

Optimization and benchmarking of the thermal cycling algorithmAmin Barzegar ^{1,2}, Anuj Kankani ¹, Salvatore Mandrà ^{3,4} and Helmut G. Katzgraber ^{5,6,7,*}¹*Department of Physics and Astronomy, Texas A&M University, College Station, Texas 77843-4242, USA*²*Microsoft Quantum, Microsoft, Redmond, Washington 98052, USA*³*Quantum Artificial Intelligence Laboratory (QuAIL), NASA Ames Research Center, Moffett Field, California 94035, USA*⁴*KBR, Inc., Houston, Texas 77002, USA*⁵*Amazon Quantum Solutions Lab, Seattle, Washington 98170, USA*⁶*AWS Intelligent and Advanced Compute Technologies, Professional Services, Seattle, Washington 98170, USA*⁷*AWS Center for Quantum Computing, Pasadena, California 91125, USA*

(Received 13 January 2021; revised 9 August 2021; accepted 26 August 2021; published 8 September 2021)

Optimization plays a significant role in many areas of science and technology. Most of the industrial optimization problems have inordinately complex structures that render finding their global minima a daunting task. Therefore, designing heuristics that can efficiently solve such problems is of utmost importance. In this paper we benchmark and improve the thermal cycling algorithm [Phys. Rev. Lett. **79**, 4297 (1997)] that is designed to overcome energy barriers in nonconvex optimization problems by temperature cycling of a pool of candidate solutions. We perform a comprehensive parameter tuning of the algorithm and demonstrate that it competes closely with other state-of-the-art algorithms such as parallel tempering with isoenergetic cluster moves, while overwhelmingly outperforming more simplistic heuristics such as simulated annealing.

DOI: [10.1103/PhysRevE.104.035302](https://doi.org/10.1103/PhysRevE.104.035302)**I. INTRODUCTION**

Optimization is ubiquitous in science and industry. From the search for the ground state of exotic states of matter such as high-temperature superconductors in physics [1], topology optimization in material science [2], lead optimization in pharmaceutical drug discovery [3], spacecraft trajectory optimization [4], portfolio optimization in finance [5], scheduling in transportation [6], and speech recognition in artificial intelligence [7] to name a few. One important category of optimization problems is combinatorial optimization, which is the search for the minima of an objective function within a finite but often large set of solutions. Paradigmatic examples are the traveling salesman problem, the geometrical packing problem [8], graph coloring, the cutting stock problem, integer linear programming, etc. Many of these problems are NP-hard, in the sense that the worst-case time to find the optimum scales worse than a polynomial in the size of the input. Moreover, these problems are not only computationally hard to solve in the worst case but also in the typical case. These problems have a rough energy (cost function) landscape, consisting of numerous metastable states. Therefore, heuristics based on local search—e.g., the greedy algorithm [9]—tend to perform poorly on these types of problems as they can easily become trapped in local minima [10].

One way to circumvent this difficulty is to use a stochastic process such as Metropolis dynamics [11] to randomly access different parts of the phase space. An example of an algorithm that utilizes such random sampling is simulated annealing

(SA) [12]. Simulated annealing is a Markov-chain Monte Carlo process where a series of quasiequilibrium states are visited by following an “annealing schedule” during which the system is gradually cooled from a sufficiently high temperature to a target low temperature. The goal is to guide the stochastic process through an occasionally complex energy landscape toward the low-lying states. At high temperatures, the random “walker” can take long strides across phase space, thus allowing for the exploration of configurations far away in Hamming distance. As the system is cooled, the exploration domain of the walker is reduced according to the Gibbs distribution, and it eventually lands in a low-lying state. There is no guarantee that this is the true optimum, unless an (impractical) infinite annealing time is used [13]. Because SA is stochastic in nature, running many such processes in parallel can increase the chance of finding the true optimum. Nevertheless, without establishing a way for the phase-space information gathered by the random walkers to be shared, mere replication of a simulated annealing process will not yield any meaningful speedup. Multiple Markov-chain algorithms such as path-integral Monte Carlo [14–19], parallel tempering (PT) [20,21], and population annealing [22–26] take advantage of this “collective knowledge” to efficiently probe the solution space of a problem.

Closely related to the genetic local search approaches [27–30], the thermal cycling algorithm (TCA) [31,32] is another heuristic that integrates the power of parallel annealing processes with the utility of local search methods. The annealing part of this algorithm ensures that the phase space can be visited ergodically, whereas the local search part biases the dynamics toward the lower-energy states. When introduced in the early 2000s, thermal cycling was shown to outperform

*helmut@katzgraber.org

simulated annealing in solving some limited instances of the traveling salesman problem. Despite the early indications that TCA might be a useful tool in dealing with hard optimization problems, it has not been carefully benchmarked and hence widely adopted by the optimization community. Here we reintroduce the thermal cycling algorithm and outline the basic pseudocode. In addition, we conduct a comprehensive parameter optimization of TCA using synthetic planted problems, where we compare the performance of TCA to a number of modern optimizers, including simulated quantum annealing (SQA). In order to quantify the efficiency of the aforementioned heuristics, we study how their time to solution (TTS) [33,34] scales with the problem size. Our results show that when optimized properly, TCA can indeed be competitive with the state-of-the-art heuristics and therefore should be included in physics-inspired optimization platforms.

The paper is structured as follows. In Sec. II we explain the analysis techniques followed by the details of the algorithm in Sec. III. Section IV is dedicated to the benchmarking results of the study. Concluding remarks are presented in Sec. V.

II. DETAILS OF ANALYSIS

The cost function that we minimize in this benchmarking study is a 2-local Ising spin system, i.e.,

$$\mathcal{H} = \sum_i^N \sum_{j \in \mathcal{N}_i} J_{ij} s_i s_j + \sum_i^N h_i s_i, \quad (1)$$

where N is the total number of variables, \mathcal{N}_i is the adjacency list of the i th lattice site, J_{ij} is the coupling between spin s_i and s_j , and, finally, h_i is an external field applied to spin $s_i \in \{\pm 1\}$.

Most algorithms involve multiple parameters that need to be carefully tuned to observe the true asymptotic scaling. As such, a comprehensive hyperparameter optimization is in order. For benchmarking, we use synthetic problems whose ground state is unique and known beforehand. Here we use the deceptive cluster loop (DCL) problems [35] that are specifically designed for testing the performance of the D-Wave [36] quantum annealer against classical algorithms. DCL's are inspired by the original frustrated cluster loop (FCL) problems [37,38], which have a ferromagnetic planted ground state defined on a chimera graph [39]. The chimera topology consists of a two-dimensional lattice of fully connected bipartite $K_{4,4}$ cells in which all qubits are coupled ferromagnetically. The entire $K_{4,4}$ unit cell can, therefore, be viewed as one virtual variable. The cells are then connected via randomly chosen frustrated loops. The magnitude of intercell couplings are capped at a finite value R that adds local "ruggedness" to the problems [40,41]. The hardness of the FCL instances can be tuned by varying the density of the frustrated loops, often denoted by parameter α . In the DCL problems, the intercell couplers are multiplied by a scaling factor λ . Depending on the value of λ , the internal structure of the cells can be masked or accentuated, thus deceiving the annealers to spend more time optimizing the local structures rather than finding the global minimum.

As the measure of performance, we use the TTS [33,34] that is defined in the following way:

$$\text{TTS}(\alpha) = n(\alpha) \tau_{\text{run}}, \quad (2)$$

where $n(\alpha)$ is the number of times that the algorithm must be repeated, for a given parameter set α , to find the ground state at least once with a desired probability of p_d . τ_{run} is the average run time, conventionally measured in microseconds. If we assume that the *success probability*, i.e., the chance of hitting the ground state in a single run of the algorithm is $p_s(\alpha)$, then one can show from the binomial distribution that

$$p_d = \sum_{k \geq 1}^n \binom{n}{k} p_s^k (1 - p_s)^{n-k} = 1 - (1 - p_s)^n. \quad (3)$$

We may now use the above expression to find $n(\alpha)$ in Eq. (2):

$$n(\alpha) = \frac{\log[1 - p_d]}{\log[1 - p_s(\alpha)]}. \quad (4)$$

It is customary to set the desired probability in Eq. (4) to a high confidence value of $p_d = 0.99$. Because the TTS is a function of the algorithm parameters, a thorough optimization of the parameters must be performed to reliably compare heuristics based on it. Note that the optimization is often multidimensional, which makes the benchmarking a relatively laborious task. For each set of parameters α and each problem instance, we repeat the runs 100 times and calculate the success probability $p_s(\alpha)$ as the percentage of the ground-state hits. This process is repeated for all instances, in this case, 100, to calculate the median TTS, and the error bars are estimated using the bootstrap method. The above procedure is carried out for many other parameter-set values, and the optimal parameters are identified as the global minimum point of the TTS function. Having calculated the optimal TTS for all problem sizes ($N = 8L^2$), we can study the scaling behavior of the algorithm, which is often an exponential, i.e.,

$$\text{TTS}_{\text{opt}} \sim 10^{a+bL}. \quad (5)$$

The scaling exponent b determines the performance of an algorithm in the asymptotic limit, whereas a is a constant offset that depends on the factors nonintrinsic to the algorithm, such as hardware speed, code efficiency, etc. Therefore, a relatively unbiased way to compare different algorithms is to focus on the scaling exponent.

III. THERMAL CYCLING ALGORITHM

The thermal cycling algorithm works by periodically heating and cooling an ensemble of states while following a decreasing temperature schedule. The ensemble is prepared by selecting N_p lowest energy states among N_0 quenched random configurations. Starting from the initial inverse temperature $\beta_i = 0$, the above pool of states is annealed toward a final inverse temperature of β_f in N_T steps. At a given temperature, some energy is deposited into the ensemble states using N_s Metropolis updates (heating) followed by an immediate quench via a local search method (cooling). If any of the resulting states are lower in energy than the original set, then they are replaced in the pool. The heating-cooling cycle is repeated N_c times at a fixed temperature. In practice, the above process steers the ensemble toward the low-lying states while ensuring that metastable configurations do not hinder the dynamics. The temperature is then reduced, and the cycles

Algorithm 1 Thermal Cycling

```

1: Randomly initialize  $N_0$  configurations of the problem.
2: Quench each of the  $N_0$  states using a local search algorithm.
3: Construct a pool of states by selecting  $N_p$  states with the lowest
   energy from the above quenched states.
4: Build a list of lattice sites by comparing the spins on a given
   site between all pools states. If all aligned, add the site to the
   list.
5: for  $N_T$  steps starting from  $\beta = 0$  until  $\beta = \beta_f$  do
6:   for  $N_c$  cycles do
7:     Pick a random state from the pool.
8:     Add heat to the pool state using  $N_s$  Metropolis sweeps
       at  $\beta$ , excluding the spins in the site list.
9:     Quench the selected state.
10:    if lower energy is achieved then
11:      Replace the old state in the pool with the new one.
12:      Rebuild the site list by comparing the spins between
        the updated pool states.
13:    end if
14:  end for
15:  Increase  $\beta \rightarrow \beta + \Delta\beta$  in which the step size  $\Delta\beta$  is
    usually constant, i.e., linear schedule.
16: end for
17: Identify the pool state with the lowest energy as the solution of
    the problem.

```

start over. In Algorithm 1, we present a concise outline of the thermal cycling algorithm.

The main advantage of thermal cycling is the possibility of using a variety of variable-update classes in the quenching phase. By using more complex updates, exponentially many smaller local minima can be skipped in favor of lower-energy and configurationally more differing ones. This, however, does not necessarily translate to increased efficiency of the algorithm as the implementation overhead associated with those complex moves can negate the overall gain. Thus, there must be a trade-off between the complexity of the moves and the speedup owing to the reduced metastability. One of the simplest updates is a single-spin greedy move (SSGM) in which the most unstable spins (i.e., spins with the largest positive local fields) are flipped in a sequential fashion until no further improvement can be made. Another subset of the move classes are the double-spin random moves (DSRM), which consist of first attempting to flip a randomly chosen spin by itself, and if this is rejected, then trying to flip it together with one of the neighboring spins (looked up sequentially) that results in lowering the overall energy. The updates stop when the rejections accumulate to the total number of bonds in the problem. Another important type of move that we have studied here is the Lin-Kernighan cluster move (LKCM) that is based on the famous Lin-Kernighan algorithm [42,43] which is considered, to date, one of the most efficient heuristics for solving the traveling salesman problem (TSP). In a LKCM, a cluster of spins of size M is constructed starting from the most unstable spin and then appending the neighboring spins to it until the total cost of flipping the cluster becomes positive. The LKCM is essentially a k -opt local search algorithm [44–46], where k is determined from a sequence of partial costs, i.e., $\{c_1, c_2, \dots, c_M\}$. Here c_1 is the cost of flipping the

first spin, c_2 is the cost of flipping the first and the second spin together, and so on. One then flips the set of k spins with the lowest partial cost, c_k . We report the performance of each of the above move classes in the next section.

Another important point that one has to bear in mind is the duration of the heating phase. As mentioned earlier, the heating part of TCA ensures that the algorithm remains dynamic in spite of being quenched to often deep local minima. This requires that the system is subjected to a sufficient number of Metropolis updates. On the other hand, to preserve the gains of the previous cycles, the equilibration must be terminated in early stages. Otherwise, the system might end up in a configuration too far away in the phase space.

During the cycling process, the states in the pool are treated independently from one another. This has the potential pitfall that some of the states in the pool might wander off to energetically unfavorable parts of the configuration space. As mentioned earlier, this shortcoming can be alleviated by establishing an interaction between the pool states. One way to do this is to freeze the variables that are common among all of the states. We can justify this reduction by realizing that if the states in the pool have a feature in common, it is very likely that the feature will also appear in the ground-state configuration. Note that this step is closely related to metaheuristics such as tabu search [47,48] as well as self-avoiding random dynamics on integer complex systems (SARDONICS) [49], search for backbones [50–53] often used in genetic type algorithms, and sample persistence [54,55] which has been used in conjunction with algorithms such as simulated annealing as well as simulated quantum annealing.

In Fig. 1, we show the percentage of frozen spins versus the inverse temperature β for the system size $L = 12$ in which different colors represent various pool sizes N_p . It is interesting to observe that the freezing mechanism is only helpful when there is a moderate number of walkers. In other words, with too few walkers the entire ensemble collapses to one state very early on in the annealing schedule, whereas with too many walkers, the probability of all the pool states agreeing on the value of a particular spin becomes exceedingly low, hence rendering the freezing practically irrelevant. On the other hand, with population size around $N_p = 16$, the percentage saturates at an optimal value of roughly 60% such that a considerable number of degrees of freedom are preserved allowing for independent random walks to continue while still restraining the walkers from spreading to far from one another in the configurational space.

IV. RESULTS

In this study, we compare TCA to SA [12], SQA [16], PT [20,21], and parallel tempering with isoenergetic cluster moves (PT+ICM) [57]. SQA is the classical implementation of the quantum annealing process [58,59] in which the system is initialized in the ground state of a simple Hamiltonian and adiabatically [60] deformed into a target Hamiltonian whose ground state is difficult to find. Parallel tempering is a Monte Carlo algorithm that efficiently samples the equilibrium configurations of a system using the replica-exchange technique. The ICM update—which consist of rearranging a large collection of variables by inspecting the overlap between two




FIG. 1. The percentage of frozen spins in TCA for the system size $L=12$ versus inverse temperature β . The colors on the lines represent different pool sizes N_p . With only a few walkers, there is more chance of agreement between the spin values at a given site, hence the system collapses to a single walker early on during the annealing process, whereas with too many walkers, the freezing mechanism ceases to function.

replicas—is extremely effective for low connectivity graphs in which the cluster percolation threshold is small. For this benchmarking study we have generated 100 DCL instances for each linear size ranging from $L = 8$ to $L = 16$. The DCL parameters are fixed to a relatively hard regime of $\alpha = 0.24$, $R = 1$, and $\lambda = 3.0$ [35]. For each studied algorithm, we optimize the parameters via a grid search within its parameter space. All of the simulations are done on a single thread using Intel Xeon E5-2680 v4 2.40-GHz and Intel Xeon E5-2673 v4

2.30-GHz processors. In Table I, we have listed the optimal parameters of the thermal cycling algorithm as well as the other studied heuristics.

Note that since we are dealing with a high-dimensional optimization space, global optimality is neither guaranteed nor necessarily unique. We observe that most of the TCA parameters are robust with respect to the problem size and the number of annealing steps N_T —much like SA—is the only varying parameter. This is valuable information as it eliminates the necessity of a full parameter optimization in a practical implementation of the algorithm. In reality, the total effort in a TCA simulation is roughly proportional to $N_T N_p N_c N_s$ with some additional overhead caused by the local search. This suggests that correlation between the above parameters can be expected. For instance, similar performance can be achieved by increasing the number of walkers or by extending the annealing schedule while having a moderate pool size.

In Figs. 2 and 3 we show some examples of such parameter optimization. Figure 2 shows time to solution versus the total number of sweeps in PT+ICM for various problem sizes. The minimum of the curve marks the optimal sweep values. Figure 3(a) illustrates a two-dimensional cross section of the parameter space of TCA for the system size $L = 12$. The rest of the parameters are fixed to the values listed in Table I. Here, the color map represents the TTS values in a logarithmic scale. The axes show the number of cycles N_c and the number of annealing steps N_T . We observe two minima with comparable depth within error bars, corroborating the fact that N_T and N_c are anticorrelated. Figure 3(b) shows the TTS as a function of the pool size N_p . Beyond the horizontal dashed line marked by the hatched region, none of the benchmark problems can be solved in 100 independent attempts. It is interesting to observe that having a sufficiently large number of walkers is essential for the efficiency of the algorithm. Although it is intuitive that having more walkers will increase the chance of finding the ground state, the additional computational effort

TABLE I. Optimal parameters of the studied algorithms for different linear problem sizes L . Here N_T represents the total number of algorithmic steps in each heuristic. For example, in SQA, one algorithmic step involves building a Wolff cluster [56] in the imaginary time direction by sweeping randomly through all lattice sites, whereas in PT, an algorithmic step is defined as a Metropolis sweep over each replica followed by a tempering exchange move. β_i and β_f are the highest and lowest temperatures that the algorithm operates between. TCA involves the additional parameters of N_p , N_c , and N_s which are the pool size, the number of heating-cooling cycles per temperature, and the number of Metropolis sweeps, respectively. In SQA, N_p determines the number of Suzuki-Trotter slices while $\beta_c = \beta_q/N_p$ sets the “classical” inverse temperature with β_q being the temperature at which the underlying quantum annealing process is performed. Finally, for PT M is the number of temperatures which are spaced as a geometric sequence. When PT is accompanied with ICM updates, $2M$ replicas are used.

L	SA		SA+DSRM		TCA					SQA			PT			
	N_T	β_f	N_T	β_f	N_T	N_p	N_c	N_s	β_f	N_T	N_p	β_c	N_T	M	β_i	β_f
8	2896	0.25	3242	0.25	128	16	32	16	0.20	65536	32	1.0	2220	8	0.10	0.20
9	8192	0.25	7529	0.25	256	16	32	16	0.25	92681	32	1.0	3246	8	0.10	0.25
10	11585	0.30	11621	0.30	512	16	64	32	0.25	185363	32	1.0	5686	16	0.10	0.25
11	46340	0.30	26989	0.30	512	16	128	32	0.30	262144	32	1.0	10787	16	0.10	0.30
12	65536	0.30	41127	0.30	724	16	128	32	0.30	262144	32	1.0	17695	16	0.10	0.30
13	92680	0.30	62675	0.30	724	16	128	32	0.30	370727	32	1.0	21899	24	0.10	0.30
14	262144	0.30	95514	0.30	1024	16	128	32	0.30	524289	32	1.0	34954	24	0.10	0.30
15	1048576	0.30	145556	0.30	1024	16	128	32	0.30	741455	64	1.0	58909	32	0.10	0.30
16	1482910	0.30	221815	0.30	1448	16	128	32	0.30	1048576	64	1.0	96183	32	0.10	0.30




FIG. 2. Optimization of total sweeps for PT+ICM using various system sizes. The minimum TTS points correspond to the optimal sweep values.

of doing so negates any potential gains as seen from the flat regions in the TTS curves.

In Fig. 4 we show the scaling results of the thermal cycling algorithm with various move classes—that is, SSGM, DSRM, and LKCM, as explained in Sec. III. The main panel of Fig. 4 displays the optimal TTS values for different system sizes L , with the lines fitted to the largest five system sizes. The inset of the figure shows the scaling exponent b in Eq. (5) obtained from the linear fits. Note that the height of the boxes represents the error bars. It is clear that the DSRM are significantly more efficient than the other two types of move classes that we have studied here. It is interesting to note that the LKCM become less efficient as the problem size increases despite giving almost two orders of magnitude in constant speedup for smaller systems sizes. We speculate that this is due to the increased overhead of constructing a long sequence of partial costs—which requires building a large cluster of spins as we explained in Sec. III—relative to the optimal subcluster that is flipped in the end. It is worth mentioning here that the efficiency of the LKCM is to some extent topology-dependent because the updates are, in essence, cluster moves, and therefore might perform better when implemented on a different set of problems.

In Fig. 5(a), we show the scaling curves of various algorithms, corresponding to the optimal values of their parameters, as a function of the linear size L . As before, the TTS values are reported in the logarithmic (base 10) scale where the linear fits are interpreted as exponential scalings. There is a considerable constant offset associated with PT and PT+ICM, which is due to the use of a highly optimized implementation by Salvatore Mandrà as a part of NASA/TAMU UFO. Figure 5(b) shows the scaling exponent b [the slope of the linear fit in Fig. 5(a)]. We observe that TCA (with DSRM) scales overwhelmingly better than SA, in agreement with previous TSP studies. It is also more efficient than PT and even competitively close to SQA. Note that PT is already established as a powerful heuristic in many optimization-related applications. SQA and PT+ICM show the best





FIG. 3. (a) Parameter tuning of the thermal cycling algorithm using DCL problems of linear size $L=12$. The figure shows the number of cycles N_c versus the number of annealing steps N_T , with the rest of the parameters fixed to the values listed in Table I. The color map shows the TTS values. Two minima with similar depths are observed. The lowest minimum corresponds to the optimal parameters. (b) Optimization of the pool size N_p of TCA for a various system sizes. Above the horizontal dashed line, none of the problems can be solved in 100 attempts. Across all problem sizes studied $N_p = 16$ is sufficient.

performances among the studied solvers. This can be ascribed to the structure of the DCL problems that involve tall but thin barriers that can be easily tunneled through using SQA. Isoenergetic cluster moves are also well suited for the DCL problems as they cause large rearrangements of the variables, resulting in an efficient sampling of the configuration space.

We have also included SA+DSRM in which simulated annealing is augmented by the double-spin random moves used in TCA as explained earlier in Sec. III. We notice that unlike TCA, the improvement is negligible which highlights the fact that the efficiency of TCA does not solely arise from the use of complex quench moves but rather because of the




FIG. 4. Main panel: TTS of the thermal cycling algorithm versus the linear problem size L . Various quenching schemes consisting of SSGM, DSRM, and LKCM are displayed. The lines represent a linear fit in the logarithmic (base 10) scale to the five largest system sizes. We observe a sizable constant speedup with the LKCM although it diminishes at larger system sizes. Inset: Scaling exponent b corresponding to the slope of the linear fits in the main panel. The height of the boxes represents the error bars. The best scaling is obtained when TCA is used in conjunction with DSRM. These moves are simple enough to cause minimal overhead, yet complex enough to considerably reduce the number of the metastable states.

population-based nature of the algorithm. With a single random walker in SA, the zero-temperature quenches tend to strongly disturb the equilibrium distribution of states obtained by the finite-temperature Metropolis moves. On the contrary, with multiple walkers, the effect of such quenches is propagated slowly through the population allowing the downhill moves to be exploratory rather than disruptive. This can be best seen in Fig. 3(b) where TTS values increase dramatically once the pool size drops down to a handful of walkers.

V. CONCLUSION

In this paper we have thoroughly benchmarked the thermal cycling algorithm. Our results demonstrate that TCA is a competitive heuristic for solving problems with complex structures as it takes advantage of repeated heating and cooling to push the system toward the lower-energy states while ensuring that the system does not get trapped in an excited state. By reducing the variables among the TCA replicas, the stochastic process can be further accelerated, and the system can be guided more effectively toward the global minimum using the collective “memory” of the solution pool. Having carefully tuned the parameters, we show that TCA can be as effective as the state-of-the-art algorithms such as SQA, while overpowering SA and PT by great margins in the asymptotic scaling.

Due to the special structure of the DCL problems, which involve tall yet narrow barriers, SQA and PT+ICM outperform TCA because they utilize quantum effects and cluster updates to bypass those barriers. The true advantage of TCA

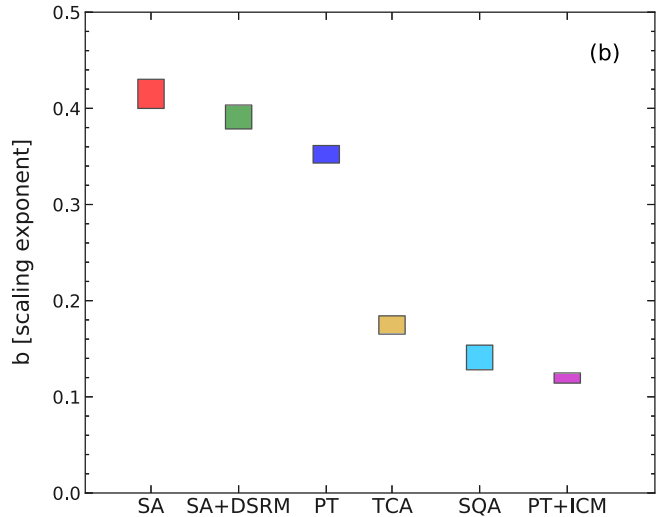



FIG. 5. Comparison between the scaling results of the studied algorithms using the DCL problems. (a) TTS versus the linear problem size L for various algorithms. Note that TTS is given on a logarithmic (base 10) scale. (b) Scaling exponent b in Eq. (5) for various algorithms. The height of the boxes represents the error bars. The TCA data points correspond to the best performing quenching scheme, i.e., DSRM, shown in Fig. 4. TCA scales better than SA and PT, and even comparable to SQA within the error bars. PT+ICM shows the best scaling. However, the latter is due to the great efficiency of the ICM updates in sampling the DCL phase space which is nonrepresentative for denser industrial problems. Unlike in TCA, the addition of DSRM quenches to SA (SA+DSRM) does not improve the performance.

might be revealed when using dense graphs with broad barriers, where PT+ICM and SQA would naturally struggle. TCA also lends itself to being integrated with ICM updates because it involves simultaneous annealing of many system replicas. This is in close analogy with to the iterative partial transcription algorithm [61,62]. It has been shown by Ochoa *et al.* [63] that a lower-energy state can be generated by overlapping two excited states via an ICM update. Therefore, one interesting addition to TCA could be trying to push the pool states further

down in energy by performing ICM updates at the end of the cycle, provided that the graph density is low enough and the clusters do not percolate.

ACKNOWLEDGMENTS

We thank A. Möbius for useful discussions regarding various aspects of the thermal cycling algorithm, and also providing code for the DSRM. The authors also acknowledge Jonathan Machta for critically reviewing the manuscript. H.G.K. thanks Dr. Pimple Popper for visualizations of energy landscapes. We thank Texas A&M University, NASA Ames Research Center, and Microsoft Quantum Group for providing access to computational resources. This work is supported in part by the Office of the Director of National Intelligence

(ODNI), Intelligence Advanced Research Projects Activity (IARPA), via MIT Lincoln Laboratory Air Force Contract No. FA8721-05-C-0002. S.M. also acknowledges the support from the Intelligence Advanced Research Projects Activity (IARPA) (IARPA IAA 1198) and the Defense Advanced Research Projects Agency (DARPA) (IAA 8839, Annex 125). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of ODNI, IARPA, DARPA or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The work of H.G.K. was performed before joining Amazon Web Services.

-
- [1] V. Stanev, C. Oses, A. G. Kusne, E. Rodriguez, S. Curtarolo, and I. Takeuchi, Machine learning modeling of superconducting critical temperature, *npj Comput. Mater.* **4**, 29 (2018).
- [2] M. P. Bendsøe, *Topology Optimization* (Springer US, Boston, MA, 2009), pp. 3928–3929.
- [3] G. M. Keserü and G. M. Makara, Hit discovery and hit-to-lead approaches, *Drug Discov. Today* **11**, 741 (2006).
- [4] A. Shirazi, J. Ceberio, and J. A. Lozano, Spacecraft trajectory optimization: A review of models, objectives, approaches and solutions, *Prog. Aeronaut. Sci.* **102**, 76 (2018).
- [5] J. Doering, A. A. Juan, R. Kizys, A. Fito, and L. Calvet, in *Modeling and Simulation in Engineering, Economics and Management*, edited by R. León, M. J. Muñoz-Torres, and J. M. Moneva (Springer International Publishing, Cham, 2016), pp. 22–30.
- [6] V. Guihaire and J.-K. Hao, Transit network design and scheduling: A global review, *Transport. Res. A: Policy Pract.* **42**, 1251 (2008).
- [7] T. Kinnunen and H. Li, An overview of text-independent speaker recognition: From features to supervectors, *Speech Commun.* **52**, 12 (2010).
- [8] R. M. Karp, *Complexity of Computer Computations* (Plenum, New York, 1972), p. 85.
- [9] T. H. Cormen, T. E. Leiserson, and R. L. Rivest, *Introduction to Algorithms* (MIT Press, Cambridge, MA, 1990).
- [10] G. Bendall and F. Margot, Greedy-type resistance of combinatorial problems, *Discr. Optimiz.* **3**, 288 (2006).
- [11] N. Metropolis and S. Ulam, The Monte Carlo method, *J. Am. Stat. Assoc.* **44**, 335 (1949).
- [12] S. Kirkpatrick, C. D. Gelatt, Jr., and M. P. Vecchi, Optimization by simulated annealing, *Science* **220**, 671 (1983).
- [13] S. Geman and D. Geman, *IEEE Trans. Pattern. Analy. Mach. Intell.* **PAMI-6**, 721 (1984).
- [14] H. F. Trotter, On the product of semi-groups of operators, *Proc. Am. Math. Soc.* **10**, 545 (1959).
- [15] M. Suzuki, Relationship among exactly soluble models of critical phenomena. I: 2D Ising model, dimer problem and the generalized XY-model, *Prog. Theor. Phys.* **46**, 1337 (1971).
- [16] M. Suzuki, Relationship between d-dimensional quantum spin systems and (d+1)-dimensional Ising systems: Equivalence, critical exponents and systematic approximants of the partition function and spin correlations, *Prog. Theor. Phys.* **56**, 1454 (1976).
- [17] M. Suzuki, *Quantum Monte Carlo Methods in Condensed Matter Physics* (World Scientific, Singapore, 1993).
- [18] D. P. Landau and K. Binder, *A Guide to Monte Carlo Simulations in Statistical Physics* (Cambridge University Press, Cambridge, UK, 2000).
- [19] M. Troyer, F. Alet, S. Trebst, and S. Wessel, in *AIP Conf. Proc. 690: The Monte Carlo Method in the Physical Sciences* (AIP, Melville, NY, 2003), pp. 156–169.
- [20] C. Geyer, in *Proceedings of the 23rd Symposium on the Interface*, edited by E. M. Keramidas (Interface Foundation, Fairfax Station, VA, 1991), p. 156.
- [21] K. Hukushima and K. Nemoto, Exchange Monte Carlo method and application to spin glass simulations, *J. Phys. Soc. Jpn.* **65**, 1604 (1996).
- [22] K. Hukushima and Y. Iba, in *The Monte Carlo Method in the Physical Sciences: Celebrating the 50th Anniversary of the Metropolis Algorithm*, Vol. 690, edited by J. E. Gubernatis (AIP, Los Alamos, NM, 2003), p. 200.
- [23] J. Machta, Population annealing with weighted averages: A Monte Carlo method for rough free-energy landscapes, *Phys. Rev. E* **82**, 026704 (2010).
- [24] W. Wang, J. Machta, and H. G. Katzgraber, Population annealing: Theory and application in spin glasses, *Phys. Rev. E* **92**, 063307 (2015).
- [25] C. Amey and J. Machta, Analysis and optimization of population annealing, *Phys. Rev. E* **97**, 033301 (2018).
- [26] A. Barzegar, C. Pattison, W. Wang, and H. G. Katzgraber, Optimization of population annealing Monte Carlo for large-scale spin-glass simulations, *Phys. Rev. E* **98**, 053308 (2018).
- [27] R. M. Brady, Optimization strategies gleaned from biological evolution, *Nature (Lond.)* **317**, 804 (1985).
- [28] H. Mühlenbein, M. Gorges-Schleuter, and O. Krämer, Evolution algorithms in combinatorial optimization, *Parallel Comput.* **7**, 65 (1988).
- [29] B. Freisleben and P. Merz, in *Proceedings of the 1996 IEEE International Conference on Evolutionary Computation*, edited by T. Fukuda and T. Furuhashi (IEEE, Los Alamitos, CA, 1996), pp. 616–621.
- [30] P. Merz and B. Freisleben, *Lecture Notes in Computer Science*, Vol. 1498 (Springer, Berlin, 1998), p. 765.

- [31] A. Möbius, A. Nekliudov, A. Díaz-Sánchez, K. H. Hoffmann, A. Fachat, and M. Schreiber, Optimization by Thermal Cycling, *Phys. Rev. Lett.* **79**, 4297 (1997).
- [32] A. Möbius, K. H. Hoffmann, and C. Schön, Optimization by thermal cycling, in *Proceedings of the 31st Workshop of the International School of Solid State Physics, Erice, Sicily, Italy, 20-26 July 2004* (World Scientific, Singapore, 2005), pp. 215–219.
- [33] S. Boixo, T. F. Rønnow, S. V. Isakov, Z. Wang, D. Wecker, D. A. Lidar, J. M. Martinis, and M. Troyer, Evidence for quantum annealing with more than one hundred qubits, *Nat. Phys.* **10**, 218 (2014).
- [34] 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).
- [35] S. Mandrà and H. G. Katzgraber, A deceptive step towards quantum speedup detection, *Quantum Sci. Technol.* **3**, 04LT01 (2018).
- [36] D-Wave Systems Inc., The D-Wave 2X Quantum Computer.
- [37] I. Hen, J. Job, T. Albash, T. F. Rønnow, M. Troyer, and D. A. Lidar, Probing for quantum speedup in spin-glass problems with planted solutions, *Phys. Rev. A* **92**, 042325 (2015).
- [38] J. King, S. Yarkoni, J. Raymond, I. Ozfidan, A. D. King, M. M. Nevisi, J. P. Hilton, and C. C. McGeoch, Quantum annealing amid local ruggedness and global frustration, *J. Phys. Soc. Jpn.* **88**, 061007 (2019).
- [39] P. Bunyk, E. Hoskinson, M. W. Johnson, E. Tolkacheva, F. Altomare, A. J. Berkley, R. Harris, J. P. Hilton, T. Lanting, and J. Whittaker, Architectural considerations in the design of a superconducting quantum annealing processor, *IEEE Trans. Appl. Supercond.* **24**, 1 (2014).
- [40] A. D. King, T. Lanting, and R. Harris, Performance of a quantum annealer on range-limited constraint satisfaction problems, (2015) [arXiv:1502.02098](https://arxiv.org/abs/1502.02098).
- [41] J. King, S. Yarkoni, J. Raymond, I. Ozfidan, A. D. King, M. M. Nevisi, J. P. Hilton, and C. C. McGeoch, Quantum annealing amid local ruggedness and global frustration, (2017), [arXiv:1701.04579](https://arxiv.org/abs/1701.04579) [quant-ph].
- [42] S. Lin and B. W. Kernighan, An effective heuristic algorithm for the traveling-salesman problem, *Operat. Res.* **21**, 498 (1973).
- [43] K. Helsgaun, An effective implementation of the linkernighan traveling salesman heuristic, *Eur. J. Operat. Res.* **126**, 106 (2000).
- [44] M. M. Flood, The traveling-salesman problem, *Operat. Res.* **4**, 61 (1956).
- [45] G. A. Croes, A method for solving traveling-salesman problems, *Operat. Res.* **6**, 791 (1958).
- [46] S. Lin, Computer solutions of the traveling salesman problem, *Bell Syst. Tech. J.* **44**, 2245 (1965).
- [47] F. Glover, Tabu searchpart i, *ORSA J. Comput.* **1**, 190 (1989).
- [48] F. Glover, Tabu searchpart ii, *ORSA J. Comput.* **2**, 4 (1990).
- [49] F. Hamze, Z. Wang, and N. de Freitas, Self-avoiding random dynamics on integer complex systems, (2011), [arXiv:1111.5379](https://arxiv.org/abs/1111.5379) [stat.CO].
- [50] J. Schneider, C. Froschhammer, I. Morgenstern, T. Husslein, and J. M. Singer, Searching for backbonesan efficient parallel algorithm for the traveling salesman problem, *Comput. Phys. Commun.* **96**, 173 (1996).
- [51] W. Zhang, Configuration landscape analysis and backbone guided local search: Part I: Satisfiability and maximum satisfiability, *Artif. Intell.* **158**, 1 (2004).
- [52] Y. Wang, Z. Lü, F. Glover, and J.-K. Hao, in *Evolutionary Computation in Combinatorial Optimization*, edited by P. Merz and J.-K. Hao (Springer, Berlin, 2011), pp. 72–83.
- [53] Y. Wang, Z. Lü, F. Glover, and J.-K. Hao, Backbone guided tabu search for solving the ubqp problem, *J. Heurist.* **19**, 679 (2013).
- [54] P. Chardaire, J. L. Lutton, and A. Sutter, Thermostatistical persistency: A powerful improving concept for simulated annealing algorithms, *Eur. J. Operat. Res.* **86**, 565 (1995).
- [55] H. Karimi, G. Rosenberg, and H. G. Katzgraber, Effective optimization using sample persistence: A case study on quantum annealers and various Monte Carlo optimization methods, *Phys. Rev. E* **96**, 043312 (2017).
- [56] U. Wolff, Collective Monte Carlo Updating for Spin Systems, *Phys. Rev. Lett.* **62**, 361 (1989).
- [57] Z. Zhu, A. J. Ochoa, and H. G. Katzgraber, Efficient Cluster Algorithm for Spin Glasses in Any Space Dimension, *Phys. Rev. Lett.* **115**, 077201 (2015).
- [58] A. B. Finnila, M. A. Gomez, C. Sebenik, C. Stenson, and J. D. Doll, Quantum annealing: A new method for minimizing multidimensional functions, *Chem. Phys. Lett.* **219**, 343 (1994).
- [59] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse Ising model, *Phys. Rev. E* **58**, 5355 (1998).
- [60] M. Born and V. Fock, Beweis des adiabatenatzes, *Z. Phys.* **51**, 165 (1928).
- [61] A. Möbius, B. Freisleben, P. Merz, and M. Schreiber, Combinatorial optimization by iterative partial transcription, *Phys. Rev. E* **59**, 4667 (1999).
- [62] A. Möbius, A. Diaz-Sanchez, B. Freisleben, M. Schreiber, A. Fachat, K. Hoffmann, P. Merz, and A. Nekliudov, Two physically motivated algorithms for combinatorial optimization: Thermal cycling and iterative partial transcription, *Comput. Phys. Commun.* **121-122**, 34 (1999).
- [63] A. J. Ochoa, D. C. Jacob, S. Mandrà, and H. G. Katzgraber, Feeding the multitude: A polynomial-time algorithm to improve sampling, *Phys. Rev. E* **99**, 043306 (2019).