×10^4	2.30×10^3	-1.76×10^5	1.38×10^5
511	7	73	9	3	7	2	20	-6.45×10^4	2.46×10^4	2.05×10^3	-2.12×10^5	9.68×10^4
899	29	31	10	5	5	0	29	-4.30×10^5	7.37×10^4	4.10×10^3	-6.43×10^5	6.45×10^5
1011	3	337	10	2	9	3	16	-2.62×10^5	1.31×10^5	2.05×10^3	-6.86×10^5	3.93×10^5
1769	29	61	11	5	6	0	35	-1.84×10^6	3.12×10^5	8.19×10^3	-2.56×10^6	2.76×10^6
2047	23	89	11	5	7	1	34	-1.78×10^6	2.95×10^5	7.14×10^3	-2.96×10^6	2.67×10^6
4009	19	211	12	5	8	1	41	-7.23×10^6	1.18×10^6	1.64×10^4	-1.25×10^7	1.08×10^7
4087	61	67	12	6	7	0	41	-7.62×10^6	1.11×10^6	1.64×10^4	-1.19×10^7	1.14×10^7
5723	59	97	13	6	7	0	48	-3.15×10^7	8.48×10^6	3.28×10^4	-3.77×10^7	4.72×10^7
8189	19	431	13	5	9	2	46	-2.91×10^7	4.72×10^6	3.28×10^4	-5.12×10^7	4.37×10^7
15 853	83	191	14	7	8	0	55	-1.28×10^8	1.73×10^7	6.55×10^4	-1.92×10^8	1.92×10^8
16 379	11	1489	14	4	11	3	44	-1.00×10^8	2.10×10^7	6.55×10^4	-2.12×10^8	1.51×10^8
21 733	103	211	15	7	8	0	63	-5.20×10^8	1.51×10^8	8.19×10^4	-6.04×10^8	7.80×10^8
32 743	137	239	15	8	8	0	63	-5.20×10^8	6.82×10^7	7.78×10^4	-7.85×10^8	7.80×10^8
49 583	179	277	16	8	9	0	71	-2.10×10^9	5.11×10^8	1.64×10^5	-2.66×10^9	3.15×10^9
59 989	239	251	16	8	8	0	71	-2.10×10^9	3.41×10^8	1.64×10^5	-3.00×10^9	3.15×10^9
65 509	109	601	16	7	10	1	69	-2.07×10^9	2.77×10^8	1.64×10^5	-3.19×10^9	3.11×10^9
94 037	271	347	17	9	9	0	80	-8.46×10^9	2.24×10^9	2.87×10^5	-1.03×10^{10}	1.27×10^{10}
131 069	53	2473	17	6	12	3	74	-8.05×10^9	1.14×10^9	3.93×10^5	-1.30×10^{10}	1.21×10^{10}
159 317	313	509	18	9	9	0	89	-3.40×10^{10}	1.09×10^{10}	5.57×10^5	-3.77×10^{10}	5.09×10^{10}
262 099	349	751	18	9	10	0	89	-3.40×10^{10}	4.33×10^9	5.57×10^5	-5.12×10^{10}	5.09×10^{10}
374 599	521	719	19	10	10	0	99	-1.36×10^{11}	3.64×10^{10}	1.08×10^6	-1.66×10^{11}	2.05×10^{11}
376 289	571	659	19	10	10	0	99	-1.36×10^{11}	3.62×10^{10}	1.08×10^6	-1.66×10^{11}	2.05×10^{11}
524 281	269	1949	19	9	11	1	98	-1.36×10^{11}	1.73×10^{10}	1.08×10^6	-2.05×10^{11}	2.04×10^{11}
881 851	727	1213	20	10	11	0	109	-5.47×10^{11}	1.11×10^{11}	2.15×10^6	-7.34×10^{11}	8.20×10^{11}
1 005 973	997	1009	20	10	10	0	109	-5.47×10^{11}	7.87×10^{10}	2.15×10^6	-8.00×10^{11}	8.20×10^{11}
1 028 171	1009	1019	20	10	10	0	109	-5.47×10^{11}	7.29×10^{10}	2.15×10^6	-8.11×10^{11}	8.20×10^{11}
1 048 561	911	1151	20	10	11	0	109	-5.47×10^{11}	6.90×10^{10}	2.13×10^6	-8.22×10^{11}	8.20×10^{11}
1 794 953	907	1979	21	10	11	0	120	-2.19×10^{12}	4.30×10^{11}	4.26×10^6	-2.97×10^{12}	3.29×10^{12}
2 093 809	1447	1447	21	11	11	0	120	-2.19×10^{12}	2.75×10^{11}	4.26×10^6	-3.28×10^{12}	3.29×10^{12}
3 819 047	1873	2039	22	11	11	0	131	-8.77×10^{12}	1.48×10^{12}	8.42×10^6	-1.24×10^{13}	1.32×10^{13}
8 219 999	251	32 749	23	8	15	3	134	-3.49×10^{13}	4.65×10^{12}	1.69×10^7	-5.20×10^{13}	5.24×10^{13}

Hamiltonian connects the trivial initial state and the nontrivial desired solution after evolution time τ . The condition on the bounding error τ^{2m} holds if the k th derivative of $f(s)$ is zero at $s = 0$ and 1 for $k = 1, 2, \dots, m - 1$; moreover, $f(s)$ is an m th differentiable function whose m th derivative is continuous. Examples of such annealing schedules are the linear one or others defined in Ref. [64]. The open-source package Openjij [46] implements such schedules until the quartic power. Inverse temperature schedules for the SA and A-PSA algorithms were computed using a geometric progression.

Starting from the quantum Hamiltonian of Eq. (9), the partition function at inverse temperature β is

$$Z = \text{Tr}[e^{-\beta H_{\text{TFF}}}] \quad (12)$$

Applying the Suzuki-Trotter decomposition with M slices [70],

$$e^{-\beta(H_{\text{Ising}}+H_{\text{TFF}})} \approx \left(e^{-\Delta\tau H_{\text{Ising}}} e^{-\Delta\tau H_{\text{TFF}}}\right)^M, \quad \Delta\tau = \frac{\beta}{M}, \quad (13)$$

where $\Delta\tau$ is the width of each imaginary-time slice. $\Delta\tau = \beta/M$ is the discretization step along the *imaginary-time* axis of length $\beta = 1/(k_B T)$. It does not represent real physical time or the annealing time. It is a numerical parameter controlling the Trotter approximation error. In natural units ($\hbar = 1$), imaginary time has the same dimension as real time, but here it is purely a mathematical construct. After inserting complete sets of σ^z eigenstates and tracing over σ^x contributions, the quantum model maps to a $(d+1)$ -dimensional classical Ising model with spins $s_{i,k} \in \{\pm 1\}$, where $i = 1, \dots, d$ labels spatial sites and $k = 1, \dots, M$ labels Trotter slices. The effective classical Hamiltonian is [71,72]

$$\mathcal{H}_{\text{cl}} = - \sum_{k=1}^M \left[\sum_{i>j} J_{ij}^{(\text{cl})} s_{i,k} s_{j,k} + \sum_{i=1}^d h_i^{(\text{cl})} s_{i,k} \right] - J_\tau \sum_{i=1}^d \sum_{k=1}^M s_{i,k} s_{i,k+1}, \quad (14)$$

with periodic boundary conditions $s_{i,M+1} = s_{i,1}$ and couplings

$$J_{ij}^{(\text{cl})} = \Delta\tau J_{ij}, \quad h_i^{(\text{cl})} = \Delta\tau h_i, \quad J_\tau = \frac{1}{2} \ln \coth(\Delta\tau \Gamma). \quad (15)$$

In simulated quantum annealing, the transverse field $\Gamma(t)$ changes according to an annealing schedule. For each fixed Γ , the classical Hamiltonian uses the same $\Delta\tau$ because β and M are fixed by the simulation temperature and discretization choice. The coupling J_τ depends on $\Gamma(t)$, so quantum fluctuations are encoded in the interslice interaction, while $\Delta\tau$ remains constant throughout the schedule.

III. BAYESIAN PIPELINE PARAMETER SEARCH

Since the performance of SA, SQA, and A-PSA is highly sensitive to the choice of parameters, identifying their optimal values is essential to provide a fair assessment of the efficiency of each algorithm. To this end, we resorted to a Bayesian optimization procedure fixing the algorithm computational budget and choosing as objective function the probability of success, i.e., the fraction of successful annealing runs over a given annealing ensemble. The use of this metric is enabled by the fact that the minimum of the prime factorization objective function [Eq. (6)] is known *a priori*, which uniquely determines the ground-state energy E_0 of the associated Ising Hamiltonian.

A run is considered successful according to solver-specific criteria, with two variants defined for SA and A-PSA. In particular, one variant evaluates success solely based on the final state of the system at the end of the annealing process, while the second—referred to as *continuous evaluation* (CE)—also accounts for intermediate configurations by monitoring the system energy throughout the annealing. This choice is motivated by the absence of convergence theorems defining annealing schedules that guarantee asymptotic ground-state occupation for the considered Ising systems. The resulting five solvers, obtained from the combination of algorithms and success criteria, are described in Table II.

Therefore, we define the probability of success computed over a sample of K runs as a function of the solver $\mathcal{S}$, namely,

$$p_s(\mathcal{S}) := \frac{1}{K} \sum_{k=1}^K \mathcal{I}_{\{0\}}(E_{\text{res}}^k(\mathcal{S})), \quad (16)$$

TABLE II. Definition of the success criteria used for each solver.

Solver	Success criterion
SA	E_0 is attained at the end of the annealing process.
SA-CE	E_0 is attained at any Monte Carlo sweep during the annealing.
SQA-CE	E_0 is attained by any Trotter slice at the end of any Monte Carlo sweep.
A-PSA	E_0 is attained by any replica at the end of the annealing process.
A-PSA-CE	E_0 is attained by any replica at the end of any Monte Carlo sweep.

where $\mathcal{I}_A(x)$ is the indicator function and $E_{\text{res}}^k(\mathcal{S})$ is the residual energy for the k th trial and solver $\mathcal{S}$:

$$E_{\text{res}}(\mathcal{S}) := \begin{cases} E(t_J) - E_0, & \mathcal{S} = \text{SA}, \\ \min_{j=1, \dots, J} E(t_j) - E_0, & \mathcal{S} = \text{SA-CE}, \\ \min_{m=1, \dots, M} E^{(m)}(t_j) - E_0, & \mathcal{S} = \text{SQA-CE}, \\ \min_{j=1, \dots, J} E^{(r)}(t_j) - E_0, & \mathcal{S} = \text{A-PSA}, \\ \min_{r=1, \dots, R} E^{(r)}(t_j) - E_0, & \mathcal{S} = \text{A-PSA-CE}, \end{cases} \quad (17)$$

where the superscripts (m) and (r) denote the m th Trotter slice and r th replica, respectively, t_j denotes the j th Monte Carlo sweep, J is the number of Monte Carlo sweeps, M is the number of Trotter slices, and R is the number of replicas.

It is worth noting that, as some of the parameters affect not only the annealing process but also the computational load of the algorithm, the metric that is typically used to compare the performance of different annealing algorithms is the time to solution, a composite metric representing the total computation time required to obtain the optimal solution with probability p_R , defined as follows:

$$\text{TTS}(\mathcal{S}, p_R) := \tau(\mathcal{S}) \frac{\ln(1 - p_R)}{\ln(1 - p_s(\mathcal{S}))}, \quad (18)$$

where $\tau(\mathcal{S})$ is the computation time for a single annealing process. This metric does not inherently impose any constraints on the computational budget and is usually computed after the parameter values have been set. As during the optimization process it is necessary to ensure a fair and controlled comparison between parameter configurations, we first fixed the computational budget, shared among algorithms, and then ran the Bayesian optimizers. In this context, the computational budget is represented by the total number of Monte Carlo steps (MCS) per spin [35] performed by the algorithm, i.e., a normalized budget with respect to the total number of spins. Specifically, this is equal to the number of Monte Carlo sweeps for SA, the product between the number of Trotter slices and the number of Monte Carlo sweeps for SQA, and the product between the number of replicas and the number of Monte Carlo sweeps for A-PSA.

The global pipeline is described in Algorithm 1. The first section shows the optimization phase to find optimal parameters, i.e., θ_{opt} for solvers fixing the budget; the second section shows the evaluation phase using different computational budgets for fixed solvers parameters determined by the first phase. θ_{opt} describes only the solver parameters; instead, the optimizer parameters are assigned their default values.

The encoding phase, formalized by Eq. (4), requires a dedicated subroutine that we show in Algorithm 2. The result is the Ising model for the different instances of N analyzed in the current study.

After fixing the computational budget, we proceeded with the optimization of the remaining parameters, each of which required a dedicated approach. The SA algorithm is governed by the

ALGORITHM 1. Meta-heuristic factorization with Bayesian optimization.

```

Input:  $N, a_{\max}, \text{optimizer\_params}, \text{solver\_params}$ 
Input:  $\text{optimizers} \leftarrow [\text{tpe}, \text{gp}]$ 
Input:  $\text{solvers} \leftarrow [\text{sa}, \text{sa\_ce}, \text{sqa\_ce}, \text{apsa}, \text{apsa\_ce}]$ 
1:  $a \leftarrow 0$ 
2:  $\text{global\_max\_ps} \leftarrow 0$ 
3: do
4:    $\text{model} \leftarrow \text{ENCODEISINGMODEL}(N, a)$ 
5:   for all  $\text{optimizer} \in \text{optimizers}$  then
6:     for all  $\text{solver} \in \text{solvers}$  do
7:        $(\theta_{\text{opt}}, \text{PS}_{\text{opt}}) \leftarrow \text{OPTIMIZE}(\text{model}, \text{optimizer}, \text{optimizer\_params}, \text{solver}, \text{solver\_params})$ 
8:        $\text{best\_params}[\text{optimizer}][\text{solver}] \leftarrow \theta_{\text{opt}}$ 
9:       if  $\text{PS}_{\text{opt}} > \text{global\_max\_ps}$  then
10:         $\text{global\_max\_ps} \leftarrow \text{PS}_{\text{opt}}$ 
11:      end if
12:    end for
13:  end for
14:   $a \leftarrow a + 1$ 
15: while  $a < a_{\max}$  and  $\text{global\_max\_ps} = 0$  then
16:   $\text{global\_max\_ps} > 0$ 
17:    for all  $\text{optimizer} \in \text{optimizers}$  do
18:      for all  $\text{solver} \in \text{solvers}$  do
19:        for all  $\text{mcs} \in \text{mcs\_list}$  do
20:           $(\text{PS}, \text{TTS}, \epsilon_{\text{res}}) \leftarrow \text{COMPUTEMETRICS}(\text{model}, \text{best\_params}[\text{optimizer}][\text{solver}], \text{solver})$ 
21:        end for
22:      end for
23:    end for
24:  end if

```

ALGORITHM 2. Encode Ising model (N, a) .

```

Input:  $N, a$ 
Output:  $\text{ISINGMODEL}(h, J, C)$   $\triangleright C := \text{Ising offset}$ 
1:  $n \leftarrow \lceil \log_2 N \rceil$ 
2: if  $n \pmod 2 = 0$  then
3:    $n_p \leftarrow n/2 - a$ 
4:    $n_q \leftarrow n/2 + 1 + a$ 
5: else
6:    $n_p \leftarrow \lfloor n/2 \rfloor + 1 - a$ 
7:    $n_q \leftarrow \lfloor n/2 \rfloor + 1 + a$ 
8: end if
9: Define binary vectors  $x_p \in \{0, 1\}^{n_p-1}, x_q \in \{0, 1\}^{n_q-1}$ 
10:  $p(x_p) \leftarrow \sum_{i=1}^{n_p-1} 2^i x_p^{(i)} + 1$ 
11:  $q(x_q) \leftarrow \sum_{i=1}^{n_q-1} 2^i x_q^{(i)} + 1$ 
12:  $F_{\text{HUBO}}(x) \leftarrow (N - p(x_p) \cdot q(x_q))^2$   $\triangleright x := [x_p, x_q]$ 
13:  $(Q, c) \leftarrow \text{HUBOTOQUBO}(F_{\text{HUBO}})$   $\triangleright c := \text{QUBO offset}$ 
14:  $(h, J, C) \leftarrow \text{QUBOTOISING}(Q, c)$ 
15: return  $\text{ISINGMODEL}(h, J, C)$ 

```

β schedule, whose functional form was fixed to be geometric, while its boundary values were estimated, for each number to be factorized, following the approach implemented in D-Wave's neal library [47]. More precisely, to allow the system to explore widely separated regions of the configuration space in the initial annealing phase, $\beta_{\min}$ was chosen such that

$$p_i^{\text{flip}}(\beta_{\min}) = e^{-\beta_{\min} \Delta E_i^{\text{flip}}} \geq 0.5, \quad \forall i, \quad (19)$$

where ΔE_i^{flip} is the variation in the system energy due to the spin flip, namely,

$$\Delta E_i^{\text{flip}} = 2s_i \left(h_i + \sum_j J_{ij} s_j \right) = 2h_i^{\text{eff}} s_i, \quad (20)$$

and

$$h_i^{\text{eff}} = h_i + \sum_j J_{ij} s_j \quad (21)$$

is the effective field for spin i .

Introducing

$$|h^{\text{eff}}|_{\max} := \max_i \left(|h_i| + \sum_j |J_{ij}| \right), \quad (22)$$

the maximum energy variation resulting from a spin flip can be approximated by

$$\Delta E_{\max}^{\text{flip}} \approx 2|h^{\text{eff}}|_{\max}. \quad (23)$$

Thus, the condition given by Eq. (19) can be rewritten as

$$e^{-\beta_{\min} \Delta E_{\max}^{\text{flip}}} = 0.5, \quad (24)$$

from which $\beta_{\min}$ can be calculated.

At the end of the annealing process, instead, to favor the convergence of the system toward the configuration corresponding to a local or global energy minimum, the value of $\beta_{\max}$ is estimated from the following condition on the total excitation rate:

$$\sum_i p_i^{\text{flip}}(\beta_{\min}) = \sum_i e^{-\beta_{\min} \Delta E_i^{\text{flip}}} \leq 0.01. \quad (25)$$

Defining

$$|h|_{\min} := \min_{i,j} \{|h_i|, |J_{ij}|\}, \quad (26)$$

the minimum energy variation resulting from a spin flip can be approximated by

$$|\Delta E^{\text{flip}}|_{\min} \approx 2|h^{\text{eff}}|_{\min}. \quad (27)$$

Considering only the contribution deriving from the spins associated with the minimum effective field and denoting with $\chi_{\min}$ the number of such spins, we obtain the following relation:

$$\sum_i p_i^{\text{flip}}(\beta_{\min}) \approx \chi_{\min} e^{-\beta_{\min} |\Delta E^{\text{flip}}|_{\min}} \leq 0.01, \quad (28)$$

from which the value of $\beta_{\max}$ can be retrieved.

Such estimate was further refined performing Bayesian optimization within the range $[\max\{\tilde{\beta}_{\min}, \tilde{\beta}_{\max}/10\}, 10\tilde{\beta}_{\max}]$, where $\tilde{\beta}_{\max}$ and $\tilde{\beta}_{\min}$ denote the parameter values estimated with the method derived from D-Wave.

The intermediate values of β were selected according to a geometric annealing schedule.

In contrast to SA, all SQA parameters, i.e., M , β , and Γ , were optimized using a Bayesian approach. Since the optimization had to be performed for different Ising systems, each of which

corresponding to the integer to be factorized, we introduced

$$\widetilde{MT} := \frac{MT}{|J|}, \quad \widetilde{\Gamma} := \frac{\Gamma}{|J|}, \quad (29)$$

where $|J|$, representing the typical spin coupling scale, was defined as

$$|J| := \text{median}\{|J_{ij}| : J_{ij} \neq 0, i, j = 1, \dots, d\}. \quad (30)$$

This choice is motivated by the high degree of sparsity characterizing the matrix J .

Bayesian optimizers were then applied to $\widetilde{MT}$, $\widetilde{\Gamma}$, and M , yielding the optimal set of SQA parameters. This approach was guided by Ref. [17], which analyzed how the interplay between Γ , MT , and the spin coupling scale affects the performance of path-integral Monte Carlo. The functional form of the annealing schedule was fixed to be quartic.

Finally, the β schedule for the A-PSA algorithm was determined in the same way as the SA algorithm. The number of replicas, instead, was set equal to the number of Trotter slices provided by the Bayesian optimization, as A-PSA was primarily studied to compare the performance of SQA to a straightforward SA replication using the same computational budget as SQA.

Given the computational cost to evaluate the objective function, parameter optimization was carried out relying on two sequential model-based global optimization algorithms: the Gaussian-process (GP)-based surrogate model, whose implementation was provided by the scikit-optimize [48] Python library, and the tree-structured Parzen estimator (TPE), whose implementation was provided by the Optuna [49] Python library. For both approaches, parameter search was based on the optimization of the expected improvement criterion [73] for the probability of success metric. The use of both methods is motivated by the fact that GP-based optimization provides a well-established baseline, whereas TPE is particularly well suited for high-dimensional parameter spaces under limited evaluation budgets.

The maximum values of p_s obtained for each optimizer and solver, for all integers N , are summarized in Table III. Each value of p_s was estimated from a sample of runs of size n_{reads} and achieved within 100 optimizer iterations. An extended version of the table, including the computational budget allocated to each solver and the corresponding optimal parameter values, is provided in Appendix D.

The SA and SA-CE solvers generally yield the lowest probabilities of success, while optimal values of p_s exhibit an overall decreasing trend across solvers as N increases. Notable exceptions to this trend are observed for specific instances, such as $N = 59\,989$, $1\,028\,171$, $8\,219\,999$, which display substantially higher values of p_s compared to semiprimes of comparable size. In the case of $N = 8\,219\,999$, this behavior can be partially attributed to the strong imbalance between its factors; however, this explanation does not apply to the remaining two instances, which correspond to balanced factorizations. The observed anomalies in the trend of p_s may indicate that $|J|$ does not accurately represent the coupling scale for all N or that the functional forms of the annealing schedules are suboptimal, pointing to potential avenues for further optimization.

IV. SIMULATION RESULTS

All simulations were run on a Google Cloud server equipped with Linux VM, Jupyter Python Environment, and Intel Xeon Gold 6242 @ 2.80 GHz CPU, 128 GB RAM. In this section, we show the simulation results for the cases $N = 59\,989$ and $N = 8\,219\,999$. The parameters and results are shown for the optimization phase, i.e., the phase where optimizers determine optimal values for solver parameters with fixed computational budget expressed in Monte Carlo steps, and for the evaluation phase, i.e., where the computational budget is variable in a given range allowing the evaluation of the different solvers using the optimal parameters. In this section, the term energy refers to the residual energy as defined in Eq. (17), and the global minimum is obtained when the energy is zero. The simulation results can be computed using the code package available as the Supplemental Material [74].

TABLE III. Optimal values of p_s obtained for each value of N , solver, and optimizer. For each optimizer, the maximum value of p_s is highlighted in bold.

N	n_N	n_{reads}	Optimizer	SA	SA-CE	SQA-CE	A-PSA	A-PSA-CE
15	4	100	TPE	1.00	1.00	1.00	1.00	1.00
			GP	1.00	1.00	1.00	1.00	1.00
21	5	100	TPE	0.74	1.00	1.00	1.00	1.00
			GP	0.81	1.00	1.00	1.00	1.00
35	6	100	TPE	0.32	1.00	1.00	1.00	1.00
			GP	0.32	1.00	1.00	1.00	1.00
57	6	100	TPE	0.52	1.00	1.00	1.00	1.00
			GP	0.46	1.00	1.00	1.00	1.00
77	7	100	TPE	0.22	1.00	1.00	1.00	1.00
			GP	0.19	1.00	1.00	1.00	1.00
123	7	100	TPE	0.31	1.00	1.00	1.00	1.00
			GP	0.29	1.00	1.00	1.00	1.00
143	8	100	TPE	0.12	1.00	1.00	1.00	1.00
			GP	0.12	1.00	1.00	1.00	1.00
253	8	100	TPE	0.14	1.00	1.00	1.00	1.00
			GP	0.14	1.00	1.00	1.00	1.00
451	9	100	TPE	0.10	0.71	1.00	0.68	0.94
			GP	0.10	0.66	1.00	0.76	0.91
511	9	100	TPE	0.14	0.94	1.00	1.00	1.00
			GP	0.11	0.96	1.00	0.99	1.00
899	10	100	TPE	0.09	0.32	1.00	1.00	1.00
			GP	0.11	0.28	1.00	1.00	1.00
1011	10	100	TPE	0.09	0.97	1.00	1.00	1.00
			GP	0.07	0.98	1.00	1.00	1.00
1769	11	100	TPE	0.06	0.39	1.00	1.00	1.00
			GP	0.05	0.37	1.00	1.00	1.00
2047	11	100	TPE	0.06	0.30	1.00	0.78	0.80
			GP	0.05	0.33	1.00	0.63	0.85
4009	12	100	TPE	0.02	0.11	0.55	0.28	0.43
			GP	0.03	0.13	0.54	0.30	0.49
4087	12	100	TPE	0.04	0.17	0.48	0.22	0.38
			GP	0.05	0.17	0.50	0.25	0.38
5723	13	100	TPE	0.02	0.06	0.37	0.22	0.25
			GP	0.02	0.10	0.37	0.29	0.28
8189	13	100	TPE	0.02	0.06	0.39	0.27	0.35
			GP	0.02	0.07	0.36	0.50	0.54
15 853	14	100	TPE	0.01	0.04	0.31	0.36	0.39
			GP	0.02	0.04	0.24	0.35	0.32
16 379	14	100	TPE	0.01	0.07	0.42	0.09	0.18
			GP	0.02	0.07	0.40	0.11	0.17
21 733	15	100	TPE	0.01	0.02	0.09	0.13	0.10
			GP	0.02	0.02	0.09	0.13	0.16
32 743	15	100	TPE	0.02	0.04	0.02	0.03	0.04
			GP	0.02	0.05	0.04	0.05	0.05

TABLE III. (Continued.)

N	n_N	n_{reads}	Optimizer	SA	SA-CE	SQA-CE	A-PSA	A-PSA-CE
49 583	16	100	TPE	0.00	0.02	0.02	0.02	0.02
			GP	0.01	0.03	0.02	0.02	0.02
59 989	16	100	TPE	0.02	0.04	0.88	0.92	0.97
			GP	0.02	0.02	0.86	0.96	0.97
65 509	16	100	TPE	0.01	0.02	0.02	0.02	0.02
			GP	0.01	0.02	0.02	0.01	0.02
94 037	17	100	TPE	0.01	0.01	0.03	0.06	0.04
			GP	0.01	0.01	0.03	0.05	0.06
131 069	17	100	TPE	0.00	0.01	0.01	0.01	0.01
			GP	0.00	0.01	0.01	0.01	0.01
159 317	18	100	TPE	0.01	0.02	0.09	0.11	0.14
			GP	0.00	0.01	0.08	0.11	0.11
262 099	18	100	TPE	0.01	0.01	0.05	0.05	0.05
			GP	0.01	0.01	0.03	0.06	0.05
374 599	19	100	TPE	0.00	0.00	0.01	0.01	0.01
			GP	0.00	0.01	0.00	0.00	0.01
376 289	19	100	TPE	0.01	0.00	0.01	0.02	0.02
			GP	0.01	0.00	0.01	0.02	0.02
524 281	19	100	TPE	0.01	0.01	0.01	0.00	0.01
			GP	0.00	0.01	0.00	0.01	0.00
881 851	20	100	TPE	0.00	0.00	0.01	0.02	0.01
			GP	0.01	0.00	0.01	0.01	0.01
1 005 973	20	100	TPE	0.01	0.00	0.02	0.01	0.02
			GP	0.00	0.00	0.01	0.02	0.02
1 028 171	20	100	TPE	0.01	0.01	0.09	0.12	0.13
			GP	0.01	0.01	0.14	0.12	0.13
1 048 561	20	100	TPE	0.01	0.01	0.00	0.01	0.01
			GP	0.01	0.00	0.01	0.02	0.02
1 794 953	21	100	TPE	0.00	0.01	0.01	0.01	0.01
			GP	0.01	0.00	0.01	0.01	0.01
2 093 809	21	100	TPE	0.00	0.00	0.00	0.00	0.01
			GP	0.00	0.01	0.00	0.00	0.00
		200	TPE	0.00	0.00	0.005	0.005	0.005
			GP	0.00	0.01	0.00	0.00	0.005
3 819 047	22	100	TPE	0.00	0.00	0.01	0.02	0.02
			GP	0.00	0.00	0.01	0.01	0.01
		100	TPE	0.00	0.01	0.20	0.11	0.12
			GP	0.01	0.01	0.17	0.10	0.11
8 219 999	23	400	TPE	0.005	0.005	0.1525	0.10	0.0875
			GP	0.000	0.0025	0.155	0.095	0.09

A. The case of $N = 59\,989$

For this case, the parameters used for the optimization phase are reported in Table IV, where the search space for the parameters of different solvers is described.

TABLE IV. Configuration parameters for the optimization phase for $N = 59\,989$.

Parameter	Value
encoding.N	59 989
encoding.max_factor_extra_bits	4
encoding.factor_extra_bits	0
optimization.num_mc_steps	20 000
optimization.optimizer.num_trials	100
optimization.sampler.num_reads	100
optimization.ca_params.static.beta_min	4.1×10^{-11}
optimization.ca_params.static.num_mc_sweeps_per_beta	1
optimization.ca_params.optimized.beta_max	[0.075, 7.49]
optimization.sqa_params.static.num_mc_sweeps_per_gamma	1
optimization.sqa_params.optimized.num_trotter_slices	[200, 2000]
optimization.sqa_params.optimized.gamma_par	[1.0, 5.0]
optimization.sqa_params.optimized.MT_par	[1.0, 5.0]

The optimization phase gives the optimal parameter results described in Table V for the different solvers.

The probability of success p_s reaches its maximum value for the A-PSA-CE solver, followed by SQA-CE and SA. The optimal parameters are kept fixed in the following phase, where the computational budget is variable. The parameters used for the evaluation phase are reported in Table VI.

In order to illustrate the metric results, we select the evaluation phase corresponding to the optimal parameters provided by the GP Optimizer. Overall, we will focus on comparing the performance of SQA and A-PSA, as they exhibit similar behavior and consistently outperform the SA solver. The first metric we analyze is connected to the frequency of obtaining the global minimum during the annealing process. For this analysis, the computational budget is the same as that reported in Table IV. In this case, the relevant quantities for statistics are the number of 100 reads, the optimal number of Trotter slices/replicas: 1249, and the optimal number of sweeps: 16. A-PSA inherits the number of replicas from the number of optimal Trotter slices. The product of these three numbers gives the total number of energy evaluations of the system that are recorded by

TABLE V. Optimal parameter values obtained for $N = 59\,989$, optimizer, and solver. For each optimizer, the maximum value of p_s is highlighted in bold.

N	n	n_{reads}	Optimizer	$\mathcal{S}$	MCS	p_s	$\beta_{\min}$	$\beta_{\max}$	R	M	β	Γ
59 989	16	100	GP	SA	20 000	0.02	4.1×10^{-11}	6.8	–	–	–	–
				SA-CE	20 000	0.02	4.1×10^{-11}	7.5	–	–	–	–
				SQA-CE	19 984	0.86	–	–	–	1249	2.7×10^{-3}	8.2×10^5
				A-PSA	19 984	0.96	4.1×10^{-11}	7.5×10^{-2}	1249	–	–	–
				A-PSA-CE	19 984	0.97	4.1×10^{-11}	7.2	1249	–	–	–
			TPE	SA	20 000	0.02	4.1×10^{-11}	5.2	–	–	–	–
				SA-CE	20 000	0.04	4.1×10^{-11}	3.3	–	–	–	–
				SQA-CE	18 760	0.88	–	–	–	1340	2.4×10^{-3}	5.9×10^5
				A-PSA	18 760	0.92	4.1×10^{-11}	6	1340	–	–	–
				A-PSA-CE	18 760	0.97	4.1×10^{-11}	5.5	1340	–	–	–


TABLE VI. Configuration parameters for the evaluation phase for $N = 59\,989$.

Parameter	Value
evaluation.num_mc_step_list.min	6000
evaluation.num_mc_step_list.max	60 000
evaluation.num_mc_step_list.num	10
evaluation.metrics.cl	0.95
evaluation.metrics.p_r	0.99
evaluation.metrics.n_resamples	1000
evaluation.sampler.num_reads_ts	100
evaluation.sampler.num_reads_perf	100

the software pipeline. The global minimum is found 913 times in the case of the SQA-CE solver. This represents a fraction of $f_{\text{minglob}} = 4.6 \times 10^{-4}$ of the total number of evaluations. This fraction should not be confused with the value of p_s reported in Table V, because we should remember the definition of success given in Table II. According to this definition, $p_s \geq f_{\text{minglob}}$, because for CE solvers one success can correspond to many different findings of the global minimum during the annealing. The global minimum is found 2187 times in the case of the A-PSA-CE solver. It is clear that these occurrences are distributed between different reads and different Trotter slices. For each read and each Trotter slice, we normalize the number of times that the global minimum is found by the number of sweeps, i.e., the total number of updates of energy during the annealing. In this way, we compute a frequency to find the global minimum. Then, we calculate the mean value of these frequencies over the Trotter slices for each read, noting that such means can have a different basis depending on different reads. These mean values are reported in Fig. 2. The values are higher for the A-PSA-CE solver relative to the SQA-CE solver.


At this point, we select different reads for each solver in order to illustrate the energy dynamics during the annealing of the system. For the case of SQA-CE exploration, we select read 0 and read 9 corresponding to frequency mean of 43.75 and 62.5, respectively. The energy dynamics is shown in Fig. 3. The energy dynamics shows boxplots computed on Trotter slices and energy is plotted with +1 bias in order to compute a log scale. The panels are coherent with zero-energy frequencies

Zero-energy frequency mean per read (N=59989), SQA-CE




(a) SQA-CE

Zero-energy frequency mean per read (N=59989), A-PSA-CE




(b) A-PSA-CE

FIG. 2. Zero-energy frequency mean per read for SQA-CE and A-PSA-CE solvers for $N = 59\,989$. The results for the SQA-CE solver are shown in panel (a) and for the A-PSA-CE solver in panel (b).



(a) Energy of the system for SQA-CE solver in the read 0.




(b) Energy of the system for SQA-CE solver in the read 9.


FIG. 3. Energy of the system explored by SQA-CE solver for $N = 59989$. The boxplot's statistics are computed over the different Trotter slices. The results corresponding to read 0 are shown in panel (a) and to read 9 in panel (b).

because the extreme values of boxplots reach the global minimum more in Fig. 3(b) compared to Fig. 3(a). The annealing process is capable to reach the global minimum as extreme energy values of Trotter slices. The energy ranges of the different Trotter slices are quite stable, as displayed by the bulk of the boxplots, demonstrating an exploration of the energy landscape in a quantum freezing regime [35,75], because different Trotter slices explore and tunnel between different energy valleys until the very end of the annealing.

For the case of A-PSA-CE exploration, we select read 3 and read 12 corresponding to mean frequencies equal to 56.25 and 68.75, respectively. The energy dynamics is shown in Fig. 4. The energy dynamics shows boxplots computed on replica slices and energy is plotted with +1 bias in order to compute a log scale. The panels are coherent with zero-energy frequencies because the extreme values of boxplots reach the global minimum more in Fig. 4(b) than in Fig. 4(a). The



(a) Energy of the system for A-PSA-CE solver in the read 3.



(b) Energy of the system for A-PSA-CE solver in the read 12.

FIG. 4. Energy of the system explored by A-PSA-CE solver for $N = 59\,989$. The boxplot's statistics are computed over the different replica slices. The results corresponding to read 3 are shown in panel (a) and to read 12 in panel (b).

annealing process is capable of reaching the global minimum as extreme energy values of replica slices. The energy ranges of the different replica slices are quite stable, as displayed by the bulk of the boxplots, demonstrating an exploration of the classical energy landscape similar to that of a quantum freezing regime.

To characterize the performance scaling of each solver as a function of the allocated computational budget, the annealing process was run 100 times for 10 different values of the total Monte Carlo steps. Three performance metrics were calculated: the probability of success p_s , the time to solution, and the mean of the minimum residual energy attained, as shown in Fig. 5.

The computational budgets were selected in the range [5000, 50 000] following a geometric progression. For each of the considered budgets, the annealing processes were performed using the optimal parameter set identified for each solver by the GP Optimizer.




FIG. 5. Performance-scaling metrics for $N = 59989$ computed from samples of size 100. Missing data points in panels (a) and (b) for the SA and SA-CE solvers correspond to samples for which $p_s = 0$. Error bars represent the 95%-confidence intervals estimated using the Wilson score interval [panel (a)], bootstrap resampling with 1000 resamples [panel (b)], and Student's t distribution [panel (c)]. The absence of error bars in panel (b) for the SA and SA-CE solvers results from the failure of the bootstrap interval estimation due to vanishing p_s .

The SQA-CE, A-PSA, and A-PSA-CE solvers exhibit similar behavior across all three performance metrics and demonstrate vastly superior performance compared to SA and SA-CE. In particular, the probability of success increases with the allocated computational budget, accompanied by a corresponding decrease in the mean residual energy. In spite of this improvement in p_s , the TTS shows an increasing trend, indicating that the additional computational cost required to run the simulations outweighs the enhanced ability of the solver to reach the global minimum.

It is worth noting that, since the parameter values for each solver were held fixed and set to those provided by the GP-based optimization, for computational budgets different from the one used during the optimization phase, the employed parameters might no longer be optimal. Furthermore, for the SQA-CE, A-PSA, and A-PSA-CE solvers, variations in the computational budget correspond exclusively to changes in the number of Monte Carlo sweeps, as the number of Trotter slices and replicas was kept constant. We show details of schedules used for the different solvers in Fig. 6.

B. The case of $N = 8\,219\,999$

For this case, the parameters used for the optimization phase are reported in Table VII, where the search space for the parameters of different solvers is described. We consider two different values




FIG. 6. Annealing schedules employed by each solver for $N = 59\,989$. Panel (a) describes the case of SA solver. Panel (b) describes the case of SA-CE solver. Panel (c) describes the case of SQA-CE solver. Panel (d) describes the case of A-PSA solver. Panel (e) describes the case of A-PSA-CE solver.

for the number of reads, i.e., 100 and 400, to obtain more precision regarding the estimate of the probability of success in the optimization phase.

The optimization phase gives the optimal parameter results described in Table VIII for the different solvers.

The probability of success p_s reaches the maximum value for the SQA-CE solver, followed by A-PSA and SA. The optimal parameters are kept fixed in the following phase, where the computational budget is variable. The parameters used for the evaluation phase are reported in Table IX.

TABLE VII. Configuration parameters for the optimization phase for $N = 8\,219\,999$.

Parameter	Value
encoding.N	8 219 999
encoding.max_factor_extra_bits	3
encoding.min_factor_extra_bits	0–3
encoding.factor_extra_bits	3
optimization.num_mc_steps	20 000
optimization.optimizer.num_trials	100
optimization.sampler.num_reads	100–400
optimization.ca_params.static.beta_min	2.5×10^{-15}
optimization.ca_params.static.num_mc_sweeps_per_beta	1
optimization.ca_params.optimized.beta_max	[0.075, 7.5]
optimization.sqa_params.static.num_mc_sweeps_per_gamma	1
optimization.sqa_params.optimized.num_trotter_slices	[200, 2000]
optimization.sqa_params.optimized.gamma_par	[1.0, 5.0]
optimization.sqa_params.optimized.MT_par	[1.0, 5.0]

TABLE VIII. Optimal parameter values obtained for $N = 8\,219\,999$, optimizer, and solver. For each optimizer, the maximum value of p_s is highlighted in bold.

N	n	n_{reads}	Optimizer	$\mathcal{S}$	MCS	p_s	$\beta_{\min}$	$\beta_{\max}$	R	M	β	Γ
8219999	23	100	GP	SA	20 000	0.01	2.5×10^{-15}	7.5	—	—	—	—
				SA-CE	20 000	0.01	2.5×10^{-15}	7.6×10^{-2}	—	—	—	—
				SQA-CE	18 914	0.17	—	—	—	1351	2.3×10^{-5}	8.4×10^7
				A-PSA	18 914	0.10	2.5×10^{-15}	1.4	1351	—	—	—
				A-PSA-CE	18 914	0.11	2.5×10^{-15}	4.3	1351	—	—	—
			TPE	SA	20 000	0.00	2.5×10^{-15}	4.8	—	—	—	—
				SA-CE	20 000	0.01	2.5×10^{-15}	3.9	—	—	—	—
				SQA-CE	18 928	0.20	—	—	—	1183	3.1×10^{-5}	8.1×10^7
				A-PSA	18 928	0.11	2.5×10^{-15}	5.3	1183	—	—	—
				A-PSA-CE	18 928	0.12	2.5×10^{-15}	1.4	1183	—	—	—
	400	400	GP	SA	20 000	0.00	2.5×10^{-15}	2.6	—	—	—	—
				SA-CE	20 000	0.0025	2.5×10^{-15}	7.5	—	—	—	—
				SQA-CE	19 012	0.155	—	—	—	1358	1.6×10^{-5}	1.8×10^7
				A-PSA	19 012	0.095	2.5×10^{-15}	2.4	1358	—	—	—
				A-PSA-CE	19 012	0.090	2.5×10^{-15}	5.2	1358	—	—	—
			TPE	SA	20 000	0.005	2.5×10^{-15}	5.6	—	—	—	—
				SA-CE	20 000	0.005	2.5×10^{-15}	6.4	—	—	—	—
				SQA-CE	19 824	0.1525	—	—	—	1239	1.5×10^{-5}	7.3×10^7
				A-PSA	19 824	0.10	2.5×10^{-15}	5.4	1239	—	—	—
				A-PSA-CE	19 824	0.0875	2.5×10^{-15}	1.5	1239	—	—	—

In order to illustrate the metric results, we select the evaluation phase corresponding to the optimal parameters identified by the TPE optimizer. Overall, we will focus on comparing the performance of SQA and A-PSA, as they exhibit similar behavior and consistently outperform the SA solver. It is important to stress that also in this case the SA solver can provide a positive probability of success. The first metric we analyze is connected to the frequency of obtaining the global minimum during the annealing process. For this analysis, the computational budget is the same as that reported in Table VII. In this case, the relevant quantities for statistics are the number of 400 reads, the optimal number of Trotter slices/replicas: 1239, and the optimal number of sweeps: 16. A-PSA inherits the number of replicas from the number of optimal Trotter slices. The product of these three numbers gives the total number of energy evaluations of the system that are recorded by the software pipeline. The global minimum is found 354 times in the case of SQA-CE solver. The global minimum is found 185 times in the case of A-PSA-CE solver. It is clear that these occurrences

TABLE IX. Configuration parameters for the evaluation phase for $N = 8\,219\,999$.

Parameter	Value
evaluation.num_mc_step_list.min	6000
evaluation.num_mc_step_list.max	60 000
evaluation.num_mc_step_list.num	10
evaluation.metrics.cl	0.95
evaluation.metrics.p_r	0.99
evaluation.metrics.n_resamples	1000
evaluation.sampler.num_reads_ts	400
evaluation.sampler.num_reads_perf	400





FIG. 7. Zero-energy frequency mean per read for SQA-CE and A-PSA-CE solvers for $N = 8\,219\,999$. The results for the SQA-CE solver are shown in panel (a) and for the A-PSA-CE solver in panel (b).

are distributed between different reads and different Trotter slices. For each read and each Trotter slice, we normalize the number of times that the global minimum is found by the number of sweeps, i.e., the total number of updates of energy during the annealing. In this way, we compute a frequency to find the global minimum. Then, we calculate the mean value of these frequencies over the Trotter slices for each read; e.g., this mean can have a different basis depending on different reads. These mean values are reported in Fig. 7. The maximum values are higher for the SQA-CE solver with respect to the A-PSA-CE solver.


At this point, we select different reads for each solver in order to illustrate the energy dynamics during the annealing of the system. For the case of SQA-CE exploration, we select read 2 and read 158 corresponding to frequency mean of 56.25 and 68.75, respectively. The energy dynamics is shown in Fig. 8. The energy dynamics shows boxplots computed on Trotter slices and energy is plotted with +1 bias in order to compute a log scale. The panels are coherent with zero-energy frequencies because the extreme values of boxplots reach the global minimum more in Fig. 8(b) compared to Fig. 8(a). The annealing process is capable to reach the global minimum as extreme energy values of Trotter slices. The energy ranges of the different Trotter slices are quite stable, as shown by most of the boxplots, demonstrating an exploration of the energy landscape in a quantum freezing regime [35,75].

For the case of A-PSA-CE exploration, we select read 65 and read 4 corresponding to frequency mean of 56.25 and 62.5, respectively. The energy dynamics is shown in Fig. 9. The energy dynamics shows boxplots computed on replica slices and energy is plotted with +1 bias in order to compute a log scale. The panels are coherent with zero-energy frequencies because the extreme values of boxplots reach the global minimum more in Fig. 9(b) compared to Fig. 9(a). The annealing process is capable to reach the global minimum as extreme energy values of replica slices. The energy ranges of the different replica slices are quite stable, as displayed by the bulk of the boxplots, demonstrating an exploration of the classical energy landscape similar to that of a quantum freezing regime.

The performance-scaling analysis for $N = 8\,219\,999$ was conducted following the same procedure described for $N = 59\,989$, as shown in Fig. 10, the sole differences being the use of a sample of size 400 and the choice of the computational budget range [6000, 60 000]. The parameters for each solver were fixed to the optimal values retrieved by the TPE-based optimization. We show details of schedules used for the different solvers in Fig. 11.



(a) Energy of the system for SQA-CE solver in the read 2.




(b) Energy of the system for SQA-CE solver in the read 158.


FIG. 8. Energy of the system explored by SQA-CE solver for $N = 8\,219\,999$. The boxplot's statistics are computed over the different Trotter slices. The results corresponding to read 2 are shown in panel (a) and to read 158 in panel (b).

Across all three computed metrics, the solvers display behavior analogous to that observed for $N = 59\,989$. Notably, increasing N from 59 989 to 8 219 999 results in a reduction of p_s by approximately one order of magnitude, while the TTS increases by roughly 2 orders of magnitude, reflecting the growing complexity of the factorization problem. Furthermore, for computational budgets exceeding 10 000 Monte Carlo steps, SQA-CE emerges as the best performing solver in terms of TTS.

We further observe that increasing the computational budget beyond 20 000 Monte Carlo steps does not result in a substantial improvement in p_s , which saturates at the optimal value achieved by the TPE optimizer, i.e., $p_s \approx 0.15$. With the fully quantum approach described in Ref. [62], the factorization of $N = 8\,219\,999$ was achieved in 329 runs out of a sample of 1000, corresponding to a probability of success that is about twice the optimal value obtained with the approach considered



(a) Energy of the system for A-PSA-CE solver in the read 65.



(b) Energy of the system for A-PSA-CE solver in the read 4.

FIG. 9. Energy of the system explored by A-PSA-CE solver for $N = 8\,219\,999$. The boxplot's statistics are computed over the different replica slices. The results corresponding to read 65 are shown in panel (a) and to read 4 in panel (b).

here. While this comparison should be interpreted with caution, given the differing samples sizes, the different architectures of the two approaches, and the dependence of the optimal parameter values on the allocated computational budget, it nevertheless suggests that quantum-inspired methods can provide a viable classical approximation of quantum annealers for this problem class [76–79].

V. CONCLUSIONS

A central outcome of our investigation is that the relative hierarchy of success probabilities is *instance dependent*, even after Bayesian optimization of solver parameters. During the optimization phase, for the medium-sized instance $N = 59\,989$ we observe $\text{A-PSA} > \text{SQA} \gg \text{SA}$, whereas for




FIG. 10. Performance-scaling metrics for $N = 8\,219\,999$ computed from samples of size 400. Missing data points in panels (a) and (b) for the SA and SA-CE solvers correspond to samples for which $p_s = 0$. Error bars represent the 95%-confidence intervals estimated using the Wilson score interval [panel (a)], bootstrap resampling with 1000 resamples [panel (b)], and Student's t distribution [panel (c)]. The absence of error bars in panel (b) for the SA and SA-CE solvers results from the failure of the bootstrap interval estimation due to vanishing p_s .

the larger instance $N = 8\,219\,999$ the hierarchy reverses to $SQA > A\text{-PSA} \gg SA$. From a replication standpoint, this instance dependence is itself an important benchmark feature: It highlights that solver comparisons based on a single instance (or a narrow regime of hyperparameters) can be misleading, and that reproducible claims require explicit reporting of the instance family, computational budget, and tuning protocol. Our A-PSA baseline is an *asynchronous multistart* strategy: multiple SA trajectories evolve independently with no inter-replica communication or coupling. Consequently, its dynamics are qualitatively different from parallel-tempering [80] and from interacting-replica approaches such as robust-ensemble/replicated simulated annealing [32,34]. In particular, the methods of Refs. [32,34] introduce *explicit* couplings among replicas (e.g., via a reference/centroid constraint) that bias the search toward dense regions of solutions; by construction, this mechanism is absent in A-PSA. In contrast, SQA—through the path-integral mapping—evolves a set of Trotter slices coupled along imaginary time, enabling tunnelinglike transitions across tall and narrow barriers in canonical examples [20,81]. The comparison between these qualitatively distinct mechanisms, performed under the same QUBO encoding and evaluation metrics, is one of the key reproducible messages of the present benchmark. The superiority of A-PSA over SQA for $N = 59\,989$ suggests that the corresponding QUBO landscape does not strongly reward multislice tunnelinglike mechanisms. In this regime, diversity provided by many independent annealing trajec-




FIG. 11. Annealing schedules employed by each solver for $N = 8\,219\,999$. Panel (a) describes the case of SA solver. Panel (b) describes the case of SA-CE solver. Panel (c) describes the case of SQA-CE solver. Panel (d) describes the case of A-PSA solver. Panel (e) describes the case of A-PSA-CE solver.

tories can be sufficient to locate the optimum more frequently than SQA’s structured updates, whose imaginary-time correlations constrain the exploration process. This interpretation is consistent with the view that SQA’s advantage is most pronounced on landscapes dominated by tall, narrow barriers [20,81], which may not be the dominant obstacle for the instance $N = 59\,989$ under the direct QUBO encoding [43]. In contrast, the reversal for $N = 8\,219\,999$ indicates that the larger instance presents a more rugged or barrier-dominated landscape where SQA’s quantum-inspired mechanism becomes beneficial. Because A-PSA uses uncoupled trajectories, it cannot exploit any collective or tunnelinglike dynamics and effectively behaves as an ensemble of independent SA-like searches. Under these conditions, SQA can surpass A-PSA, reinforcing that the qualitative SQA-A-PSA relationship depends on the landscape structure induced by the chosen QUBO encoding and by the instance size. We also observe that SQA and A-PSA can display superficially similar *late-anneal* phenomenology (e.g., plateaus in residual energy or in progress toward the ground state), but the underlying causes differ. In SQA, the late-stage slowdown is associated with the effective strengthening of interslice coupling as the transverse field Γ decreases, which suppresses configurational changes and leads to a freeze-out of the imaginary-time dynamics [17,18]. In A-PSA, any analogous plateau originates from classical thermal effects: as the temperature decreases, Metropolis acceptance for uphill moves diminishes and each independent trajectory becomes increasingly trapped in the basin it has already entered [27]. The similarity of curves should therefore be interpreted cautiously: it does not imply identical dynamics, but rather that both methods can become dominated by basin-selection effects late in the schedule. A plausible explanation for the observed similarity between SQA and A-PSA in some regimes for QUBO-based factorization is the following. The direct mapping yields dense, highly constrained QUBOs where success at finite budget is often governed less by sustained “barrier traversal” and more by whether the solver *enters* a basin connected to the true factorization early enough to survive the subsequent cooling/freeze-out. In that situation, increasing the number of effectively independent “attempts” becomes a primary driver of success probability: A-PSA provides such attempts explicitly via independent restarts, whereas

SQA can provide an *implicit* multiplicity through its multislice representation and repeated sampling during the anneal, even if its slices are correlated. When this basin-selection mechanism dominates, both methods can yield comparable success probabilities once normalized by computational effort, despite their different microscopic update rules. This perspective is compatible with the resource-normalized conclusions reported in hardware-oriented studies that compare SQA against optimized classical baselines for factorization workloads [44]. Finally, SA consistently performs worst in both instances, in line with the well-known limitations of single-trajectory thermal annealing in escaping deep local minima. Within this replication study, SA therefore serves as a baseline confirming that, for the considered factorization QUBOs and computational budgets, additional mechanisms—either multistart diversity (A-PSA) or quantum-inspired multislice dynamics (SQA)—are necessary to obtain competitive success probabilities. Overall, the benchmark supports three reproducible conclusions: (1) An asynchronous multistart baseline (A-PSA) can outperform SQA when the induced landscape does not strongly reward tunnelinglike barrier crossing; (2) SQA can become superior for instances whose induced landscapes contain more significant barrier structure [20,81]; and (3) SQA and A-PSA can exhibit similar late-anneal phenomenology under fixed budgets, albeit for fundamentally different reasons rooted in quantum freeze-out versus classical thermal trapping [17,18,27]. By explicitly documenting solver settings and the evaluation pipeline, this work is intended to provide a stable reference point for future comparisons, including alternative QUBO mappings, improved schedules, and different computational budgets. The evaluation phase with varying computational budgets further indicates that scaling to larger semiprimes rapidly increases computational effort, emphasizing the practical importance of transparent time-to-solution reporting. Moreover, as discussed in Appendix C 6, more accurate and instance-tailored annealing schedules may improve performance in future studies. Additional algorithmic improvements may also arise from more efficient treatments of slack variables [82], which could reduce overhead in constraint handling and improve effective landscape navigation. From a broader perspective, these results reinforce that *reproducible, CPU-based simulation benchmarks* can meaningfully support the assessment of annealing-inspired optimization behavior in the near term, complementing hardware-specific evaluations. As quantum technologies mature, such reproducible benchmarks can serve as a common baseline for evaluating solver improvements, comparing implementations across platforms and clarifying when observed performance differences originate from algorithmic mechanisms versus tuning, budget, or instance effects.

ACKNOWLEDGMENTS

This research used classical computing resources provided by Intesa Sanpaolo Bank, specifically by “Cybersecurity Architecture & Innovation” structure within “Cybersecurity, Antifraud, Business Continuity Management” department, part of the “Chief Security Officer.” We would like to thank all the people involved within Cyberlab team, in particular Alberto Crippa and Scarongella Francesco, for their support and useful discussions on linkages between security topics and quantum threat arguments. We would like to thank Davide Corbelleto, lead of the Quantum Competence Center team in Intesa Sanpaolo, for its mentorship on quantum computing arguments.

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

The views and opinions expressed are those of the contributors and do not necessarily reflect the views of Intesa Sanpaolo, its affiliates, or its employees.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available because of legal restrictions preventing unrestricted public distribution. The data are available from the authors upon reasonable request.

APPENDIX A: SEARCH-SPACE REDUCTION IN QUBO-BASED INTEGER FACTORIZATION

In the QUBO formulation of integer factorization, one typically considers the objective

$$O(p, q) = (N - pq)^2, \quad (\text{A1})$$

where p and q are binary-encoded integers. The objective function obeys the permutation symmetry

$$O(p, q) = O(q, p), \quad (\text{A2})$$

which induces a twofold degeneracy in the energy landscape for all pairs (p, q) with $p \neq q$. Such degeneracy can adversely affect the performance of optimization procedures such as Simulated Classical Annealing (SA) or SQA. It is therefore useful to quantify the exact number of redundant configurations arising from this symmetry, and how much can be eliminated by breaking it.

In this Appendix, we analyze two symmetry-breaking mechanisms:

- (1) enforcing the ordering constraint $p \geq q$;
- (2) using asymmetric bit lengths for the two factors.

Both analyses include exact formulas valid for finite bit lengths (e.g., 2–20 bits), followed by asymptotic simplifications.

1. Case 1: enforcing the ordering constraint

Assume that p and q use the same number of bits,

$$n_p = n_q = n, \quad p, q \in \{0, \dots, 2^n - 1\}.$$

The full configuration space is the square

$$\Omega = \{0, \dots, 2^n - 1\}^2, \quad |\Omega| = 2^{2n}.$$

The symmetry $O(p, q) = O(q, p)$ identifies all pairs (p, q) and (q, p) with $p \neq q$. Thus, the redundant configurations correspond to the strict upper (or lower) triangular part of the $2^n \times 2^n$ square, whose exact cardinality is

$$R_{\text{sym}}^{(1)} = \frac{2^{2n} - 2^n}{2} = 2^{2n-1} - 2^{n-1}. \quad (\text{A3})$$

This counts every pair (p, q) with $p < q$ (or $p > q$), each of which has a symmetric partner. The surviving configurations are all pairs on or above the diagonal:

$$|\Omega_{p \geq q}| = 2^{2n} - R_{\text{sym}}^{(1)} = 2^{2n-1} + 2^{n-1}. \quad (\text{A4})$$

The exact reduction in search space is therefore

$$\mathcal{R}_{\text{ord}} = \frac{R_{\text{sym}}^{(1)}}{|\Omega|} = \frac{2^{2n-1} - 2^{n-1}}{2^{2n}} = \frac{1}{2} - 2^{-n-1}. \quad (\text{A5})$$

For large n , the asymptotic expression is

$$\mathcal{R}_{\text{ord}} \rightarrow \frac{1}{2}.$$

2. Case 2: asymmetric encoding ($n_p > n_q$)

Suppose

$$p \in \{0, \dots, 2^{n_p} - 1\}, \quad q \in \{0, \dots, 2^{n_q} - 1\}, \quad n_p > n_q.$$

Then,

$$|\Omega| = 2^{n_p + n_q}.$$

The permutation $(p, q) \leftrightarrow (q, p)$ is valid only when both coordinates lie inside the domain, i.e.,

$$p, q < 2^{n_q}.$$

Thus, the symmetry acts only on the square

$$S = \{0, \dots, 2^{n_q} - 1\}^2, \quad |S| = 2^{2n_q}.$$

Inside S , all off-diagonal pairs (p, q) with $p \neq q$ are redundant:

$$R_{\text{sym}}^{(2)} = \frac{2^{2n_q} - 2^{n_q}}{2} = 2^{2n_q-1} - 2^{n_q-1}. \quad (\text{A6})$$

The exact reduction relative to the whole rectangular domain is therefore

$$\mathcal{R}_{\text{asym}} = \frac{R_{\text{sym}}^{(2)}}{|\Omega|} = \frac{2^{2n_q-1} - 2^{n_q-1}}{2^{n_p + n_q}} \quad (\text{A7})$$

$$= \frac{2^{n_q} - 1}{2^{n_p+1}}. \quad (\text{A8})$$

For $n_q \gg 1$, the asymptotic form is

$$\mathcal{R}_{\text{asym}} \approx \frac{2^{n_q}}{2^{n_p+1}} = \frac{1}{2} 2^{n_q - n_p}.$$

If $n_p = n_q + k$, then

$$\mathcal{R}_{\text{asym}} \approx 2^{-(k+1)}.$$

3. Example reductions for small and midsize bit lengths

The following tables illustrate the exact reduction fractions for practical bit lengths.

n	Exact $\mathcal{R}_{\text{ord}} = \frac{1}{2} - 2^{-n-1}$	Numeric value
2	$\frac{1}{2} - \frac{1}{8} = 0.375$	0.375
4	$\frac{1}{2} - \frac{1}{32} = 0.46875$	0.46875
8	$\frac{1}{2} - \frac{1}{512} \approx 0.49805$	0.49805
12	$\frac{1}{2} - \frac{1}{8192} \approx 0.49988$	0.49988
20	$\frac{1}{2} - \frac{1}{2^{21}} \approx 0.4999995$	0.4999995

(n_p, n_q)	Exact $\mathcal{R}_{\text{asym}} = \frac{2^{n_q}-1}{2^{n_p+1}}$	Approx. $\frac{1}{2} 2^{n_q - n_p}$	Numeric value
(3, 2)	$\frac{4-1}{16} = 0.1875$	0.25	0.1875
(4, 3)	$\frac{8-1}{32} = 0.21875$	0.25	0.21875
(5, 3)	$\frac{8-1}{64} = 0.109375$	0.125	0.109375
(6, 4)	$\frac{16-1}{128} = 0.1171875$	0.125	0.1171875
(8, 4)	$\frac{16-1}{512} = 0.0292969$	0.03125	0.0292969

TABLE X. Statistics of the Ising model corresponding to different PUBO and PCBO encodings for $N = 59989$.

Encoding	n_{spins}	$J_{\min}$	$J_{\max}$	$ J $	$h_{\min}$	$h_{\max}$
PUBO (asymmetric)	71	-2.10×10^9	3.41×10^8	1.64×10^5	-3.00×10^9	3.15×10^9
PUBO (symmetric)	80	-8.46×10^9	3.35×10^9	2.87×10^5	-8.07×10^9	1.27×10^{10}
PCBO ($\lambda = 1$)	89	-8.46×10^9	3.35×10^9	2.62×10^5	-8.07×10^9	1.27×10^{10}
PCBO ($\lambda = 10$)	89	-8.46×10^9	3.35×10^9	2.62×10^5	-8.07×10^9	1.27×10^{10}
PCBO ($\lambda = 10^6$)	89	-3.28×10^{10}	3.28×10^{10}	1.05×10^6	-7.35×10^{10}	6.54×10^{10}
PCBO ($\lambda = 10^9$)	89	-3.28×10^{13}	3.28×10^{13}	1.05×10^6	-6.54×10^{13}	6.54×10^{13}
PCBO ($\lambda = 10^{10}$)	89	-3.28×10^{14}	3.28×10^{14}	1.05×10^6	-6.54×10^{14}	6.54×10^{14}

APPENDIX B: POLYNOMIAL CONSTRAINED BOOLEAN OPTIMIZATION FOR INTEGER FACTORIZATION

The encoding defined by Eq. (4), adopted for the numerical simulations, allows for different bit lengths for the factors p and q , but does not otherwise restrict the degeneracy of the search space. To assess its efficiency, we compare the solver performance with that obtained using a symmetric encoding, whereby

$$L_p = L_q = \left\lceil \frac{L_N}{2} \right\rceil + 1 + a, \quad (\text{B1})$$

in which the constraint $q \geq p$ is enforced through an additional penalty term in the QUBO formulation of the objective function [Eq. (6)]. This construction leads to a polynomial constrained Boolean optimization [83], where the magnitude of the constraint term is controlled by a the penalty parameter λ .

Finally, to assess the impact on the solver performance when transitioning from QUBO to PCBO, we also considered a QUBO formulation based on Eq. (B1).

The Ising-model statistics and the corresponding optimal probabilities of success obtained for $N = 59989$ are summarized in Tables X and XI.

TABLE XI. Maximum p_s values. Statistics of the Ising model corresponding to different PUBO and PCBO encodings. For each optimizer, the maximum value of p_s is highlighted in bold.

Encoding	Optimizer	SA	SA-CE	SQA-CE	A-PSA	A-PSA-CE
PUBO	TPE	0.02	0.04	0.88	0.92	0.97
(asymmetric)	GP	0.02	0.02	0.86	0.96	0.97
PUBO	TPE	0.02	0.01	0.64	0.79	0.85
(symmetric)	GP	0.01	0.01	0.66	0.82	0.82
PCBO	TPE	0.01	0.01	0.04	0.21	0.25
($\lambda = 1$)	GP	0.01	0.01	0.05	0.44	0.52
PCBO	TPE	0.01	0.02	0.08	0.17	0.23
($\lambda = 10$)	GP	0.01	0.01	0.08	0.41	0.45
PCBO	TPE	0.01	0.01	0.27	0.26	0.28
($\lambda = 10^6$)	GP	0.01	0.01	0.28	0.39	0.35
PCBO	TPE	0.00	0.01	0.06	0.05	0.03
($\lambda = 10^9$)	GP	0.01	0.00	0.04	0.04	0.04
PCBO	TPE	0.00	0.00	0.03	0.03	0.04
($\lambda = 10^{10}$)	GP	0.00	0.00	0.06	0.04	0.04

We note that, although PCBO-based encodings reduce the redundancy of the configuration space, they give rise to a substantially larger search space than PUBO-based encodings, as the corresponding Ising models involve a larger number of spins.

For PUBO-based encodings, the higher efficiency of the asymmetric formulation results in an increased probability of success. In contrast, for PCBO encodings, the probability of success is strongly dependent on the choice of the penalty parameter λ . Nevertheless, even the highest optimal probability of success achieved with PCBO-based encodings remains significantly lower than that obtained with PUBO-based encodings. This indicates that the reduction of configurational redundancy alone is insufficient to compensate for the increase in search-space size arising from the additional constraint encoded in the PCBO formulation.

It should be emphasized, however, that determining the optimal value of λ is, in general, a challenging task. In the present study, the range of λ values explored was not subject to extensive fine tuning. Consequently, a definitive assessment of the effectiveness of the PCBO-based encoding can only be made following a more systematic investigation of the penalty parameter optimization [76].

APPENDIX C: COMPARATOR PIPELINE: CONSTRUCTIONS AND METRICS

In this Appendix, we present a comprehensive quantitative analysis of the Ising model derived from the QUBO formulation of the integer factorization problem. The model is specified by the set of couplings J_{ij} , local fields h_i , and a constant offset. Since the structure of the interaction graph plays a central role in determining dynamical behavior, algorithmic hardness, and similarity to known physical systems, we compute structural metrics motivated by statistical physics, spectral graph theory [84], and spin-glass theory.

We furthermore compare the empirical structure of the QUBO instance to canonical benchmark models in the physics literature, namely, (1) the Sherrington-Kirkpatrick (SK) spin glass [85–87], (2) the Hopfield associative-memory model [88], (3) a sparse Edwards-Anderson-like (EA-like) model based on the DATA sparsity pattern [89,90], and (4) a canonical two-dimensional (2D) EA lattice model. To these four standard comparators, we add three *advanced comparators* motivated by geometrically frustrated magnets and artificial spin-ice (ASI) systems: (5) a nearest-neighbor kagome antiferromagnet (KAGOME_NN), (6) a nearest-neighbor triangular-lattice antiferromagnet (TRIANGULAR_NN), and (7) a kagome dipolar artificial spin-ice analog (ASI_KAGOME_DIP) [91,92].

All comparators [93–95] are rescaled to match the off-diagonal coupling variance of the DATA instance so that comparisons reflect structural differences rather than trivial global energy units.

1. Symmetric reconstruction and Hamiltonian conventions

The QUBO-derived Ising instance is given as an upper-triangular dictionary $J^{\text{DATA}} = \{(i, j) \mapsto J_{ij} \ (i < j)\}$, local fields h^{DATA} , and a constant offset. We reconstruct the symmetric coupling matrix $J^{\text{DATA}}_{ji} = J^{\text{DATA}}_{ij}$ and set $J^{\text{DATA}}_{ii} = 0$, inferring the system size N from the largest index occurring in J or h . The Hamiltonian is written over unordered pairs as

$$\mathcal{H}(s) = - \sum_{i < j} J_{ij} s_i s_j - \sum_i h_i s_i - \text{offset} = - \frac{1}{2} \sum_{i \neq j} J_{ij} s_i s_j - \sum_i h_i s_i - \text{offset}, \quad (\text{C1})$$

with the factor $\frac{1}{2}$ avoiding double counting. This symmetrized matrix representation guarantees real spectra and orthogonal eigenvectors, consistent with standard treatments of mean-field and short-range spin glasses [85,88,89].

Let

$$\hat{\sigma}_{\text{DATA}} = \text{std}(\{J^{\text{DATA}}_{ij} : i < j, J_{ij} \neq 0\})$$

denote the empirical off-diagonal standard deviation. This quantity serves as the *reference energy scale* for calibrating all comparators.

2. Classical comparators matched to DATA scale

(1) *Sherrington-Kirkpatrick, fully connected mean field*: The SK model [85] is the canonical infinite-range spin glass:

$$J_{ij}^{(\text{SK})} \sim \mathcal{N}(0, \hat{\sigma}_{\text{DATA}}^2)(i < j), \quad J_{ji}^{(\text{SK})} = J_{ij}^{(\text{SK})}, \quad J_{ii}^{(\text{SK})} = 0. \quad (\text{C2})$$

This produces fully connected, mean-field disorder with a Wigner-like spectral density in the large- N limit [86].

(2) *Hopfield model (Hebbian low rank)*: Given p random patterns $\xi^\mu \in \{\pm 1\}^N$, the Hebbian matrix is

$$\tilde{J}_{ij}^{(\text{HOP})} = \frac{1}{N} \sum_{\mu=1}^p \xi_i^\mu \xi_j^\mu, \quad \tilde{J}_{ii}^{(\text{HOP})} = 0. \quad (\text{C3})$$

We rescale

$$J^{(\text{HOP})} = \tilde{J}^{(\text{HOP})} \frac{\hat{\sigma}_{\text{DATA}}}{\tilde{\sigma}}, \quad \tilde{\sigma} = \text{std}(\{\tilde{J}_{ij}^{(\text{HOP})} : i < j\}), \quad (\text{C4})$$

so that the typical coupling magnitude matches DATA. This model exhibits low-rank structure and p macroscopic outliers [88]. We draw $p = (N/10)$ random patterns.

(3) *EA-like (instance-matched short-range surrogate)*: We preserve the DATA sparsity pattern and magnitudes while randomizing signs:

$$J_{ij}^{(\text{EA})} = \sigma_{ij} |J_{ij}^{\text{DATA}}| (i < j), \quad \sigma_{ij} \in \{\pm 1\} \text{ i.i.d.} \quad (\text{C5})$$

This removes logical correlations while isolating geometry-driven frustration in the spirit of short-range EA glasses [90,96].

(4) *EA_LATTICE (canonical 2D nearest-neighbor EA)*: We embed N spins on a square $L_x \times L_y$ lattice, assign nearest-neighbor bonds

$$J_{ij}^{(\text{EA_LAT})} \in \{\pm 1\} \quad \text{or} \quad \mathcal{N}(0, 1), \quad (\text{C6})$$

and rescale nonzero bonds so that

$$\text{std}(\{J_{ij}^{(\text{EA_LAT})} \neq 0\}) = \hat{\sigma}_{\text{DATA}}.$$

This produces a canonical short-range spin glass with regular geometry [89].

3. Advanced comparators: frustrated lattices and dipolar spin ice

(5) *KAGOME_NN (nearest-neighbor kagome antiferromagnet)*: We construct a kagome lattice of N sites and include only nearest-neighbor bonds. The Ising Hamiltonian is

$$H_{\text{kagome}} = - \sum_{\langle i, j \rangle_{\text{kag}}} J_{ij} s_i s_j, \quad J_{ij}^{(0)} = -1 \text{ on kagome NN edges}, \quad (\text{C7})$$

and $J_{ij} = 0$ otherwise. Thus, each bond is antiferromagnetic and the model realizes the classical Ising kagome frustrated magnet composed of corner-sharing triangles. After construction, the nonzero couplings are rescaled so that $\text{std}(J_{ij} \neq 0) = \hat{\sigma}_{\text{DATA}}$.

(6) *TRIANGULAR_NN (nearest-neighbor triangular antiferromagnet)*: On a triangular lattice, the Ising Hamiltonian is

$$H_{\text{tri}} = - \sum_{\langle i, j \rangle_{\Delta}} J_{ij} s_i s_j, \quad J_{ij}^{(0)} = -1 \text{ on NN edges}, \quad (\text{C8})$$

and zero otherwise. This generates the canonical geometrically frustrated triangular Ising AFM, in which each triangle contains one unsatisfied bond. As above, all nonzero couplings are finally rescaled so that $\text{std}(J_{ij} \neq 0) = \hat{\sigma}_{\text{DATA}}$.

(7) *ASI_KAGOME_DIP* (*kagome artificial spin-ice dipolar model*): ASI systems exhibit long-range dipolar interactions, ice-rule manifolds, emergent monopoles, and rich frustration physics [91,92]. We construct a kagome lattice, assign each site an easy axis (three directions at 120°), and compute a dipolar kernel:

$$J_{ij}^{\text{dip}} \propto \frac{\mathbf{e}_i \cdot \mathbf{e}_j - 3(\mathbf{e}_i \cdot \hat{\mathbf{r}}_{ij})(\mathbf{e}_j \cdot \hat{\mathbf{r}}_{ij})}{r_{ij}^3}. \quad (\text{C9})$$

We restrict to interactions within a cutoff radius or to the k nearest neighbors. The final matrix is symmetrized and rescaled so that its nonzero entries have standard deviation $\hat{\sigma}_{\text{DATA}}$.

4. Calibration summary: the key uniformity constraint

All comparators $J^{(X)}$ (classical and advanced) are rescaled so that the standard deviation of their off-diagonal nonzero couplings matches that of the DATA instance:

$$\text{std}(\{J_{ij}^{(X)} : i < j, J_{ij}^{(X)} \neq 0\}) = \hat{\sigma}_{\text{DATA}}. \quad (\text{C10})$$

This ensures that similarity measurements reflect *structural* differences—geometry, frustration patterns, spectral structure, and motif densities—rather than trivial global rescaling of energies.

The only exception is the EA-like comparator, which already inherits bond magnitudes from DATA and therefore naturally shares its scale.

5. Why this design matters

By embedding DATA into a unified family of models of increasing structural specificity—from fully connected random glasses (SK) to low-rank models (Hopfield), to irregular short-range glasses (EA-like), and finally to geometrically frustrated lattices and dipolar ASI systems—we obtain a robust “structural phase portrait” of the QUBO instance.

This broad comparator ensemble allows one to diagnose

- (1) whether DATA behaves like a mean-field glass (SK-like),
- (2) whether it contains low-rank coherent structure (Hopfield-like),
- (3) whether its complexity is graph-geometric (EA-like),
- (4) whether it resembles ideal short-range glasses (EA_LATTICE),
- (5) whether frustration arises primarily from triangular motifs (TRIANGULAR_NN),
- (6) whether DATA exhibits kagome-like degeneracies (KAGOME_NN),
- (7) whether DATA mimics long-range dipolar frustration (ASI_KAGOME_DIP).

Because all comparators share both the same number of spins and the same effective coupling scale, this comparison is structurally meaningful and enables principled similarity analysis as described in Appendix C 6.

a. Structural metrics (DATA and comparators)

To characterize the structure of the DATA Ising instance and to enable a meaningful comparison with the classical and advanced comparators, we compute a set of metrics rooted in statistical physics, spectral graph theory, and the mathematical theory of spin glasses. These metrics probe disorder distributions, graph geometry, frustration, spectral structure, and dynamical features of the corresponding energy landscape.

Let the set of unique off-diagonal couplings be

$$\mathcal{J} = \{J_{ij} : i < j\}, \quad (\text{C11})$$

and define the unweighted interaction graph by the adjacency matrix

$$A_{ij} = \mathbb{I}(|J_{ij}| > J_{\text{tol}}), \quad J_{\text{tol}} \geq 0. \quad (\text{C12})$$

Distributional statistics. We compute summary statistics of $\mathcal{J}$: mean, standard deviation, skewness, excess kurtosis, minimum and maximum, sign fractions $\mathbb{P}(J > 0)$ and $\mathbb{P}(J < 0)$, and robust scales such as $\mathbb{E}(|J|)$ and $\text{median}(|J|)$. These quantities are routinely used in the study of mean-field spin glasses [85], short-range EA glasses [89,90], and low-rank associative-memory models [88].

Graph measures. The adjacency A induces a structural graph $G = (V, E)$ with degrees $k_i = \sum_j A_{ij}$. We compute

- (1) connected components;
- (2) local clustering coefficients $C_i = 2T_i/(k_i(k_i - 1))$, where T_i is the number of triangles incident on i ;
- (3) global transitivity

$$C_{\text{glob}} = \frac{3 \text{ No. of \{triangles\} }}{\sum_i k_i(k_i - 1)}, \quad (\text{C13})$$

which measures the overall density of triangular motifs.

Clustering and motif densities are central diagnostics for geometric frustration [91] and irregular random graphs.

Triangle frustration. For each triangle (i, j, k) in the interaction graph, define the loop product $\Pi_{ijk} = J_{ij}J_{jk}J_{ki}$; the triangle is *frustrated* if $\text{sign}(\Pi_{ijk}) = -1$. We report both an unweighted and an interaction-weighted fraction:

$$F = \frac{\text{No. of \{sign}(\Pi_{ijk}) = -1\}}{\text{No. of \{triangles\}}}, \quad F_{\text{wt}} = \frac{\sum_{\text{frustrated}} |\Pi_{ijk}|}{\sum_{\text{all}} |\Pi_{ijk}|}. \quad (\text{C14})$$

Frustration is a defining signature of competing interactions and glassy landscapes [91,96].

Spectral diagnostics. Since J is symmetric, the spectrum consists of real eigenvalues $\lambda_1 \geq \dots \geq \lambda_N$ with orthogonal eigenvectors. We compute

- (1) the empirical spectral density (histogram of eigenvalues),
- (2) the scree curve (sorted eigenvalues),
- (3) participation ratios (PRs) of the leading eigenvectors $\text{PR}(v) = 1/\sum_i v_i^4$, probing localization.

The participation ratio (PR) is a standard diagnostic of eigenvector localization used across Anderson localization theory [97,98], random-matrix theory and quantum chaos [99,100], spectral graph theory [101], and spin-glass models [102]. Extended eigenvectors have large PR, while localized modes have small PR. Mean-field SK exhibits a semicircular spectral bulk [85,86], Hopfield models exhibit p macroscopic outliers [88], and sparse EA-like models often show partial localization.

Optional landscape probe: overlap distribution $P(q)$. Using a Metropolis sampler at inverse temperature β , independent replicas $s^{(1)}$ and $s^{(2)}$ can be produced. The spin overlap

$$q = \frac{1}{N} \sum_{i=1}^N s_i^{(1)} s_i^{(2)}$$

estimates the distribution $P(q)$: sharp peaks at $q \neq 0$ indicate retrieval/memory structure (as in Hopfield systems [88]), while broad or multimodal $P(q)$ indicates glassy landscapes (as in SK/EA glasses [85,90]).

Its distribution $P(q)$ distinguishes phases: A single sharp peak indicates a paramagnet, multiple peaks correspond to retrieval states in associative-memory models, and a broad or continuous distribution reflects replica-symmetry breaking in spin glasses. The concept originates from Parisi's solution of the SK model [103] and is reviewed extensively in Refs. [104,105]. In Hopfield networks, peaks in $P(q)$ signal memory attractors [106,107]. Modern numerical studies use Monte Carlo sampling to estimate $P(q)$ in EA and SK glasses.

TABLE XII. Similarity components and mappings $S_\alpha \in [0, 1]$. JS distance is mapped via $1 - \text{JS}/\ln 2$, KS via $1 - \text{KS}$.

Component	Definition	Structural intent
J_dist	JS/KS on off-diagonal J histograms	Local coupling statistics (bimodal versus Gaussian).
eig_dist	JS/KS on eigenvalue density	Bulk spectra: Wigner-like (SK) versus sparse (EA).
scree	Cosine similarity of top- m eigenvalues	Outliers (Hopfield) versus bulk-dominated profiles.
pr	Cosine similarity of top- k PR vectors	Mode localization across graph.
F, Fw	$1 - F_{\text{DATA}} - F_X $, $1 - F_{\text{wt}} $	Spin-glass frustration in triangular loops.
Cglob	$1 - C_{\text{glob,DATA}} - C_{\text{glob,X}} $	Triangle motif density/geometry.
deg_dist	JS/KS on degree histogram	Graph sparsity structure.
pq_dist	JS/KS on $P(q)$	Landscape signature (retrieval versus glassy).

6. Similarity analysis methodology

The goal of the similarity analysis is to quantify how close a comparator X (e.g., SK, HOP, EA, EA_LATTICE, or any advanced comparator) is to the DATA instance in an interpretable, multicomponent manner as is summarized in Table XII. Each component is mapped to a unit interval $[0, 1]$ (higher = more similar), and a final composite score $S_{\text{comp}}(X) \in [0, 1]$ summarizes the overall resemblance.

If a component is unavailable for a comparator [e.g., $P(q)$ not computed], its weight is redistributed across the remaining components.

a. Distributional similarities: Jensen-Shannon and KS

Given normalized histograms P and Q over common bins, the Jensen-Shannon divergence [108]

$$\text{JS}(P\|Q) = \frac{1}{2}\text{KL}(P\|M) + \frac{1}{2}\text{KL}(Q\|M), \quad M = \frac{1}{2}(P + Q), \quad (\text{C15})$$

is finite, symmetric, and takes values in $[0, \ln 2]$. We convert it to similarity via

$$S_{\text{JS}} = 1 - \text{JS}/\ln 2.$$

We apply this to the distributions of off-diagonal couplings (J_dist) and eigenvalues (eig_dist). Optionally, the Kolmogorov-Smirnov test statistic [109] is used instead, with $S = 1 - \text{KS}$.

b. Spectral-shape similarities

Let the top- m sorted eigenvalues of DATA and model X be $\lambda_{(1)}^{(\text{DATA})} \geq \dots \geq \lambda_{(m)}^{(\text{DATA})}$ and $\lambda_{(1)}^{(X)}, \dots, \lambda_{(m)}^{(X)}$.

The scree similarity compares the leading m sorted eigenvalues $\lambda_{(1)} \geq \dots \geq \lambda_{(m)}$ of the DATA matrix to those of the comparator model. Since the upper tail of the spectrum encodes global structure—mean-field randomness (SK), low-rank “pattern” outliers (Hopfield), or sparse short-range geometry (EA-like)—the shape of the scree curve provides a robust structural fingerprint. This idea originates from classical scree plots in multivariate analysis [110] and is widely used in random-matrix theory [86, 111] and the spectral theory of spin glasses [107, 112]. We measure the cosine similarity of the leading eigenvalue vectors and map it to $[0, 1]$.

Define the scree similarity

$$S_{\text{scree}}(X) = \frac{1}{2} \left(1 + \frac{\sum_{k=1}^m \lambda_{(k)}^{(\text{DATA})} \lambda_{(k)}^{(X)}}{\|\lambda_{(1:m)}^{(\text{DATA})}\| \|\lambda_{(1:m)}^{(X)}\|} \right) \in [0, 1]. \quad (\text{C16})$$

For the top- k participation ratios $\text{PR}(v)$ of DATA and model X , we similarly compute a cosine similarity $S_{\text{pr}}(X) \in [0, 1]$.

c. Frustration and geometric similarities

We define

$$S_F(X) = 1 - |F_{\text{DATA}} - F_X|, \quad S_{F_w}(X) = 1 - |F_{\text{wt,DATA}} - F_{\text{wt,X}}|,$$

$$S_C(X) = 1 - |C_{\text{glob,DATA}} - C_{\text{glob,X}}|,$$

each in $[0, 1]$. Degree distributions are compared via JS (or KS) to obtain $S_{\text{deg}}(X)$. These metrics capture graph-level sparsity, motif structure, and loop frustration, which play central roles in physical spin glasses [91,96].

d. Dynamical similarity via $P(q)$ (optional)

If $P(q)$ is available for DATA and model X , we compute

$$S_q(X) = 1 - \frac{\text{JS}(P_{\text{DATA}}(q) \| P_X(q))}{\ln 2}$$

(or the KS analog). This component detects whether the comparator replicates the dynamical/glassy landscape structure: sharp peaks (retrieval), broad multimodal profiles (spin glass), or near-zero peaks (paramagnet) [85,88,90].

e. Composite score

Let $\mathcal{A}_X$ be the components available for model X , with non-negative weights $\{w_\alpha\}$ summing to 1. We define the composite similarity as

$$S_{\text{comp}}(X) = \frac{\sum_{\alpha \in \mathcal{A}_X} w_\alpha S_\alpha(X)}{\sum_{\alpha \in \mathcal{A}_X} w_\alpha}, \quad (\text{C17})$$

with the denominator ensuring consistent renormalization when some components [e.g., $P(q)$] are missing.

Default weight choice. Our default weights reflect a balanced combination of structural signatures:

- (1) distributional statistics: $w_{\text{J_dist}} = 0.15$, $w_{\text{eig_dist}} = 0.15$;
- (2) spectral-shape diagnostics: $w_{\text{scree}} = 0.10$, $w_{\text{pr}} = 0.10$;
- (3) frustration and geometry: $w_F = 0.15$, $w_{F_w} = 0.05$, $w_C = 0.05$, $w_{\text{deg}} = 0.10$;
- (4) dynamical information (optional): $w_{\text{pq_dist}} = 0.15$.

This weighting scheme reflects established practice in benchmarking classical and quantum annealing performance using spectral densities, frustration measures, and graph motifs [113]. It ensures that no single diagnostic dominates unless its similarity is genuinely anomalous.

7. Pipeline inputs, metrics, and similarities for the $N = 59\,989$ instance

We executed the complete structural analysis pipeline on the number 59 989 of 16 bits length, represented by $d = 71$ spins in Fig. 12, using the following call to Python routines:

```
manifest = run_full_ising_analysis(
    ising_dict,
    outdir='ising_qubo_structure_3',
    hopfield_p=None,          # defaults to N//10
    compute_overlap=True,     # compute P(q)
    beta_for_Pq=0.15,
    sweeps_for_Pq=300,
    topk_PR=10,
    J_tol_edge=0.0,
    random_seed=12345,
```




FIG. 12. Visualization of the Ising model for the 16-bit number $N = 59989$, consisting of the upper-triangular interaction matrix J , the local fields h , and the global energy offset, all plotted with a shared color scale centered at zero. Asymmetric encoding is used.

```

build_advanced=True,          # build advanced comparators
advanced_specs=None,          # defaults for 3 advanced comparators
ea_lattice_periodic=False,
ea_lattice_dist='pmJ'
)
    
```

The factorization problem is represented using the asymmetric encoding of Eq. (4).

Meaning of the main parameters.

- (1) `hopfield_p`: number p of Hopfield patterns; by default $p = \lfloor d/10 \rfloor$.
- (2) `compute_overlap`: If True, compute the overlap distribution $P(q)$ via a Metropolis sampler at inverse temperature $\beta = \text{beta_for_Pq}$ with `sweeps_for_Pq` Monte Carlo sweeps per replica pair.
- (3) `topk_PR`: number k of leading eigenvectors used to form the PR vector for spectral-shape comparisons.
- (4) `J_tol_edge`: edge threshold on $|J_{ij}|$ for building the unweighted interaction graph (triangles, clustering).
- (5) `build_advanced`: If True, generate the three advanced comparators (KAGOME_NN, TRIANGULAR_NN, ASI_KAGOME_DIP) using `advanced_specs` or defaults sized to N ; ASI uses a dipolar kernel on a kagome lattice.

(6) `ea_lattice_periodic`, `ea_lattice_dist`: boundary condition and bond distribution ($\pm J/\text{Gaussian}$) for the canonical EA lattice comparator; nonzero bonds are rescaled to match DATA's off-diagonal standard deviation $\hat{\sigma}_{\text{DATA}}$.

All comparators are then analyzed with the suite of structural (distributional, graph, frustration) and spectral metrics, and the results exported to CSV files for *post hoc* similarity analysis (Appendix C 6).

a. Summary statistics discussion

The DATA instance for $N = 59\,989$ produces an Ising model with $d = 71$ spins. The metrics summary regarding DATA and comparators is shown in Table XIII. Its coupling distribution has heavy positive bias (90.4% positive), strong heavy tails (excess kurtosis ≈ 136), and moderate frustration ($F \approx 0.244$). The global clustering $C_{\text{glob}} \approx 0.269$ indicates significant triangle density.

The column `offdiag_sigma_ref_DATA` reports the DATA reference scale $\hat{\sigma}_{\text{DATA}} = \text{std}(J_{ij} : i < j, J_{ij} \neq 0)$. All comparators attempt to match this via multiplicative rescaling. For models with genuine coupling variance (SK, HOP, EA, EA_LATTICE, ASI_KAGOME_DIP), the empirical J_{std} aligns with $\hat{\sigma}_{\text{DATA}}$ up to $\mathcal{O}(1/\sqrt{M})$ Monte Carlo fluctuations, where M is the number of nonzero edges. This explains the small ($\lesssim 1\%$) differences for SK and EA, both consistent with sampling error.

For KAGOME_NN and TRIANGULAR_NN, all nonzero bonds are identically -1 . Therefore, $\text{std}(J) = 0$ and no meaningful rescaling can be applied. The pipeline properly skips rescaling, producing $J_{\text{std}} = 0$.

SK and HOP are dense; EA matches DATA sparsity exactly (0.4704). EA_LATTICE, KAGOME_NN, and TRIANGULAR_NN show very low densities, consistent with their NN geometries; ASI_KAGOME_DIP shows intermediate density due to dipolar neighborhoods.

The global clustering C_{glob} reveals triangle content: EA and DATA match exactly (0.2691); triangular/kagome comparators exhibit expected geometric triangle density (0.17–0.21); ASI dipolar has higher clustering (0.315). EA_LATTICE is triangle-free (square NN lattice); hence $C_{\text{glob}} = 0$ and frustration metrics are undefined.

DATA shows moderate frustration ($F = 0.244$). SK and EA show $F \approx 0.5$, matching random-sign expectations. Triangular and kagome AFM comparators show $F = 1$, as every NN triangle is frustrated. EA_LATTICE yields blanks because no triangles are present. The ASI dipolar model shows intermediate frustration ($F \approx 0.507$), consistent with its mixed-sign, anisotropic couplings.

The metrics collectively indicate that the $d = 71$ DATA instance is best described as a *sparse, irregular, triangle-rich spin glass with heavy-tailed couplings and moderate frustration*, closely captured by the EA-like comparator, with ASI_KAGOME_DIP and frustrated-lattice comparators reproducing only portions of its structural signature.

b. Similarity analysis and implications for simulated Quantum Annealing

We compute componentwise similarities (each mapped to $[0, 1]$) and a composite similarity score using the weighted scheme of Appendix C 6. Table XIV reports the results for the $N = 59\,989$ instance (DATA size $d = 71$ spins). Higher values indicate closer structural proximity to the DATA Ising model.

Key observations.

(1) EA (instance matched) remains the closest comparator ($S_{\text{comp}} \approx 0.752$). Its perfect degree distribution match and identical clustering ($C_{\text{glob}} = 1$ after normalization) confirm that the DATA graph is sparsity driven and structurally similar to an EA-like short-range spin glass [89,90]. These high proximities reflect that DATA's topology dominates its complexity.

(2) ASI_KAGOME_DIP is consistently the second-closest structure (≈ 0.639). Its dipolar weights produce mixed signs, anisotropic magnitudes, and moderate frustration ($F \approx 0.507$), all qualitatively similar to DATA. This suggests the DATA instance possesses interactions whose *structure* is closer to weakly long-range frustrated geometries than to ideal NN lattices [92].

TABLE XIII. Metrics summary for the $N = 59\,989$ factorization instance (DATA size: $d = 71$ spins).

Model	No zero off diag	Density	J_{mean}	J_{std}	Skew	Kurt	frac_+ + frac_-	F	F_{wt}	C_{glob}	No. of Comp.	Comp. size	$\hat{\sigma}_{\text{DATA}}$
DATA	2338	0.4704	-1.4047×10^7	1.2888×10^8	-10.26	136.21	0.904	0.096	0.2444	0.1668	1	71	1.2888×10^8
SK	4970	1.0000	8.435×10^5	1.2961×10^8	0.0026	-0.071	0.497	0.503	0.4992	0.4996	1	71	1.2888×10^8
HOP	4970	1.0000	-1.771×10^5	1.2888×10^8	0.0630	-0.351	0.497	0.503	0.4076	0.1885	1	71	1.2888×10^8
EA	2338	0.4704	-8.079×10^6	1.2939×10^8	-8.17	136.44	0.478	0.522	0.4913	0.7151	1	71	1.2888×10^8
EA_LATTICE	250	0.0503	-7.214×10^6	1.2888×10^8	0.113	-2.003	0.472	0.528	-	0.0000	1	71	1.2888×10^8
KAGOME_NN	246	0.0495	-1.0	0.0	0.0	0.0	0.0	1.0	1.0	0.1759	1	71	1.2888×10^8
TRIANGULAR_NN	362	0.0728	-1.0	0.0	0.0	0.0	0.0	1.0	1.0	0.2116	1	71	1.2888×10^8
ASL_KAGOME_DIP	1552	0.3123	4.701×10^7	1.2888×10^8	1.674	1.179	0.577	0.423	0.5071	0.4005	1	71	1.2888×10^8

TABLE XIV. Similarity to DATA for the $N = 59\,989$ instance. The composite score is a weighted average of distributional (JS/KS), spectral-shape, geometric/frustration, and optional $P(q)$ components, renormalized to the available metrics.

Model	J_dist	eig_dist	scree	pr	F	F_w	C_{glob}	deg_dist	Composite
EA	0.9971	0.8282	0.9638	0.9617	0.7530	0.4518	1.0000	1.0000	0.752
ASI_KAGOME_DIP	0.8201	0.5341	0.8307	0.9786	0.7373	0.7663	0.9543	0.5795	0.639
SK	0.5749	0.5503	0.8606	0.9836	0.7452	0.6672	0.7691	0.0931	0.546
TRIANGULAR_NN	0.9494	0.7870	0.8196	0.9838	0.2444	0.1668	0.9424	≈ 0	0.533
KAGOME_NN	0.9494	0.7870	0.7790	0.9677	0.2444	0.1668	0.9068	≈ 0	0.525
HOP	0.1749	0.1245	0.9669	0.9302	0.8367	0.9783	0.7691	0.0931	0.475
EA_LATTICE	0.0549	0.6452	0.8362	0.9640	—	—	0.7309	≈ 0	0.402

(3) SK reaches $S_{\text{comp}} \approx 0.546$. SK matches some spectral-shape features (scree, PR), consistent with its mean-field nature [85], but differs sharply in degree distribution and local geometry. As in many QUBO-derived sparse Ising models, full connectivity is a poor approximation.

(4) KAGOME_NN and TRIANGULAR_NN cluster near 0.53. Both reproduce good coupling and eigenvalue distribution matches, but their frustration fractions ($F=1$) strongly disagree with DATA’s moderate value ($F \approx 0.244$). These comparators capture geometric frustration but lack the heterogeneity and asymmetric coupling statistics of DATA (cf. frustrated magnetism [91]).

(5) Hopfield is significantly mismatched (≈ 0.475). Despite excellent scree/PR similarity (due to globally extended top modes), its poor J and eigenvalue-distribution similarities indicate that DATA does not exhibit low rank or pattern-retrieval structure [88].

(6) EA_LATTICE is the most dissimilar (≈ 0.402). Its triangle-free, quasiregular 2D geometry fails to capture DATA’s motif density and frustration profile. This confirms that DATA is *not* well modeled by a simple short-range regular lattice [89].

Implications for SQA schedules. The similarity structure—EA $\gg$ ASI $\gg$ (SK, frustrated lattices) $\gg$ Hopfield, EA_LATTICE—points to a *sparse, moderately frustrated, heterogeneous Ising landscape*. Such landscapes are well known to benefit from schedule controls [114–116] developed for glassy models [113,117,118].

(1) Midanneal plateaus (pauses). EA- and ASI-like systems have small-gap regions where diabatic excitations are common. Pausing shortly after the avoided crossing significantly improves ground-state recovery [117,118]. In SQA, this corresponds to inserting a plateau in $\Gamma(t)$ or $\beta(t)$ before the final descent.

(2) Inhomogeneous transverse-field driving. Because DATA’s geometry is irregular (perfectly matched by EA), per-qubit anneal offsets (softening the field more slowly on high-degree or highly frustrated nodes) help mitigate local first-order transitions [119,120]. This is particularly relevant given the similarity between DATA and ASI-like couplings, which exhibit local anisotropies.

(3) Reverse-anneal refinement. For EA-like landscapes, reverse annealing (increasing Γ around a seed, brief hold, slow return) efficiently refines good classical states [121,122]. The similarity scores suggest DATA belongs to exactly this class.

(4) Multiple short anneals with restarts. Sparse frustrated systems often exhibit wide basin structures; repeated short runs outperform single slow runs at fixed compute budget [113]. Given DATA’s EA-like profile, restart strategies are strongly recommended.

(5) Freeze-out mitigation. Moderate frustration and heavy tails (seen in metrics summary) imply a freeze-out region. Pauses and restarts are significantly more effective than simply increasing inverse temperature β [117].

(6) (Advanced) Controlled diabatic features. Small diabatic “bumps”—temporary increases in Γ before a pause—can exploit nonadiabatic pathways in glassy landscapes [81,123]. The similarity to ASI-like models supports experimenting with such features numerically.

TABLE XV. Optimal parameter values obtained for each semiprime integer, optimizer, and solver. For each optimizer, the maximum value of p_s is highlighted in bold.

N	n	n_{reads}	Optimizer	$\mathcal{S}$	MCS	p_s	$\beta_{\min}$	$\beta_{\max}$	R	M	β	Γ
15	4	100	GP	SA	4000	1.00	1.3×10^{-3}	2.7×10^{-1}	—	—	—	—
				SA-CE	4000	1.00	1.3×10^{-3}	5.5×10^{-1}	—	—	—	—
				SQA-CE	3268	1.00	—	—	—	817	6.3	9.6×10^1
				A-PSA	3268	1.00	1.3×10^{-3}	1.8	817	—	—	—
				A-PSA-CE	3268	1.00	1.3×10^{-3}	4.8×10^{-1}	817	—	—	—
			TPE	SA	4000	1.00	1.3×10^{-3}	1.1	—	—	—	—
				SA-CE	4000	1.00	1.3×10^{-3}	2	—	—	—	—
				SQA-CE	3798	1.00	—	—	—	422	3.1	1.6×10^2
				A-PSA	3798	1.00	1.3×10^{-3}	8×10^{-1}	422	—	—	—
				A-PSA-CE	3798	1.00	1.3×10^{-3}	6.3×10^{-1}	422	—	—	—
21	5	100	GP	SA	4000	0.81	2.2×10^{-4}	4.2	—	—	—	—
				SA-CE	4000	1.00	2.2×10^{-4}	2.6	—	—	—	—
				SQA-CE	3934	1.00	—	—	—	562	3.8	3.3×10^2
				A-PSA	3934	1.00	2.2×10^{-4}	3.5×10^{-1}	562	—	—	—
				A-PSA-CE	3934	1.00	2.2×10^{-4}	2.3×10^{-1}	562	—	—	—
			TPE	SA	4000	0.74	2.2×10^{-4}	3.6	—	—	—	—
				SA-CE	4000	1.00	2.2×10^{-4}	3.9	—	—	—	—
				SQA-CE	3680	1.00	—	—	—	920	1.2×10^1	2.8×10^2
				A-PSA	3680	1.00	2.2×10^{-4}	1.9×10^{-1}	920	—	—	—
				A-PSA-CE	3680	1.00	2.2×10^{-4}	1.2×10^{-1}	920	—	—	—
35	6	100	GP	SA	4000	0.32	5.9×10^{-5}	6.7×10^{-2}	—	—	—	—
				SA-CE	4000	1.00	5.9×10^{-5}	1.5	—	—	—	—
				SQA-CE	3976	1.00	—	—	—	56	6.1×10^{-2}	1.1×10^3
				A-PSA	3976	1.00	5.9×10^{-5}	3.5×10^{-1}	56	—	—	—
				A-PSA-CE	3976	1.00	5.9×10^{-5}	2.4	56	—	—	—
			TPE	SA	4000	0.32	5.9×10^{-5}	5.6	—	—	—	—
				SA-CE	4000	1.00	5.9×10^{-5}	6.1	—	—	—	—
				SQA-CE	3952	1.00	—	—	—	208	3.6×10^{-1}	6.3×10^2
				A-PSA	3952	1.00	5.9×10^{-5}	3.3	208	—	—	—
				A-PSA-CE	3952	1.00	5.9×10^{-5}	6.5	208	—	—	—
57	6	100	GP	SA	4000	0.46	6.6×10^{-5}	4.9	—	—	—	—
				SA-CE	4000	1.00	6.6×10^{-5}	6.5	—	—	—	—
				SQA-CE	3468	1.00	—	—	—	867	1.6	3.2×10^2
				A-PSA	3468	1.00	6.6×10^{-5}	5.5	867	—	—	—
				A-PSA-CE	3468	1.00	6.6×10^{-5}	6.4	867	—	—	—
			TPE	SA	4000	0.52	6.6×10^{-5}	5.1	—	—	—	—
				SA-CE	4000	1.00	6.6×10^{-5}	5.7×10^{-1}	—	—	—	—
				SQA-CE	3704	1.00	—	—	—	463	2.1	5.1×10^2
				A-PSA	3704	1.00	6.6×10^{-5}	4.6	463	—	—	—
				A-PSA-CE	3704	1.00	6.6×10^{-5}	4	463	—	—	—
77	7	100	GP	SA	4000	0.19	1.7×10^{-5}	5.8	—	—	—	—
				SA-CE	4000	1.00	1.7×10^{-5}	6.4	—	—	—	—
				SQA-CE	3736	1.00	—	—	—	467	1.7×10^{-1}	2×10^3
				A-PSA	3736	1.00	1.7×10^{-5}	6.5	467	—	—	—
				A-PSA-CE	3736	1.00	1.7×10^{-5}	1.6	467	—	—	—
			TPE	SA	4000	0.22	1.7×10^{-5}	2.1	—	—	—	—
				SA-CE	4000	1.00	1.7×10^{-5}	1.3	—	—	—	—
				SQA-CE	3850	1.00	—	—	—	350	1.9×10^{-1}	2.2×10^3
				A-PSA	3850	1.00	1.7×10^{-5}	3.6	350	—	—	—
				A-PSA-CE	3850	1.00	1.7×10^{-5}	2.8	350	—	—	—

TABLE XV. (Continued.)

N	n	n_{reads}	Optimizer	$\mathcal{S}$	MCS	p_s	$\beta_{\min}$	$\beta_{\max}$	R	M	β	Γ
123	7	100	GP	SA	6000	0.29	1.5×10^{-5}	6.6	—	—	—	—
				SA-CE	6000	1.00	1.5×10^{-5}	3	—	—	—	—
				SQA-CE	5670	1.00	—	—	—	378	4.3×10^{-1}	1.1×10^3
				A-PSA	5670	1.00	1.5×10^{-5}	5.1	378	—	—	—
				A-PSA-CE	5670	1.00	1.5×10^{-5}	4.7	378	—	—	—
			TPE	SA	6000	0.31	1.5×10^{-5}	2.4	—	—	—	—
				SA-CE	6000	1.00	1.5×10^{-5}	2.6	—	—	—	—
				SQA-CE	5814	1.00	—	—	—	323	3.5×10^{-1}	4.8×10^2
				A-PSA	5814	1.00	1.5×10^{-5}	5	323	—	—	—
				A-PSA-CE	5814	1.00	1.5×10^{-5}	2.9	323	—	—	—
143	8	100	GP	SA	6000	0.12	3.5×10^{-6}	2.6	—	—	—	—
				SA-CE	6000	1.00	3.5×10^{-6}	7.6×10^{-2}	—	—	—	—
				SQA-CE	5178	1.00	—	—	—	863	1.7×10^{-1}	5.1×10^3
				A-PSA	5178	1.00	3.5×10^{-6}	6.6	863	—	—	—
				A-PSA-CE	5178	1.00	3.5×10^{-6}	3.7	863	—	—	—
			TPE	SA	6000	0.12	3.5×10^{-6}	3.9	—	—	—	—
				SA-CE	6000	1.00	3.5×10^{-6}	1.3×10^{-1}	—	—	—	—
				SQA-CE	5715	1.00	—	—	—	635	3.4×10^{-1}	6×10^3
				A-PSA	5715	1.00	3.5×10^{-6}	1.6×10^{-1}	635	—	—	—
				A-PSA-CE	5715	1.00	3.5×10^{-6}	3.2	635	—	—	—
253	8	100	GP	SA	6000	0.14	3.4×10^{-6}	3.5	—	—	—	—
				SA-CE	6000	1.00	3.4×10^{-6}	7.5×10^{-2}	—	—	—	—
				SQA-CE	5985	1.00	—	—	—	285	2.5×10^{-1}	1.8×10^3
				A-PSA	5985	1.00	3.4×10^{-6}	4.1	285	—	—	—
				A-PSA-CE	5985	1.00	3.4×10^{-6}	1.7	285	—	—	—
			TPE	SA	6000	0.14	3.4×10^{-6}	4.6	—	—	—	—
				SA-CE	6000	1.00	3.4×10^{-6}	8.3×10^{-2}	—	—	—	—
				SQA-CE	5960	1.00	—	—	—	745	2×10^{-1}	3×10^3
				A-PSA	5960	1.00	3.4×10^{-6}	5.7	745	—	—	—
				A-PSA-CE	5960	1.00	3.4×10^{-6}	2.5	745	—	—	—
451	9	100	GP	SA	6000	0.10	7.3×10^{-7}	2.6	—	—	—	—
				SA-CE	6000	0.66	7.3×10^{-7}	5.9×10^{-1}	—	—	—	—
				SQA-CE	5984	1.00	—	—	—	187	4.2×10^{-2}	6.5×10^3
				A-PSA	5984	0.76	7.3×10^{-7}	6.6	187	—	—	—
				A-PSA-CE	5984	0.91	7.3×10^{-7}	6.6×10^{-2}	187	—	—	—
			TPE	SA	6000	0.10	7.3×10^{-7}	4.3	—	—	—	—
				SA-CE	6000	0.71	7.3×10^{-7}	7.3×10^{-2}	—	—	—	—
				SQA-CE	5980	1.00	—	—	—	92	2.3×10^{-2}	9.8×10^3
				A-PSA	5980	0.68	7.3×10^{-7}	5.3	92	—	—	—
				A-PSA-CE	5980	0.94	7.3×10^{-7}	9.6×10^{-2}	92	—	—	—
511	9	100	GP	SA	6000	0.11	6.9×10^{-7}	5	—	—	—	—
				SA-CE	6000	0.96	6.9×10^{-7}	6.7×10^{-2}	—	—	—	—
				SQA-CE	5773	1.00	—	—	—	251	8.9×10^{-2}	3.1×10^3
				A-PSA	5773	0.99	6.9×10^{-7}	5.1	251	—	—	—
				A-PSA-CE	5773	1.00	6.9×10^{-7}	3.8	251	—	—	—
			TPE	SA	6000	0.14	6.9×10^{-7}	4.2	—	—	—	—
				SA-CE	6000	0.94	6.9×10^{-7}	2.7×10^{-1}	—	—	—	—
				SQA-CE	5882	1.00	—	—	—	346	6.5×10^{-2}	7.6×10^3
				A-PSA	5882	1.00	6.9×10^{-7}	6.2	346	—	—	—
				A-PSA-CE	5882	1.00	6.9×10^{-7}	2.3	346	—	—	—

TABLE XV. (Continued.)

N	n	n_{reads}	Optimizer	$\mathcal{S}$	MCS	p_s	$\beta_{\min}$	$\beta_{\max}$	R	M	β	Γ
899	10	100	GP	SA	6000	0.11	1.9×10^{-7}	3.4	—	—	—	—
				SA-CE	6000	0.28	1.9×10^{-7}	7.5×10^{-2}	—	—	—	—
				SQA-CE	5980	1.00	—	—	—	299	6.1×10^{-2}	4.7×10^3
				A-PSA	5980	1.00	1.9×10^{-7}	8.5×10^{-1}	299	—	—	—
				A-PSA-CE	5980	1.00	1.9×10^{-7}	2.8	299	—	—	—
			TPE	SA	6000	0.09	1.9×10^{-7}	1.3	—	—	—	—
				SA-CE	6000	0.32	1.9×10^{-7}	1.1×10^{-1}	—	—	—	—
				SQA-CE	5712	1.00	—	—	—	816	5.8×10^{-2}	1.4×10^4
				A-PSA	5712	1.00	1.9×10^{-7}	6.3	816	—	—	—
				A-PSA-CE	5712	1.00	1.9×10^{-7}	7.5	816	—	—	—
1011	10	100	GP	SA	6000	0.07	2.1×10^{-7}	6.1×10^{-1}	—	—	—	—
				SA-CE	6000	0.98	2.1×10^{-7}	6.6×10^{-2}	—	—	—	—
				SQA-CE	5640	1.00	—	—	—	564	1.2×10^{-1}	6.8×10^3
				A-PSA	5640	1.00	2.1×10^{-7}	1.8	564	—	—	—
				A-PSA-CE	5640	1.00	2.1×10^{-7}	4.1	564	—	—	—
			TPE	SA	6000	0.09	2.1×10^{-7}	6.6	—	—	—	—
				SA-CE	6000	0.97	2.1×10^{-7}	2.7	—	—	—	—
				SQA-CE	5964	1.00	—	—	—	497	1×10^{-1}	9.4×10^3
				A-PSA	5964	1.00	2.1×10^{-7}	1.9	497	—	—	—
				A-PSA-CE	5964	1.00	2.1×10^{-7}	4	497	—	—	—
1769	11	100	GP	SA	20000	0.05	4.6×10^{-8}	3.7	—	—	—	—
				SA-CE	20000	0.37	4.6×10^{-8}	7.5×10^{-2}	—	—	—	—
				SQA-CE	19516	1.00	—	—	—	1394	7.7×10^{-2}	2.6×10^4
				A-PSA	19516	1.00	4.6×10^{-8}	2.5	1394	—	—	—
				A-PSA-CE	19516	1.00	4.6×10^{-8}	3.9	1394	—	—	—
			TPE	SA	20000	0.06	4.6×10^{-8}	7.9×10^{-2}	—	—	—	—
				SA-CE	20000	0.39	4.6×10^{-8}	9.6×10^{-2}	—	—	—	—
				SQA-CE	18447	1.00	—	—	—	1677	1.9×10^{-1}	1.5×10^4
				A-PSA	18447	1.00	4.6×10^{-8}	6.4	1677	—	—	—
				A-PSA-CE	18447	1.00	4.6×10^{-8}	1.6	1677	—	—	—
2047	11	100	GP	SA	20000	0.05	4.5×10^{-8}	2.1	—	—	—	—
				SA-CE	20000	0.33	4.5×10^{-8}	7.5	—	—	—	—
				SQA-CE	19872	1.00	—	—	—	207	6.7×10^{-3}	3.3×10^4
				A-PSA	19872	0.63	4.5×10^{-8}	7.6×10^{-1}	207	—	—	—
				A-PSA-CE	19872	0.85	4.5×10^{-8}	7.6×10^{-2}	207	—	—	—
			TPE	SA	20000	0.06	4.5×10^{-8}	7	—	—	—	—
				SA-CE	20000	0.30	4.5×10^{-8}	1.6	—	—	—	—
				SQA-CE	19668	1.00	—	—	—	894	2.7×10^{-2}	3.4×10^4
				A-PSA	19668	0.78	4.5×10^{-8}	1.6×10^{-1}	894	—	—	—
				A-PSA-CE	19668	0.80	4.5×10^{-8}	1.3×10^{-1}	894	—	—	—
4009	12	100	GP	SA	20000	0.03	1.1×10^{-8}	4.2	—	—	—	—
				SA-CE	20000	0.13	1.1×10^{-8}	6.7×10^{-2}	—	—	—	—
				SQA-CE	19998	0.54	—	—	—	202	6.3×10^{-3}	8.1×10^4
				A-PSA	19998	0.30	1.1×10^{-8}	4.2	202	—	—	—
				A-PSA-CE	19998	0.49	1.1×10^{-8}	6.7×10^{-2}	202	—	—	—
			TPE	SA	20000	0.02	1.1×10^{-8}	4.2	—	—	—	—
				SA-CE	20000	0.11	1.1×10^{-8}	3.9	—	—	—	—
				SQA-CE	19926	0.55	—	—	—	243	5.6×10^{-3}	6.2×10^4
				A-PSA	19926	0.28	1.1×10^{-8}	4.3	243	—	—	—
				A-PSA-CE	19926	0.43	1.1×10^{-8}	6.7×10^{-2}	243	—	—	—

TABLE XV. (Continued.)

N	n	n_{reads}	Optimizer	$\mathcal{S}$	MCS	p_s	$\beta_{\min}$	$\beta_{\max}$	R	M	β	Γ
4087	12	100	GP	SA	20 000	0.05	1.1×10^{-8}	6.1	—	—	—	—
				SA-CE	20 000	0.17	1.1×10^{-8}	4.4×10^{-1}	—	—	—	—
				SQA-CE	19 899	0.50	—	—	—	201	4.5×10^{-3}	6.7×10^4
				A-PSA	19 899	0.25	1.1×10^{-8}	2.3	201	—	—	—
				A-PSA-CE	19 899	0.38	1.1×10^{-8}	7.5×10^{-2}	201	—	—	—
			TPE	SA	20 000	0.04	1.1×10^{-8}	5	—	—	—	—
				SA-CE	20 000	0.17	1.1×10^{-8}	5.5	—	—	—	—
				SQA-CE	19 838	0.48	—	—	—	218	4.6×10^{-3}	6.8×10^4
				A-PSA	19 838	0.22	1.1×10^{-8}	1.2	218	—	—	—
				A-PSA-CE	19 838	0.38	1.1×10^{-8}	9×10^{-2}	218	—	—	—
5723	13	100	GP	SA	20 000	0.02	2.7×10^{-9}	1.9	—	—	—	—
				SA-CE	20 000	0.10	2.7×10^{-9}	5.5	—	—	—	—
				SQA-CE	19 572	0.37	—	—	—	932	5.7×10^{-3}	1.6×10^5
				A-PSA	19 572	0.29	2.7×10^{-9}	7.6×10^{-2}	932	—	—	—
				A-PSA-CE	19 572	0.28	2.7×10^{-9}	7.5×10^{-2}	932	—	—	—
			TPE	SA	20 000	0.02	2.7×10^{-9}	6	—	—	—	—
				SA-CE	20 000	0.06	2.7×10^{-9}	4.3	—	—	—	—
				SQA-CE	19 952	0.37	—	—	—	688	4.3×10^{-3}	1.5×10^5
				A-PSA	19 952	0.22	2.7×10^{-9}	4.6	688	—	—	—
				A-PSA-CE	19 952	0.25	2.7×10^{-9}	4.8	688	—	—	—
8189	13	100	GP	SA	20 000	0.02	2.6×10^{-9}	6.5	—	—	—	—
				SA-CE	20 000	0.07	2.6×10^{-9}	1.8	—	—	—	—
				SQA-CE	19 220	0.36	—	—	—	1922	5.8×10^{-2}	1.6×10^5
				A-PSA	19 220	0.50	2.6×10^{-9}	6.2	1922	—	—	—
				A-PSA-CE	19 220	0.54	2.6×10^{-9}	6.6	1922	—	—	—
			TPE	SA	20 000	0.02	2.6×10^{-9}	4.6	—	—	—	—
				SA-CE	20 000	0.06	2.6×10^{-9}	6.5	—	—	—	—
				SQA-CE	19 800	0.39	—	—	—	264	4.7×10^{-3}	1.3×10^5
				A-PSA	19 800	0.27	2.6×10^{-9}	6.2	264	—	—	—
				A-PSA-CE	19 800	0.35	2.6×10^{-9}	6.5	264	—	—	—
15 853	14	100	GP	SA	20 000	0.02	6.7×10^{-10}	4.2	—	—	—	—
				SA-CE	20 000	0.04	6.7×10^{-10}	6.2	—	—	—	—
				SQA-CE	19 800	0.24	—	—	—	1100	1.7×10^{-2}	2.4×10^5
				A-PSA	19 800	0.35	6.7×10^{-10}	3.2	1100	—	—	—
				A-PSA-CE	19 800	0.32	6.7×10^{-10}	7.3	1100	—	—	—
			TPE	SA	20 000	0.01	6.7×10^{-10}	3.9	—	—	—	—
				SA-CE	20 000	0.04	6.7×10^{-10}	3.6	—	—	—	—
				SQA-CE	19 770	0.31	—	—	—	1977	1.7×10^{-2}	2.6×10^5
				A-PSA	19 770	0.36	6.7×10^{-10}	3.5	1977	—	—	—
				A-PSA-CE	19 770	0.39	6.7×10^{-10}	7.6×10^{-2}	1977	—	—	—
16 379	14	100	GP	SA	20 000	0.02	6.6×10^{-10}	5.1	—	—	—	—
				SA-CE	20 000	0.07	6.6×10^{-10}	4.1	—	—	—	—
				SQA-CE	20 000	0.40	—	—	—	200	6.5×10^{-4}	3.3×10^5
				A-PSA	20 000	0.11	6.6×10^{-10}	2.8	200	—	—	—
				A-PSA-CE	20 000	0.17	6.6×10^{-10}	6.6	200	—	—	—
			TPE	SA	20 000	0.01	6.6×10^{-10}	2	—	—	—	—
				SA-CE	20 000	0.07	6.6×10^{-10}	9.5×10^{-1}	—	—	—	—
				SQA-CE	19 872	0.42	—	—	—	207	7.3×10^{-4}	3.2×10^5
				A-PSA	19 872	0.09	6.6×10^{-10}	4.8	207	—	—	—
				A-PSA-CE	19 872	0.18	6.6×10^{-10}	1×10^{-1}	207	—	—	—

TABLE XV. (Continued.)

N	n	n_{reads}	Optimizer	$\mathcal{S}$	MCS	p_s	$\beta_{\min}$	$\beta_{\max}$	R	M	β	Γ
21 733	15	100	GP	SA	20 000	0.02	1.7×10^{-10}	6.7×10^{-1}	—	—	—	—
				SA-CE	20 000	0.02	1.7×10^{-10}	3.1	—	—	—	—
				SQA-CE	18 468	0.09	—	—	—	1539	7.6×10^{-3}	2.6×10^5
				A-PSA	18 468	0.13	1.7×10^{-10}	2.9	1539	—	—	—
				A-PSA-CE	18 468	0.16	1.7×10^{-10}	4.2	1539	—	—	—
			TPE	SA	20 000	0.01	1.7×10^{-10}	7.2	—	—	—	—
				SA-CE	20 000	0.02	1.7×10^{-10}	7.1	—	—	—	—
				SQA-CE	19 568	0.09	—	—	—	1223	1.1×10^{-2}	3.8×10^5
				A-PSA	19 568	0.13	1.7×10^{-10}	1.4	1223	—	—	—
				A-PSA-CE	19 568	0.10	1.7×10^{-10}	1.4	1223	—	—	—
32 743	15	100	GP	SA	20 000	0.02	1.7×10^{-10}	2.8	—	—	—	—
				SA-CE	20 000	0.05	1.7×10^{-10}	6.2	—	—	—	—
				SQA-CE	19 875	0.04	—	—	—	375	1.5×10^{-3}	1.9×10^5
				A-PSA	19 875	0.05	1.7×10^{-10}	5.9	375	—	—	—
				A-PSA-CE	19 875	0.05	1.7×10^{-10}	3.2	375	—	—	—
			TPE	SA	20 000	0.02	1.7×10^{-10}	5.1	—	—	—	—
				SA-CE	20 000	0.04	1.7×10^{-10}	3.9	—	—	—	—
				SQA-CE	19 024	0.02	—	—	—	1189	4.3×10^{-3}	2.4×10^5
				A-PSA	19 024	0.03	1.7×10^{-10}	3.1	1189	—	—	—
				A-PSA-CE	19 024	0.04	1.7×10^{-10}	1.8	1189	—	—	—
49 583	16	100	GP	SA	20 000	0.01	4.1×10^{-11}	6.4	—	—	—	—
				SA-CE	20 000	0.03	4.1×10^{-11}	5.4	—	—	—	—
				SQA-CE	19 894	0.02	—	—	—	203	2.5×10^{-4}	5.4×10^5
				A-PSA	19 894	0.02	4.1×10^{-11}	2	203	—	—	—
				A-PSA-CE	19 894	0.02	4.1×10^{-11}	3.8	203	—	—	—
			TPE	SA	20 000	0.00	4.1×10^{-11}	4.9	—	—	—	—
				SA-CE	20 000	0.02	4.1×10^{-11}	1.7×10^{-1}	—	—	—	—
				SQA-CE	19 680	0.02	—	—	—	492	9.8×10^{-4}	3.8×10^5
				A-PSA	19 680	0.02	4.1×10^{-11}	1.8	492	—	—	—
				A-PSA-CE	19 680	0.02	4.1×10^{-11}	1.6	492	—	—	—
59 989	16	100	GP	SA	20 000	0.02	4.1×10^{-11}	6.8	—	—	—	—
				SA-CE	20 000	0.02	4.1×10^{-11}	7.5	—	—	—	—
				SQA-CE	19 984	0.86	—	—	—	1249	2.7×10^{-3}	8.2×10^5
				A-PSA	19 984	0.96	4.1×10^{-11}	7.5×10^{-2}	1249	—	—	—
				A-PSA-CE	19 984	0.97	4.1×10^{-11}	7.2	1249	—	—	—
			TPE	SA	20 000	0.02	4.1×10^{-11}	5.2	—	—	—	—
				SA-CE	20 000	0.04	4.1×10^{-11}	3.3	—	—	—	—
				SQA-CE	18 760	0.88	—	—	—	1340	2.4×10^{-3}	5.9×10^5
				A-PSA	18 760	0.92	4.1×10^{-11}	6	1340	—	—	—
				A-PSA-CE	18 760	0.97	4.1×10^{-11}	5.5	1340	—	—	—
65 509	16	100	GP	SA	20 000	0.01	4.1×10^{-11}	7.5	—	—	—	—
				SA-CE	20 000	0.02	4.1×10^{-11}	4.7	—	—	—	—
				SQA-CE	18 928	0.02	—	—	—	1183	1.7×10^{-3}	3.4×10^5
				A-PSA	18 928	0.01	4.1×10^{-11}	7.5×10^{-2}	1183	—	—	—
				A-PSA-CE	18 928	0.02	4.1×10^{-11}	6.5	1183	—	—	—
			TPE	SA	20 000	0.01	4.1×10^{-11}	3.6	—	—	—	—
				SA-CE	20 000	0.02	4.1×10^{-11}	3	—	—	—	—
				SQA-CE	19 778	0.02	—	—	—	319	4.8×10^{-4}	1.8×10^5
				A-PSA	19 778	0.02	4.1×10^{-11}	1.1	319	—	—	—
				A-PSA-CE	19 778	0.02	4.1×10^{-11}	3.3×10^{-1}	319	—	—	—

TABLE XV. (Continued.)

N	n	n_{reads}	Optimizer	$\mathcal{S}$	MCS	p_s	$\beta_{\min}$	$\beta_{\max}$	R	M	β	Γ
94 037	17	100	GP	SA	20 000	0.01	1×10^{-11}	7.5	—	—	—	—
				SA-CE	20 000	0.01	1×10^{-11}	7.5	—	—	—	—
				SQA-CE	19 762	0.03	—	—	—	241	4.6×10^{-4}	1.4×10^6
				A-PSA	19 762	0.05	1×10^{-11}	6.5	241	—	—	—
				A-PSA-CE	19 762	0.06	1×10^{-11}	3.6	241	—	—	—
			TPE	SA	20 000	0.01	1×10^{-11}	3.7	—	—	—	—
				SA-CE	20 000	0.01	1×10^{-11}	2.6	—	—	—	—
				SQA-CE	19 437	0.03	—	—	—	589	1.3×10^{-3}	8.6×10^5
				A-PSA	19 437	0.06	1×10^{-11}	4	589	—	—	—
				A-PSA-CE	19 437	0.04	1×10^{-11}	4.4	589	—	—	—
131 069	17	100	GP	SA	20 000	0.00	1×10^{-11}	2.9	—	—	—	—
				SA-CE	20 000	0.01	1×10^{-11}	6.6×10^{-2}	—	—	—	—
				SQA-CE	19 800	0.01	—	—	—	220	1.1×10^{-4}	4.5×10^5
				A-PSA	19 800	0.01	1×10^{-11}	6.6	220	—	—	—
				A-PSA-CE	19 800	0.01	1×10^{-11}	6.6×10^{-2}	220	—	—	—
			TPE	SA	20 000	0.00	1×10^{-11}	4.4	—	—	—	—
				SA-CE	20 000	0.01	1×10^{-11}	2.1	—	—	—	—
				SQA-CE	19 866	0.01	—	—	—	231	1.6×10^{-4}	6.6×10^5
				A-PSA	19 866	0.01	1×10^{-11}	4.8×10^{-1}	231	—	—	—
				A-PSA-CE	19 866	0.01	1×10^{-11}	6.8×10^{-1}	231	—	—	—
159 317	18	100	GP	SA	20 000	0.00	2.5×10^{-12}	3.1	—	—	—	—
				SA-CE	20 000	0.01	2.5×10^{-12}	7.5	—	—	—	—
				SQA-CE	19 690	0.08	—	—	—	1969	9.6×10^{-4}	2.2×10^6
				A-PSA	19 690	0.11	2.5×10^{-12}	7.5×10^{-2}	1969	—	—	—
				A-PSA-CE	19 690	0.11	2.5×10^{-12}	1.5	1969	—	—	—
			TPE	SA	20 000	0.01	2.5×10^{-12}	4.7×10^{-1}	—	—	—	—
				SA-CE	20 000	0.02	2.5×10^{-12}	1.9	—	—	—	—
				SQA-CE	18 950	0.09	—	—	—	1895	9.6×10^{-4}	1.1×10^6
				A-PSA	18 950	0.11	2.5×10^{-12}	3.8	1895	—	—	—
				A-PSA-CE	18 950	0.14	2.5×10^{-12}	6.3	1895	—	—	—
262 099	18	100	GP	SA	20 000	0.01	2.5×10^{-12}	7.5	—	—	—	—
				SA-CE	20 000	0.01	2.5×10^{-12}	7.5	—	—	—	—
				SQA-CE	19 648	0.03	—	—	—	614	2.6×10^{-4}	2.5×10^6
				A-PSA	19 648	0.06	2.5×10^{-12}	1.4	614	—	—	—
				A-PSA-CE	19 648	0.05	2.5×10^{-12}	6.3	614	—	—	—
			TPE	SA	20 000	0.01	2.5×10^{-12}	2.1	—	—	—	—
				SA-CE	20 000	0.01	2.5×10^{-12}	9.7×10^{-1}	—	—	—	—
				SQA-CE	19 810	0.05	—	—	—	566	2.3×10^{-4}	2.5×10^6
				A-PSA	19 810	0.05	2.5×10^{-12}	1.3×10^{-1}	566	—	—	—
				A-PSA-CE	19 810	0.05	2.5×10^{-12}	6×10^{-1}	566	—	—	—
374 599	19	100	GP	SA	20 000	0.00	6.3×10^{-13}	4.5	—	—	—	—
				SA-CE	20 000	0.01	6.3×10^{-13}	7.5	—	—	—	—
				SQA-CE	19 740	0.00	—	—	—	282	2.3×10^{-4}	1.4×10^6
				A-PSA	19 740	0.00	6.3×10^{-13}	1.2	282	—	—	—
				A-PSA-CE	19 740	0.01	6.3×10^{-13}	7.5	282	—	—	—
			TPE	SA	20 000	0.00	6.3×10^{-13}	5.2	—	—	—	—
				SA-CE	20 000	0.00	6.3×10^{-13}	2.8	—	—	—	—
				SQA-CE	19 941	0.01	—	—	—	391	2.2×10^{-4}	4.7×10^6
				A-PSA	19 941	0.01	6.3×10^{-13}	5.2	391	—	—	—
				A-PSA-CE	19 941	0.01	6.3×10^{-13}	3.5	391	—	—	—

TABLE XV. (Continued.)

N	n	n_{reads}	Optimizer	S	MCS	p_s	$\beta_{\min}$	$\beta_{\max}$	R	M	β	Γ
376 289	19	100	GP	SA	20 000	0.01	6.3×10^{-13}	7.5	—	—	—	—
				SA-CE	20 000	0.00	6.3×10^{-13}	2.1	—	—	—	—
				SQA-CE	19 872	0.01	—	—	—	288	5.4×10^{-5}	1.1×10^6
				A-PSA	19 872	0.02	6.3×10^{-13}	2.3×10^{-1}	288	—	—	—
				A-PSA-CE	19 872	0.02	6.3×10^{-13}	1.6	288	—	—	—
			TPE	SA	20 000	0.01	6.3×10^{-13}	6	—	—	—	—
				SA-CE	20 000	0.00	6.3×10^{-13}	2.6	—	—	—	—
				SQA-CE	19 942	0.01	—	—	—	338	8×10^{-5}	4.8×10^6
				A-PSA	19 942	0.02	6.3×10^{-13}	2.3	338	—	—	—
				A-PSA-CE	19 942	0.02	6.3×10^{-13}	2.8×10^{-1}	338	—	—	—
524 281	19	100	GP	SA	20 000	0.00	6.3×10^{-13}	6.9	—	—	—	—
				SA-CE	20 000	0.01	6.3×10^{-13}	7.7×10^{-2}	—	—	—	—
				SQA-CE	19 320	0.00	—	—	—	690	3.1×10^{-4}	1.1×10^6
				A-PSA	19 320	0.01	6.3×10^{-13}	7.5×10^{-2}	690	—	—	—
				A-PSA-CE	19 320	0.00	6.3×10^{-13}	4.2	690	—	—	—
			TPE	SA	20 000	0.01	6.3×10^{-13}	1.9	—	—	—	—
				SA-CE	20 000	0.01	6.3×10^{-13}	5.7	—	—	—	—
				SQA-CE	18 976	0.01	—	—	—	1186	4.3×10^{-4}	5×10^6
				A-PSA	18 976	0.00	6.3×10^{-13}	5.7×10^{-1}	1186	—	—	—
				A-PSA-CE	18 976	0.01	6.3×10^{-13}	3.5	1186	—	—	—
881 851	20	100	GP	SA	20 000	0.00	1.6×10^{-13}	7.6×10^{-2}	—	—	—	—
				SA-CE	20 000	0.00	1.6×10^{-13}	2.2	—	—	—	—
				SQA-CE	19 800	0.01	—	—	—	220	9.8×10^{-5}	1.1×10^7
				A-PSA	19 800	0.01	1.6×10^{-13}	7.2	220	—	—	—
				A-PSA-CE	19 800	0.01	1.6×10^{-13}	7.5	220	—	—	—
			TPE	SA	20 000	0.00	1.6×10^{-13}	4.2	—	—	—	—
				SA-CE	20 000	0.00	1.6×10^{-13}	2.9	—	—	—	—
				SQA-CE	19 215	0.01	—	—	—	915	1.5×10^{-4}	1×10^7
				A-PSA	19 215	0.02	1.6×10^{-13}	1.5	915	—	—	—
				A-PSA-CE	19 215	0.01	1.6×10^{-13}	6.3×10^{-1}	915	—	—	—
1 005 973	20	100	GP	SA	20 000	0.00	1.6×10^{-13}	2	—	—	—	—
				SA-CE	20 000	0.00	1.6×10^{-13}	5.4	—	—	—	—
				SQA-CE	19 320	0.01	—	—	—	1610	2.6×10^{-4}	2.8×10^6
				A-PSA	19 320	0.02	1.6×10^{-13}	3.1	1610	—	—	—
				A-PSA-CE	19 320	0.02	1.6×10^{-13}	7.6×10^{-2}	1610	—	—	—
			TPE	SA	20 000	0.01	1.6×10^{-13}	4.8	—	—	—	—
				SA-CE	20 000	0.00	1.6×10^{-13}	2.4	—	—	—	—
				SQA-CE	19 844	0.02	—	—	—	242	2.3×10^{-5}	6.1×10^6
				A-PSA	19 844	0.01	1.6×10^{-13}	2.3	242	—	—	—
				A-PSA-CE	19 844	0.02	1.6×10^{-13}	4.2	242	—	—	—
1 028 171	20	100	GP	SA	20 000	0.01	1.6×10^{-13}	7.7×10^{-2}	—	—	—	—
				SA-CE	20 000	0.01	1.6×10^{-13}	3.7	—	—	—	—
				SQA-CE	19 040	0.14	—	—	—	1360	1.3×10^{-4}	2.3×10^6
				A-PSA	19 040	0.12	1.6×10^{-13}	6.8	1360	—	—	—
				A-PSA-CE	19 040	0.13	1.6×10^{-13}	6.3	1360	—	—	—
			TPE	SA	20 000	0.01	1.6×10^{-13}	4.9	—	—	—	—
				SA-CE	20 000	0.01	1.6×10^{-13}	4.5	—	—	—	—
				SQA-CE	19 334	0.09	—	—	—	1381	1.8×10^{-4}	2.7×10^6
				A-PSA	19 334	0.12	1.6×10^{-13}	2.9	1381	—	—	—
				A-PSA-CE	19 334	0.13	1.6×10^{-13}	3.3	1381	—	—	—

TABLE XV. (Continued.)

N	n	n_{reads}	Optimizer	$\mathcal{S}$	MCS	p_s	$\beta_{\min}$	$\beta_{\max}$	R	M	β	Γ
1 048 561 20	100		GP	SA	20 000	0.01	1.6×10^{-13}	7.6×10^{-2}	—	—	—	—
				SA-CE	20 000	0.00	1.6×10^{-13}	3.3	—	—	—	—
				SQA-CE	19 920	0.01	—	—	—	249	2.4×10^{-5}	1.1×10^7
				A-PSA	19 920	0.02	1.6×10^{-13}	6.1	249	—	—	—
				A-PSA-CE	19 920	0.02	1.6×10^{-13}	5.2	249	—	—	—
			TPE	SA	20 000	0.01	1.6×10^{-13}	6.9	—	—	—	—
				SA-CE	20 000	0.01	1.6×10^{-13}	1.5	—	—	—	—
				SQA-CE	19 129	0.00	—	—	—	1739	1.8×10^{-4}	5×10^6
				A-PSA	19 129	0.01	1.6×10^{-13}	6.9	1739	—	—	—
				A-PSA-CE	19 129	0.01	1.6×10^{-13}	4.6	1739	—	—	—
1 794 953 21	100		GP	SA	20 000	0.01	4×10^{-14}	7.5	—	—	—	—
				SA-CE	20 000	0.00	4×10^{-14}	7.1×10^{-1}	—	—	—	—
				SQA-CE	19 890	0.01	—	—	—	221	5.1×10^{-5}	2.1×10^7
				A-PSA	19 890	0.01	4×10^{-14}	7.5	221	—	—	—
				A-PSA-CE	19 890	0.01	4×10^{-14}	7.5×10^{-2}	221	—	—	—
			TPE	SA	20 000	0.00	4×10^{-14}	7.3	—	—	—	—
				SA-CE	20 000	0.01	4×10^{-14}	1.1	—	—	—	—
				SQA-CE	19 936	0.01	—	—	—	1424	2.5×10^{-4}	1.4×10^7
				A-PSA	19 936	0.01	4×10^{-14}	1.2×10^{-1}	1424	—	—	—
				A-PSA-CE	19 936	0.01	4×10^{-14}	7.2	1424	—	—	—
2 093 809 21	100		GP	SA	20 000	0.00	4×10^{-14}	1.6	—	—	—	—
				SA-CE	20 000	0.01	4×10^{-14}	7.5	—	—	—	—
				SQA-CE	18 440	0.00	—	—	—	1844	1.4×10^{-4}	1.4×10^7
				A-PSA	18 440	0.00	4×10^{-14}	5.6	1844	—	—	—
				A-PSA-CE	18 440	0.00	4×10^{-14}	1.1	1844	—	—	—
			TPE	SA	20 000	0.00	4×10^{-14}	1.2	—	—	—	—
				SA-CE	20 000	0.00	4×10^{-14}	1	—	—	—	—
				SQA-CE	19 305	0.00	—	—	—	715	4.2×10^{-5}	5.7×10^6
				A-PSA	19 305	0.00	4×10^{-14}	3.8	715	—	—	—
				A-PSA-CE	19 305	0.01	4×10^{-14}	3.2	715	—	—	—
3 819 047 22	100		GP	SA	20 000	0.00	4×10^{-14}	2.9	—	—	—	—
				SA-CE	20 000	0.00	4×10^{-14}	8.3×10^{-1}	—	—	—	—
				SQA-CE	18 744	0.00	—	—	—	1562	7.4×10^{-5}	1.6×10^7
				A-PSA	18 744	0.00	4×10^{-14}	1.2	1562	—	—	—
				A-PSA-CE	18 744	0.005	4×10^{-14}	7.5	1562	—	—	—
			TPE	SA	20 000	0.00	4×10^{-14}	5	—	—	—	—
				SA-CE	20 000	0.00	4×10^{-14}	1.3	—	—	—	—
				SQA-CE	19 894	0.005	—	—	—	203	1.1×10^{-5}	4.6×10^6
				A-PSA	19 894	0.005	4×10^{-14}	1×10^{-1}	203	—	—	—
				A-PSA-CE	19 894	0.005	4×10^{-14}	3.9×10^{-1}	203	—	—	—
3 819 047 22	100		GP	SA	20 000	0.00	9.9×10^{-15}	4.5	—	—	—	—
				SA-CE	20 000	0.00	9.9×10^{-15}	3.8	—	—	—	—
				SQA-CE	19 838	0.01	—	—	—	218	1.3×10^{-5}	2.8×10^7
				A-PSA	19 838	0.01	9.9×10^{-15}	7.5	218	—	—	—
				A-PSA-CE	19 838	0.01	9.9×10^{-15}	7.5	218	—	—	—
			TPE	SA	20 000	0.00	9.9×10^{-15}	2.6	—	—	—	—
				SA-CE	20 000	0.00	9.9×10^{-15}	6	—	—	—	—
				SQA-CE	19 520	0.01	—	—	—	976	3.5×10^{-5}	3.7×10^7
				A-PSA	19 520	0.02	9.9×10^{-15}	5	976	—	—	—
				A-PSA-CE	19 520	0.02	9.9×10^{-15}	2	976	—	—	—

TABLE XV. (Continued.)

N	n	n_{reads}	Optimizer	$\mathcal{S}$	MCS	p_s	$\beta_{\min}$	$\beta_{\max}$	R	M	β	Γ
8 219 999 23	100	GP	SA	20 000	0.01	2.5×10^{-15}	7.5	—	—	—	—	—
			SA-CE	20 000	0.01	2.5×10^{-15}	7.6×10^{-2}	—	—	—	—	—
			SQA-CE	18 914	0.17	—	—	—	1351	2.3×10^{-5}	8.4×10^7	—
			A-PSA	18 914	0.10	2.5×10^{-15}	1.4	1351	—	—	—	—
		TPE	A-PSA-CE	18 914	0.11	2.5×10^{-15}	4.3	1351	—	—	—	—
			SA	20 000	0.00	2.5×10^{-15}	4.8	—	—	—	—	—
			SA-CE	20 000	0.01	2.5×10^{-15}	3.9	—	—	—	—	—
			SQA-CE	18 928	0.20	—	—	—	1183	3.1×10^{-5}	8.1×10^7	—
	400	GP	A-PSA	18 928	0.11	2.5×10^{-15}	5.3	1183	—	—	—	—
			A-PSA-CE	18 928	0.12	2.5×10^{-15}	1.4	1183	—	—	—	—
			SA	20 000	0.00	2.5×10^{-15}	2.6	—	—	—	—	—
			SA-CE	20 000	0.0025	2.5×10^{-15}	7.5	—	—	—	—	—
		TPE	SQA-CE	19 012	0.155	—	—	—	1358	1.6×10^{-5}	1.8×10^7	—
			A-PSA	19 012	0.095	2.5×10^{-15}	2.4	1358	—	—	—	—
			A-PSA-CE	19 012	0.090	2.5×10^{-15}	5.2	1358	—	—	—	—
			SA	20 000	0.005	2.5×10^{-15}	5.6	—	—	—	—	—
8 219 999 23	100	GP	SA-CE	20 000	0.005	2.5×10^{-15}	6.4	—	—	—	—	—
			SQA-CE	19 824	0.1525	—	—	—	1239	1.5×10^{-5}	7.3×10^7	—
			A-PSA	19 824	0.10	2.5×10^{-15}	5.4	1239	—	—	—	—
			A-PSA-CE	19 824	0.0875	2.5×10^{-15}	1.5	1239	—	—	—	—

Taken together, the $N = 59\,989$ similarity profile strongly supports SQA schedules designed for *sparse, frustration-dominated glasses*: EA-like offsets, ASI-inspired plateau timing, and restart-heavy protocols are expected to yield the best performance.

APPENDIX D: OPTIMAL HYPERPARAMETER VALUES FROM BAYESIAN OPTIMIZATION

The optimal parameter values obtained from Bayesian optimization are summarized in Table XV. The nominal amount of Monte Carlo steps assigned to each optimizer is the same and it is equal to the number of Monte Carlo sweeps performed by the SA and SA-CE solvers. Differences in the number of the effective number of Monte Carlo steps reported in the table arise from the integer-valued nature of parameters M , R , and J .

- [1] A. J. Menezes, P. C. Van Oorschot, and S. A. Vanstone, *Handbook of Applied Cryptography* (CRC Press, Boca Raton, FL, 2018).
- [2] M. Tesoro, I. Siloi, D. Jaschke, G. Magnifico, and S. Montangero, Integer factorization via tensor-network Schnorr’s sieving, *Phys. Rev. A* **113**, 032418 (2026).
- [3] W. Stallings, *Cryptography and Network Security—Principles and Practice* (Pearson Education India, New Delhi, India, 2017), 7th ed.
- [4] D. J. Bernstein and T. Lange, Post-quantum cryptography, *Nature (London)* **549**, 188 (2017).
- [5] K. Rabah, Review of methods for integer factorization applied to cryptography, *J. Appl. Sci.* **6**, 458 (2006).
- [6] H. Riesel, J. Oesterlé, and A. Weinstein, *Prime Numbers and Computer Methods for Factorization* (Springer, Boston, MA, 1994), Vol. 2.
- [7] F. K. Wilhelm and R. Steinwandt, Status of quantum computer development, Technical Report 477, 2024, https://www.bsi.bund.de/SharedDocs/Downloads/DE/BSI/Publikationen/Studien/Quantencomputer/Entwicklungsstand_QC_V_2_1.html.

- [8] R. Mengoni, D. Ottaviani, and P. Iorio, Breaking RSA security with a low noise D-Wave 2000Q quantum annealer: Computational times, limitations and prospects, [arXiv:2005.02268](#).
- [9] P. W. Shor, Algorithms for quantum computation: Discrete logarithms and factoring, in *Proceedings 35th Annual Symposium on Foundations of Computer Science* (IEEE, New York, 1994), p. 124.
- [10] R. Babbush, R. King, S. Boixo, W. Huggins, T. Khattar, G. H. Low, J. R. McClean, T. O’Brien, and N. C. Rubin, Grand challenge of quantum applications, [PRX Quantum](#) **7**, 020101 (2026).
- [11] U. Skosana and M. Tame, Demonstration of Shor’s factoring algorithm for $N = 21$ on IBM quantum processors, [Sci. Rep.](#) **11**, 16599 (2021).
- [12] J. Palsberg, J. Cong, Y. Ding, B. Fefferman, M. Qureshi, G. S. Ravi, K. N. Smith, H. Wang, X. Wu, and H. Yuen, Computer science challenges in quantum computing: Early fault-tolerance and beyond, [arXiv:2601.20247](#).
- [13] J. Preskill, Quantum Computing in the NISQ era and beyond, [Quantum](#) **2**, 79 (2018).
- [14] D. Willsch, P. Hanussek, G. Hoever, M. Willsch, F. Jin, H. De Raedt, and K. Michielsen, The state of factoring on quantum computers, [NIC Symposium Proc.](#) **52**, 423 (2025).
- [15] B. Wang, X. Yang, and D. Zhang, Research on quantum annealing integer factorization based on different columns, [Front. Phys.](#) **10**, 914578 (2022).
- [16] Š. Grigaliūnas and R. Brūzgienė, Towards a unified quantum risk assessment, [Electronics](#) **14**, 3338 (2025).
- [17] R. Martoňák, G. E. Santoro, and E. Tosatti, Quantum annealing by the path-integral Monte Carlo method: The two-dimensional random Ising model, [Phys. Rev. B](#) **66**, 094203 (2002).
- [18] T. Albash and D. A. Lidar, Adiabatic quantum computation and quantum annealing: Theory and practice, [Rev. Mod. Phys.](#) **90**, 015002 (2018).
- [19] V. Bapst and G. Semerjian, Thermal, quantum and simulated quantum annealing: Analytical comparisons for simple models, [J. Phys.: Conf. Ser.](#), **473**, 012011 (2013).
- [20] E. Crosson and A. W. Harrow, Simulated quantum annealing can be exponentially faster than classical simulated annealing, in *Proceedings of the 2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, Piscataway, NJ, 2016), pp. 714–723.
- [21] C. Gomes, J. P. Fernandes, G. Falcao, S. Kar, and S. Tayur, A systematic mapping study on quantum and quantum-inspired algorithms in operations research, [ACM Comput. Surv.](#) **57**, 1 (2025).
- [22] H. M. Waidyasooriya and M. Hariyama, Temporal and spatial parallel processing of simulated quantum annealing on a multicore CPU, [J. Supercomput.](#) **78**, 8733 (2022).
- [23] H. M. Waidyasooriya and M. Hariyama, A GPU-based quantum annealing simulator for fully-connected Ising models utilizing spatial and temporal parallelism, [IEEE Access](#) **8**, 67929 (2020).
- [24] B. Tasseff, T. Albash, Z. Morrell, M. Vuffray, A. Y. Lokhov, S. Misra, and C. Coffrin, On the emerging potential of quantum annealing hardware for combinatorial optimization, [J. Heuristics](#) **30**, 325 (2024).
- [25] H. Oshiyama and M. Ohzeki, Benchmark of quantum-inspired heuristic solvers for quadratic unconstrained binary optimization, [Sci. Rep.](#) **12**, 2146 (2022).
- [26] G. Schaller and R. Schützhold, The role of symmetries in adiabatic quantum algorithms, [Quantum Inf. Comput.](#) **10**, 109 (2010).
- [27] S. Kirkpatrick, C. Gelatt, and M. Vecchi, Optimization by simulated annealing, [Science](#) **220**, 671 (1983).
- [28] E. Onbaşoğlu and L. Özdamar, Parallel simulated annealing algorithms in global optimization, [J. Global Optim.](#) **19**, 27 (2001).
- [29] A. M. Ferreiro, J. García, J. G. López-Salas, and C. Vázquez, An efficient implementation of parallel simulated annealing algorithm in GPUs, [J. Global Optim.](#) **57**, 863 (2013).
- [30] A. L. B. Almeida, J. de C. Lima, and M. A. M. Carvalho, Systematic literature review on parallel trajectory-based metaheuristics, [ACM Comput. Surv.](#) **55**, 1 (2023).
- [31] P. Siarry, *Metaheuristics* (Springer, Cham, 2016), p. 19.
- [32] C. Baldassi, C. Borgs, J. T. Chayes, A. Ingrosso, C. Lucibello, L. Saglietti, and R. Zecchina, Unreasonable effectiveness of learning neural networks: From accessible states and robust ensembles to basic algorithmic schemes, [Proc. Natl. Acad. Sci. USA](#) **113**, E7655 (2016).
- [33] C. Baldassi and R. Zecchina, Efficiency of quantum vs. classical annealing in nonconvex learning problems, [Proc. Natl. Acad. Sci. USA](#) **115**, 1457 (2018).

- [34] V. Gripon, M. Löwe, and F. Vermet, Some remarks on replicated simulated annealing, *J. Stat. Phys.* **182**, 51 (2021).
- [35] G. E. Santoro, R. Martonák, E. Tosatti, and R. Car, Theory of quantum annealing of an Ising spin glass, *Science* **295**, 2427 (2002).
- [36] T. F. Rønnow, Z. Wang, J. Job, S. Boixo, S. V. Isakov, D. Wecker, J. M. Martinis, D. A. Lidar, and M. Troyer, Defining and detecting quantum speedup, *Science* **345**, 420 (2014).
- [37] S. V. Isakov, G. Mazzola, V. N. Smelyanskiy, Z. Jiang, S. Boixo, H. Neven, and M. Troyer, Understanding quantum tunneling through quantum Monte Carlo simulations, *Phys. Rev. Lett.* **117**, 180402 (2016).
- [38] E. Raimondo, E. Garzón, Y. Shao, A. Grimaldi, S. Chiappini, R. Tomasello, N. Davila-Melendez, J. A. Katine, M. Carpentieri, M. Chiappini, *et al.*, High-performance and reliable probabilistic Ising machine based on simulated quantum annealing, *Phys. Rev. X* **15**, 041001 (2025).
- [39] S. Chowdhury, N. A. Aadit, A. Grimaldi, E. Raimondo, A. Raut, P. A. Lott, J. H. Mentink, M. M. Rams, F. Ricci-Tersenghi, M. Chiappini, *et al.*, Pushing the boundary of quantum advantage in hard combinatorial optimization with probabilistic computers, *Nat. Commun.* **16**, 9193 (2025).
- [40] T. Okuyama, M. Hayashi, and M. Yamaoka, An Ising computer based on simulated quantum annealing by path integral Monte Carlo method, in *2017 IEEE International Conference on Rebooting Computing (ICRC)* (IEEE, Washington, DC, 2017), p. 1.
- [41] L. T. Brady and W. van Dam, Quantum Monte Carlo simulations of tunneling in quantum adiabatic optimization, *Phys. Rev. A* **93**, 032304 (2016).
- [42] L. Hormozi, E. W. Brown, G. Carleo, and M. Troyer, Nonstoquastic Hamiltonians and quantum annealing of an Ising spin glass, *Phys. Rev. B* **95**, 184416 (2017).
- [43] C. J. Burges, Factoring as optimization, Technical Report MSR-TR-2002-83, 2002, <https://www.microsoft.com/en-us/research/publication/factoring-as-optimization/>.
- [44] A. Grimaldi, K. Selcuk, N. A. Aadit, K. Kobayashi, Q. Cao, S. Chowdhury, G. Finocchio, S. Kanai, H. Ohno, S. Fukami, and K. Y. Camsari, Experimental evaluation of simulated quantum annealing with MTJ-augmented p-bits, in *Proceedings of the 2022 IEEE International Electron Devices Meeting (IEDM)* (IEEE, Piscataway, NJ, 2022), pp. 22.4.1–22.4.4.
- [45] D. Bernal, Qubovert: A Python library for QUBO conversion and manipulation (2022), <https://qubovert.readthedocs.io/en/latest/>.
- [46] K. Nishimura, Y. Sakamoto, T. Shimizu, K. Suzuki, and Y. Yamashiro, OpenJij: An open-source framework for Ising and QUBO optimization (2021), <https://jij-inc.github.io/OpenJij-Reference-Page/>.
- [47] D-Wave, Dwave-neal: D-Wave systems Python library for classical sampling methods (2022), <https://docs.ocean.dwavesys.com/projects/neal/en/latest/>.
- [48] Scikit-optimize: Sequential model-based optimization in Python (2016), <https://scikit-optimize.github.io/stable/>.
- [49] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, Optuna: A next-generation hyperparameter optimization framework, in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (ACM, New York, NY, 2019), pp. 2623–2631.
- [50] D. Dobrynin, A. Renaudineau, M. Hizzani, D. Strukov, M. Mohseni, and J. P. Strachan, Energy landscapes of combinatorial optimization in Ising machines, *Phys. Rev. E* **110**, 045308 (2024).
- [51] National Institute of Standards and Technology, *Digital Signature Standard (DSS)*, *Federal Information Processing Standards Publication FIPS 186-5* (U.S. Department of Commerce, Washington, DC, 2023).
- [52] Z. Luo, Y. Yuan, Z. Wang, M. Lai, and P. Lu, General integer factorization algorithm based on Ising machine, *EPJ Quantum Technol.* **13**, 1 (2026).
- [53] D. Zhang, H. Wang, S. Li, and B. Wang, Progress in the prime factorization of large numbers, *J. Supercomput.* **80**, 11382 (2024).
- [54] S. Jiang, K. A. Britt, A. J. McCaskey, T. S. Humble, and S. Kais, Quantum annealing for prime factorization, *Sci. Rep.* **8**, 17667 (2018).
- [55] K. Jun and H. Lee, HUBO and QUBO models for prime factorization, *Sci. Rep.* **13**, 10080 (2023).
- [56] F. Glover, G. Kochenberger, and Y. Du, Quantum Bridge Analytics I: A tutorial on formulating and using QUBO models, *Ann. Oper. Res.* **314**, 141 (2022).

- [57] Q.-G. Zeng, X.-P. Cui, B. Liu, Y. Wang, P. Mosharev, and M.-H. Yung, Performance of quantum annealing inspired algorithms for combinatorial optimization problems, *Commun. Phys.* **7**, 249 (2024).
- [58] M. Aramon, G. Rosenberg, E. Valiante, T. Miyazawa, H. Tamura, and H. G. Katzgraber, Physics-inspired optimization for quadratic unconstrained problems using a digital annealer, *Front. Phys.* **7**, 48 (2019).
- [59] M. Sobhani, Y. Chai, T. Hartung, and K. Jansen, Variational quantum eigensolver approach to prime factorization on IBM's noisy intermediate scale quantum computer, *Phys. Rev. A* **111**, 042413 (2025).
- [60] S. McGarraghy and M. Venkova, Quantum-inspired heuristics, in *Handbook of Heuristics*, edited by R. Martí, P. M. Pardalos, and M. G. Resende (Springer Nature Switzerland, Cham, 2025), p. 259.
- [61] C. Hong, Z. Pei, Q. Wang, S. Yang, J. Yu, and C. Wang, Quantum attack on RSA by D-Wave Advantage: A first break of 80-bit RSA, *Sci. China Inf. Sci.* **68**, 129501 (2025).
- [62] J. Ding, G. Spallitta, and R. Sebastiani, Effective prime factorization via quantum annealing by modular locally-structured embedding, *Sci. Rep.* **14**, 3518 (2024).
- [63] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse Ising model, *Phys. Rev. E* **58**, 5355 (1998).
- [64] S. Morita and H. Nishimori, Mathematical foundation of quantum annealing, *J. Math. Phys.* **49**, 125210 (2008).
- [65] H. Nishimori and Y. Nonomura, Quantum effects in neural networks, *J. Phys. Soc. Jpn.* **65**, 3780 (1996).
- [66] A. Das and B. K. Chakrabarti, *Quantum Annealing And Related Optimization Methods* (Springer, Berlin, 2005), Vol. 679.
- [67] A. Rajak, S. Suzuki, A. Dutta, and B. K. Chakrabarti, Quantum annealing: An overview, *Philos. Trans. R. Soc. A* **381**, 20210417 (2023).
- [68] V. Mukherjee and B. K. Chakrabarti, Multivariable optimization: Quantum annealing and computation, *Eur. Phys. J.: Spec. Top.* **224**, 17 (2015).
- [69] S. Tanaka, R. Tamura, and B. K. Chakrabarti, *Quantum Spin Glasses, Annealing and Computation* (Cambridge University Press, Cambridge, 2017), 1st ed.
- [70] B. K. Chakrabarti, A. Dutta, and P. Sen, *Quantum Ising Phases and Transitions in Transverse Ising Models* (Springer, Berlin, 2008), Chap. 3, pp. 50–57.
- [71] Y. Okabe and M. Kikuchi, *Quantum Monte Carlo Methods in Equilibrium and Nonequilibrium Systems*, edited by M. Suzuki (Springer, Berlin, 1987), Chap. 3, pp. 34–40.
- [72] J. E. Gubernatis, N. Kawashima, and P. Werner, *Quantum Monte Carlo Methods: Algorithms for Lattice Models* (Cambridge University Press, Cambridge, 2016), Chap. 5, pp. 84–118.
- [73] D. R. Jones, A taxonomy of global optimization methods based on response surfaces, *J. Global Optim.* **21**, 345 (2001).
- [74] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/hj46-qw5t> for details on the QUBO construction, solver benchmarks, and convergence analysis of the path-integral Monte Carlo simulations.
- [75] G. E. Santoro and E. Tosatti, Optimization using quantum mechanics: Quantum annealing through adiabatic evolution, *J. Phys. A: Math. Gen.* **39**, R393 (2006).
- [76] T. Albash and J. Marshall, Comparing relaxation mechanisms in quantum and classical transverse-field annealing, *Phys. Rev. Appl.* **15**, 014029 (2021).
- [77] M. Bernaschi, I. González-Adalid Pemartín, V. Martín-Mayor, and G. Parisi, The quantum transition of the two-dimensional Ising spin glass, *Nature (London)* **631**, 749 (2024).
- [78] B. Heim, T. F. Rønnow, S. V. Isakov, and M. Troyer, Quantum versus classical annealing of Ising spin glasses, *Science* **348**, 215 (2015).
- [79] C. R. McLeod and M. Sasdelli, Benchmarking D-Wave quantum annealers: Spectral gap scaling of maximum cardinality matching problems, in *Proceedings of the International Conference on Computational Science* (Springer, Cham, 2022), p. 150.
- [80] D. Volpe, G. A. Cirillo, M. Zamboni, and G. Turvani, Integration of simulated quantum annealing in parallel tempering and population annealing for heterogeneous-profile QUBO exploration, *IEEE Access* **11**, 30390 (2023).
- [81] E. Crosson and D. Lidar, Prospects for quantum enhancement with diabatic quantum annealing, *Nat. Rev. Phys.* **3**, 466 (2021).

- [82] D. De Santis, S. Tirone, S. Marmi, and V. Giovannetti, Optimized QUBO formulation methods for quantum computing, *Quantum Sci. Technol.* **11**, 015056 (2026).
- [83] D. Volpe, G. Orlandi, and G. Turvani, Improving the solving of optimization problems: A comprehensive review of quantum approaches, *Quantum Rep.* **7**, 3 (2025).
- [84] F. R. K. Chung, *Spectral Graph Theory* (American Mathematical Society, Providence, 1997).
- [85] D. Sherrington and S. Kirkpatrick, Solvable model of a spin-glass, *Phys. Rev. Lett.* **35**, 1792 (1975).
- [86] L. Petrov, Lectures on random matrices (Spring 2025), Lecture 2: Wigner semicircle law, course notes (2025), <https://lpetrov.cc/rmt25/rmt25-notes/rmt2025-l02.pdf>.
- [87] D. Sherrington and S. Kirkpatrick, 50 years of spin glass theory, *Nat. Rev. Phys.* **7**, 528 (2025).
- [88] J. J. Hopfield, Neural networks and physical systems with emergent collective computational abilities, *Proc. Natl. Acad. Sci. USA* **79**, 2554 (1982).
- [89] S. F. Edwards and P. W. Anderson, Theory of spin glasses, *J. Phys. F: Met. Phys.* **5**, 965 (1975).
- [90] S. F. Edwards and P. W. Anderson, Theory of spin glasses. II, *J. Phys. F: Met. Phys.* **6**, 1927 (1976).
- [91] R. Moessner and A. P. Ramirez, Geometrical frustration, *Phys. Today* **59**(2), 24 (2006).
- [92] C. Nisoli, R. Moessner, and P. Schiffer, Colloquium: Artificial spin ice: Designing and imaging magnetic frustration, *Rev. Mod. Phys.* **85**, 1473 (2013).
- [93] B. B. Bhattacharya and S. Mukherjee, Inference in Ising models, *Bernoulli* **24**, 493 (2018).
- [94] S. Chatterjee, Features of a spin glass in the random field Ising model, *Commun. Math. Phys.* **405**, 93 (2024).
- [95] A. R. Feier, Methods of proof in random matrix theory, Master's thesis, Harvard University, 2012.
- [96] L. F. Cugliandolo, Advanced statistical physics: Frustration, Lecture Notes (2017), <https://www.lpthe.jussieu.fr/~leticia/TEACHING/Master2017/frustration.pdf>.
- [97] F. Wegner, Inverse participation ratio in $2 + \varepsilon$ dimensions, *Z. Phys. B* **36**, 209 (1980).
- [98] B. Kramer and A. MacKinnon, Localization: Theory and experiment, *Rep. Prog. Phys.* **56**, 1469 (1993).
- [99] F. Evers and A. D. Mirlin, Anderson transitions, *Rev. Mod. Phys.* **80**, 1355 (2008).
- [100] F. Haake, *Quantum Signatures of Chaos* (Springer, Berlin, Heidelberg, 2010), 3rd ed.
- [101] I. J. Farkas, I. Derényi, A.-L. Barabási, and T. Vicsek, Spectra of “real-world” graphs: Beyond the semicircle law, *Phys. Rev. E* **64**, 026704 (2001).
- [102] A. J. Bray and M. A. Moore, Replica theory of quantum spin glasses, *J. Phys. C* **14**, 2629 (1981).
- [103] G. Parisi, A sequence of approximated solutions to the S-K model for spin glasses, *J. Phys. A: Math. Gen.* **13**, L115 (1980).
- [104] M. Mézard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).
- [105] K. Binder and A. P. Young, Spin glasses: Experimental facts, theoretical concepts, and open questions, *Rev. Mod. Phys.* **58**, 801 (1986).
- [106] D. J. Amit, H. Gutfreund, and H. Sompolinsky, Storing infinite numbers of patterns in a spin-glass model of neural networks, *Phys. Rev. Lett.* **55**, 1530 (1985).
- [107] D. J. Amit, H. Gutfreund, and H. Sompolinsky, Statistical mechanics of neural networks near saturation, *Ann. Phys.* **173**, 30 (1987).
- [108] J. Lin, Divergence measures based on the Shannon entropy, *IEEE Trans. Inf. Theory* **37**, 145 (2002).
- [109] F. J. Massey, Jr., The Kolmogorov-Smirnov test for goodness of fit, *J. Am. Stat. Assoc.* **46**, 68 (1951).
- [110] R. B. Cattell, The scree test for the number of factors, *Multivar. Behav. Res.* **1**, 245 (1966).
- [111] Z. Bai and J. W. Silverstein, *Spectral Analysis of Large Dimensional Random Matrices* (Springer, New York, 2010).
- [112] A. J. Bray and S. A. Roberts, Replica theory of the spherical model, *J. Phys. C* **13**, 5405 (1980).
- [113] D. Herr, E. Brown, B. Heim, M. Könz, G. Mazzola, and M. Troyer, Optimizing schedules for quantum annealing, [arXiv:1705.00420](https://arxiv.org/abs/1705.00420).
- [114] G. Quiroz, Robust quantum control for adiabatic quantum computation, *Phys. Rev. A* **99**, 062306 (2019).
- [115] W. Vinci, T. Albash, G. Paz-Silva, I. Hen, and D. A. Lidar, Quantum annealing correction with minor embedding, *Phys. Rev. A* **92**, 042310 (2015).
- [116] E. Grant, T. S. Humble, and B. Stump, Benchmarking quantum annealing controls with portfolio optimization, *Phys. Rev. Appl.* **15**, 014012 (2021).

- [117] J. Marshall, E. G. Rieffel, and I. Hen, Thermalization, freeze-out, and noise: Deciphering experimental quantum annealers, [Phys. Rev. Appl. **8**, 064025 \(2017\)](#).
- [118] J. Marshall, D. Venturelli, I. Hen, and E. G. Rieffel, Power of pausing: Advancing understanding of thermalization in experimental quantum annealers, [Phys. Rev. Appl. **11**, 044083 \(2019\)](#).
- [119] Y. Susa, Y. Yamashiro, M. Yamamoto, I. Hen, D. A. Lidar, and H. Nishimori, Quantum annealing of the p -spin model under inhomogeneous transverse field driving, [Phys. Rev. A **98**, 042326 \(2018\)](#).
- [120] J. I. Adame and P. L. McMahon, Inhomogeneous driving in quantum annealers can result in orders-of-magnitude improvements in performance, [Quantum Sci. Technol. **5**, 035011 \(2020\)](#).
- [121] 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 \(2025\)](#).
- [122] D. Venturelli and A. Kondratyev, Reverse quantum annealing approach to portfolio optimization problems, [Quantum Mach. Intell. **1**, 17 \(2019\)](#).
- [123] B. Yan and N. A. Sinitsyn, Analytical solution for nonadiabatic quantum annealing to arbitrary Ising spin Hamiltonian, [Nat. Commun. **13**, 2212 \(2022\)](#).