Quantum Annealing Sampling with a Bias Field

Tobias Grass[1,2,3,*]

[1] *ICFO—Institut de Ciencies Fotoniques, The Barcelona Institute of Science and Technology, Av. Carl Friedrich Gauss 3, 08860 Castelldefels (Barcelona), Spain*

[2] *DIPC—Donostia International Physics Center, Paseo Manuel de Lardizábal 4, 20018 San Sebastián, Spain*

[3] *Ikerbasque—Basque Foundation for Science, Maria Diaz de Haro 3, 48013 Bilbao, Spain*

(Received 31 May 2022; revised 30 July 2022; accepted 13 September 2022; published 14 October 2022)

The presence of a bias field, encoding some information about a target state, can enhance the performance of quantum optimization methods. Here we investigate the effect of such a bias field on the outcome of quantum annealing sampling, with the example of the exact cover problem. The sampling is carried out on a D-Wave machine, and different bias configurations are benchmarked against the unbiased sampling procedure. It is found that the biased annealing algorithm works particularly well for larger problem sizes, where the Hamming distance between bias and target configuration becomes less important. This work motivates future research efforts for finding good bias configurations, either on the quantum machine itself or in a hybrid fashion via classical algorithms.

DOI: 10.1103/PhysRevApplied.18.044036

I. INTRODUCTION

Quantum annealing is a computational strategy for solving complex optimization problems on a programmable quantum device [1,2]. To this end, the optimization problem is formulated as a ground-state search problem, by expressing its cost function as a Hamiltonian H_P, typically an Ising model. Then, ground-state search is carried out by producing quantum fluctuations via a driver Hamiltonian $H_{\rm drive}$, which does not commute with H_P, and by reducing these fluctuations during an annealing procedure [3]. Two different shapes of quantum annealing can be distinguished. (i) Adiabatic quantum annealing (AQA), proposed in Refs. [4,5], demands that initially the system is prepared in the ground state of, typically, the pure driver Hamiltonian. By slowly ramping $H_{\rm drive}$ down and H_P up, the state is adiabatically transformed into the ground state of H_P. The adiabatic theorem is able to guarantee the success of this procedure, but to this end, it demands that the ramp speed is sufficiently slow (depending on the gap above the actual ground state). (ii) The other approach, quantum annealing sampling (QAS), typically performs the same actions as AQA, but at a ramp speed that invalidates the adiabatic theorem. Therefore, the annealer is not at all guaranteed to reach the ground state of H_P. However, through the weak coupling to a cold environment, the annealer is still expected to end up in a superposition of low-energy states. Then, a projective measurement at the end of the annealing process randomly selects one state out of this manifold, and repeated annealing runs sample over this manifold. The success of this strategy is heuristically based on the expectation that, due to the coupling to the cold bath, the ground state is still prominently represented in the final superposition.

Adiabatic state preparation, the key element of the AQA algorithm, is routinely performed in the context of quantum simulation [6]; see also Ref. [7] for an early realization of a small-scale quantum annealer and Ref. [8] reporting the early implementation of the Ising model in a trapped ion quantum simulator. For the more complicated case of glassy Ising couplings, which are relevant in the context of *NP*-hard optimization problems [9], the feasibility of adiabatic state preparation in trapped ion systems has been demonstrated theoretically [10–12]. However, several studies have argued that the scaling of the annealing gap with system size might be a prohibitive bottleneck to the AQA method [13–17]. Various strategies to overcome this bottleneck have been considered in the literature: inhomogeneous driver fields [18–26], nonstochastic driver Hamiltonians [27–31], tailored ramp protocols [32–39], reverse or biased annealing, and combinations thereof [40–49]. Another strategy to avoid the annealing bottleneck is to apply shortcuts to adiabaticity, such as counterdiabatic driving [50–58].

From the point of view of practical implementation, QAS has been more relevant than AQA. In particular, there exists a programmable quantum device that allows one to carry out QAS operations with thousands of qubits, produced and commercialized by the D-Wave company [59]. Unfortunately, QAS also suffers from exponentially

*tobias.grass@dipc.org

small annealing gaps, as they suppress the weight of the ground state in the final state. Accordingly, QAS also requires strategies that are designed to suppress the population of higher levels. In contrast to AQA, the QAS method may tolerate a certain level of noise in the system. In fact, in an early experiment with the D-Wave device [60], it has been demonstrated that thermal fluctuations at the start of the annealing can enhance the QAS success rate. Another strategy for performance enhancement of QAS is the nonadiabatic version of reverse annealing [61], which has recently been implemented experimentally [62]. Although this experiment demonstrates the success of the method, it also shows that the success rates are reduced by noise. Theoretically, it was shown that noise can be so harmful to reverse annealing that standard QAS produces better results [63].

In the present paper, we consider biased annealing as another QAS strategy. In the context of AQA, this strategy is proposed in Ref. [45]. It consists of adding a longitudinal bias field to the driver Hamiltonian, such that an infinitely slow switching of the Hamiltonian is still guaranteed to reach the ground state of the problem Hamiltonian. Theoretical simulations have shown that, for finite annealing times, the presence of the bias field allows for reaching the desired ground state on time scales where AQA without bias would fail, at least for small systems and for the choice of a good bias field. In Ref. [64], simulations have shown that a bias field also enhances the performance in the ultrafast limit where the annealing procedure is replaced by a sudden quench. Motivated by these theoretical studies, the present paper reports the experimental implementation of a biased QAS algorithm in a D-Wave machine. Hence, our study benchmarks the biasing algorithm under realistic conditions, which include also noisy incoherent processes.

Specifically, we perform biased QAS for a *NP*-hard optimization problem (exact cover), using the same small instances (up to $N = 14$ spins) as in the theoretical work from Ref. [45], and extending the system size to $N = 26$. While solving the classical optimization problem of this size is still no serious computational problem for classical computers, simulating the quantum annealing dynamics of such a large system would be extremely challenging on existing classical hardware. We use different bias fields that differ from the correct ground state by a Hamming distance of 0 to 13 spins, and in most cases, the bias field can improve the success rate considerably. Notably, the perspective of a performance enhancement becomes better as the system size grows. These results are in line with the results obtained for biased AQA reported in Ref. [45], as well as for biased quantum approximate optimization algorithm in Ref. [65]. This is a clear demonstration of the potential benefit due to a bias field, and it motivates future work on strategies of how to find suitable bias configurations.

The paper is organized as follows. In Sec. II, we describe the nature of the optimization problem (Sec. II A) and the biased sampling algorithm used to solve it (Sec. II B). In Sec. III, we present our results, and the most relevant information is captured in Fig. 3, showing the scaling of the success probabilities of the biased scheme with the problem size. In Sec. IV, we discuss the need for future research in order to develop feasible schemes for obtaining good bias configurations.

II. SETUP

A. Optimization problem

We study random instances of the exact cover problem. In this *NP*-hard problem, a manifold X of M elements is given, as well as N subsets S_i of this manifold. We focus on the special case, called EC3, in which every element of X appears in three subsets. For each element $x \in X$, these three subsets define a clause $C_x = \{i_1, i_2, i_3\}$. The task is to look for a union of subsets that exactly covers the manifold X. In a spin formulation of this problem (see, e.g., Ref. [66]), every subset is represented by a spin-1/2 variable, so every clause selects three spins from which one spin has to be singled out. Choosing σ_i^z as the computational basis, the cost function of each clause is given by $h_{C_x} = (\sigma^z_{C_x(1)} + \sigma^z_{C_x(2)} + \sigma^z_{C_x(3)} - 1)^2$, which becomes zero if one spin or subset is singled out by taking the value -1, while the other two spins or subsets take the value $+1$. Any other configuration is associated with a positive energy cost. An exact cover of X exists if the following Ising Hamiltonian has zero-energy eigenstates:

$$H_P = \sum_{x=1}^{M} h_{C_x} = \sum_{x=1}^{M} (\sigma^z_{C_x(1)} + \sigma^z_{C_x(2)} + \sigma^z_{C_x(3)} - 1)^2. \quad (1)$$

At this point, we keep the Hamiltonian dimensionless, but we specify units of energy below. An example for an EC3 problem with seven spins (subsets) and five clauses (dimension of X) is given in Fig. 1. From randomly generated problem instances, we select only those instances that have exactly one configuration fulfilling all clauses, as these instances are particularly hard to solve (cf. Ref. [5]). The list of used instances together with the code of the biased annealing on the D-Wave is provided in Ref. [67].

Given a set of clauses, the problem instance can be rephrased in terms of a longitudinal Ising Hamiltonian:

$$H_P = \sum_{i<j} J_{ij} \sigma_i^z \sigma_j^z + \sum_i h_i \sigma_i^z. \quad (2)$$

When the so-defined Ising problem is implemented on the D-Wave machine, the device's operating system Ocean automatically scales the parameters J_{ij} and h_i in such

FIG. 1. Example for exact cover problem (EC3). A manifold X (here spanned by five elements $1, 2, 3, 4, 5$) and subsets S_i of this manifold are given; for instance $S_1 = \{1, 2, 4\}$, $S_2 = \{1, 2, 5\}$, etc., as shown in the table. In EC3, every element of X appears in three subsets. A clause C_x is then given by the three subsets in which the element x appears; see last column of the table. In the spin Hamiltonian H_P, defined in Eq. (1), every spin corresponds to one subset, and if a zero-energy configuration exists, the subsets associated with a down spin form an exact cover of X.

a way that they remain within the accessible parameter range, and fully exploit the range. On the `D-Wave 2000Q` device used in this study, the parameter ranges are $J_{ij} \in [-1, 1]$ and $h_i \in [-2, 2]$. The corresponding frequency units are specified by the annealing schedule (see below). To encode the problem Hamiltonian on the chimera geometry of the `D-Wave 2000Q` device, a minor embedding procedure is carried out [68,69], using the function `EmbeddingComposite()`, provided by the `Ocean` software.

B. Biased sampling algorithm

For the solution of the optimization problem via QAS, we define a homogeneous driver Hamiltonian $H_{\text{drive}} = \sum_i \sigma_i^x$, which does not commute with H_P, as well as a longitudinal bias Hamiltonian $H_{\text{bias}} = -\sum_i \mu_i \sigma_i^z$. This bias term energetically favors the spin configuration where $\langle \sigma_i^z \rangle = \text{sign}(\mu_i)$ for all i, and, therefore, if $\text{sign}(\mu_i)$ agrees with the solution of H_P in all or many values of i, the presence of such a bias Hamiltonian enhances the annealer's success rate. In the following, we concentrate on the case where $\mu_i = \pm 1$. However, we note that another interesting scenario might include also the case $\mu_i = 0$. This could be relevant in situations where the estimate of the target state can only make a prediction over a subset of spins, as recently discussed in Ref. [70].

In the AQA version of biased annealing, as described in Ref. [45], the bias field is included in the driver Hamiltonian, i.e., it is switched off during the annealing. The implementation of biased QAS, presented here, differs from the earlier work, as the bias is switched on simultaneously with the problem Hamiltonian. Accordingly, the full annealing Hamiltonian reads

$$H(t) = A(t)H_{\text{drive}} + B(t)[H_P + H_{\text{bias}}]. \tag{3}$$

In the beginning (at $t = 0$), $A(0) \gg B(0)$, whereas at the end of the annealing, at time $t = \tau$, A has decreased to

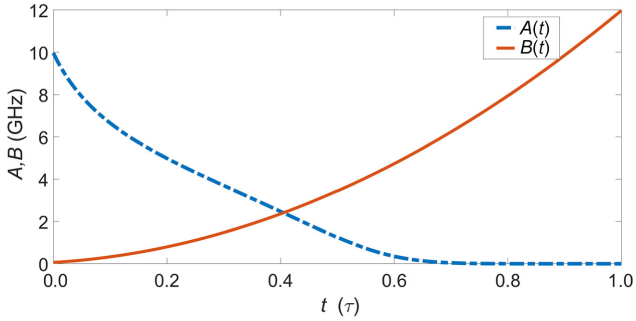

FIG. 2. Annealing schedule used for the sampling.

$A(\tau) \to 0$ and B has grown to $B(\tau) \approx A(0)$. The explicit schedules $A(t)$ and $B(t)$ are shown in Fig. 2. For the sampling, we choose different annealing times τ in the range between 1 μs (the fastest possible choice) and 15 μs.

From the perspective of AQA, it is an advantage to include the bias field in the driver, because with this choice, the adiabatic theorem still can guarantee that the annealer ends up in the ground state of H_P, even for a bad choice of μ_i. On the other hand, for the QAS method, the adiabatic theorem is irrelevant. Including the bias field into the problem part of the Hamiltonian is, first and foremost, motivated by practical considerations. Although the `D-Wave` allows control of the time dependence of the coupling terms (i.e., the terms of type $J_{ij}\sigma_i^z\sigma_j^z$) and the longitudinal fields (i.e., the terms of type $h_i\sigma_i^z$) independently from each other through the function `h_gain_schedule`, we cannot independently control the problem and the bias Hamiltonians, because both parts contribute to the longitudinal terms. However, from the point of view of QAS, the presence of even a bad bias term at the end of the sampling dynamics seems relatively harmless. At that stage of the annealing schedule, the quantum fluctuations are too weak to produce significant changes anymore, and the system thermalizes with respect to H_P (in the standard case) or with respect to $(H_{\text{bias}} + H_P)$ in the case with bias [71]. In both Hamiltonians, the target state is expected to have a relatively large thermal weight, if the Hamming distance between bias and target is small. In this case, generic low-energy excitations of H_P (with an energy of 4 for every violated clause) are likely to have a larger Hamming distance from the bias than the target state; hence the thermal weight of the target state in the Boltzmann distribution with respect to $(H_{\text{bias}} + H_P)$ is even enhanced as compared with the weight in the Boltzmann distribution with respect to H_P.

III. RESULTS

We perform QAS for the exact cover problem, as specified in the previous section, for problem instances denoted (N, α), where N defines the number of spins or qubits and α is an index for each of the random instances. We

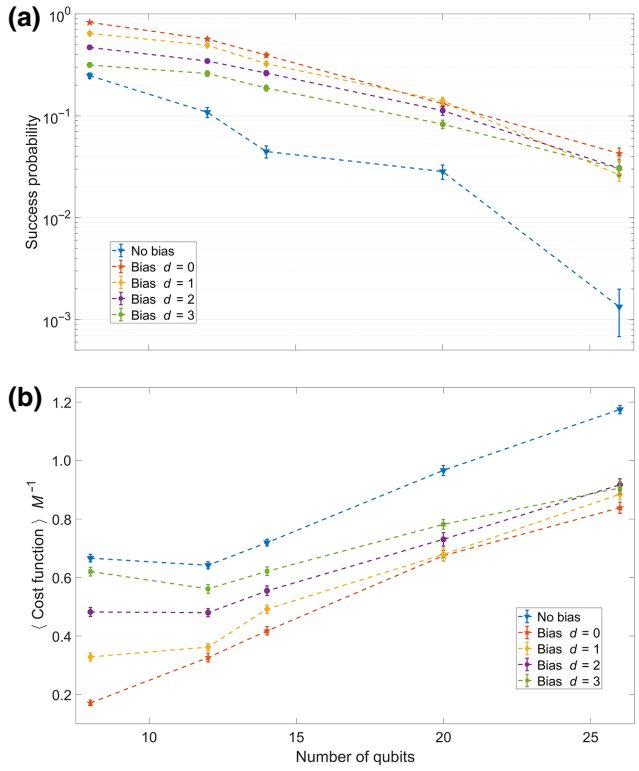

FIG. 3. The performance of biased QAS is benchmarked and compared with unbiased QAS, by considering (a) success probability and (b) average cost function value at the end of the anneal (divided by the number M of clauses) as a function of system size, and for different biases, distinguished by their Hamming distance d to the optimal solution. Here, the annealing time $\tau = 1\ \mu$s is fixed. The error bars represent the standard error of the mean.

consider sizes between $N = 8$ and $N = 26$, and for each size, we consider $N_{\text{in}} = 100$ random instances, that is, $\alpha \in [1, 100]$. From the solutions of the problem Hamiltonian, which in our case are known *a priori* and that we denote by $\{s_i\}$, we construct bias fields $\{\mu_i\}$, such that the Hamming distance $d = (1/2)\sum_i(|\mu_i - s_i|)$ between bias and target is fixed. For every instance and for each value of d, we perform $N_{\text{anneal}} = 30$ anneals, and we count the number of times N_{success} in which the final measurement outcome agrees with the target configuration. Then, the corresponding success probability $p_{(N,\alpha)}(d) = N_{\text{success}}/N_{\text{anneal}}$ is averaged over all instances of a given size, $p_N(d) = \sum_\alpha p_{(N,\alpha)}(d)/N_{\text{in}}$, and the result is plotted in Fig. 3(a), with error bars representing the standard error of the mean. We also determine, after every anneal, the cost function corresponding to the obtained configuration. The average over all anneals and all instances, as a function of system size N and for different Hamming distance d of the bias, is shown in Fig. 3(b). For the annealing time τ in Fig. 3, we use the fastest possible choice, $\tau = 1\ \mu$s. For comparison, we also perform QAS without bias field.

As expected, the obtained data provide evidence that any of the chosen bias fields enhances the performance of the annealer. Importantly, the enhancement becomes more significant as the system size is increased. To some extent, this behavior is also expected from the study of biased AQA [45], and can be attributed to the fact that, at a fixed Hamming distance d, the error ratio of the bias, $r \equiv d/(N-d)$, that is, the ratio of "bad" to "good" bias terms, decreases with N. However, let us compare, for instance, the results at $N = 8$ and $d = 1$ with the results at $N = 26$ and $d = 3$. Although in these two configurations the bias fields have a similar error ratio r, the enhancement turns out to be quite different. To make the comparison quantitative, we define the enhancement factor $\gamma_N(d)$ for a given problem size N and a given bias distance d as the success probability with bias, $p_N(d)$, divided by success probability without bias, p_N, that is, $\gamma_N(d) = p_N(d)/p_N$. At $N = 8$ and $d = 1$, the enhancement is $\gamma_8(1) = 2.6 \pm 0.2$, whereas at $N = 26$ and $d = 3$, we obtain $\gamma_{26}(3) = 23 \pm 16$. Despite the large statistical uncertainty in the latter number, owing to the very low success rate of the unbiased annealing at $N = 26$, these numbers suggest that the biased QAS method is particularly strong for larger problem instances. To further appreciate this strength of the biased method, it is also illustrative to look at the number of instances that, after 30 anneals, are solved correctly. At $N = 26$ and $\tau = 1\ \mu$s, there are only 4 (out of 100) instances that are solved without bias. This shows that unbiased QAS is essentially not able to find the solution for this problem size at this short annealing time. On the other hand, even a $d = 3$ bias elevates the percentage of correctly solved instances to almost 50% (47 out of 100 instances). Another interesting observation that can be made for large system sizes is the fact that, within the studied range of d, the Hamming distance d of the bias affects the annealing outcome only weakly. While at small N both success probability and average cost function appear ordered according to d, for $N = 26$ the bias with $d = 1$ achieves a slightly lower success probability than for biases with $d = 2$ and $d = 3$. As seen from the error bars in the plot, this "under"-performance can be attributed to statistical fluctuations, and in fact at large problem sizes, all the different bias choices perform similarly well.

An important practical question is the amount of errors that a bias field may contain before it starts to negatively affect the performance. To systematically investigate this question we analyze the performance, again in terms of success probability and average cost function, as a function of the Hamming distance of the bias (see Fig. 4) for three different system sizes ($N = 14, 20, 26$). Improvements of both success probability and cost function are obtained for bias fields with Hamming distances up to $d = 4, 7, 12$ for the three different system sizes, respectively. Hence, the data demonstrate that larger systems can accommodate a larger *relative* amount of errors in the bias field, which

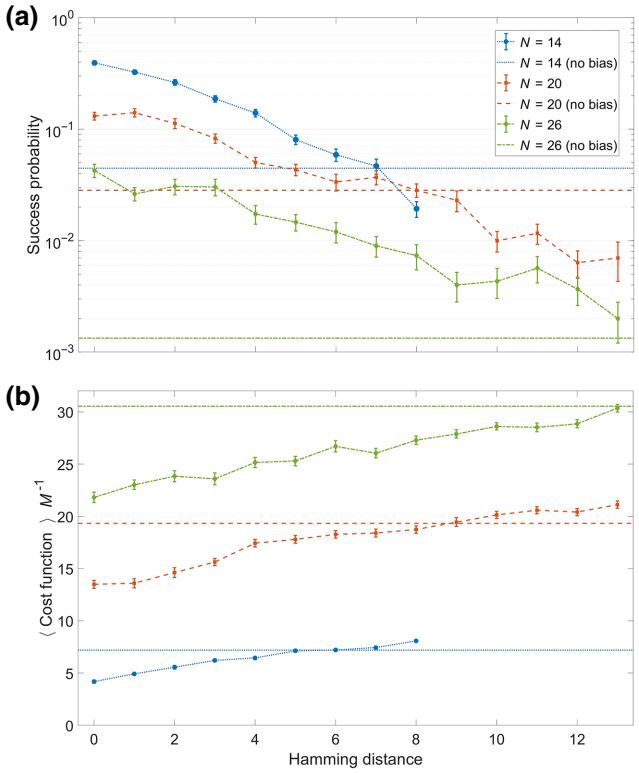

FIG. 4. The performance of biased QAS is benchmarked by considering (a) success probability and (b) average cost function value as a function of Hamming distance of the bias field, for different system sizes. The horizontal lines correspond to the values obtained without any bias field. Improvements for both success probability and cost function values are obtained for biases with maximal Hamming distances $d = 4, 7, 12$ for system sizes of $N = 14, 20, 26$ spins. This demonstrates that the larger systems can profit from worse bias fields (in terms of both absolute Hamming distance d as well as relative Hamming distance d/N.

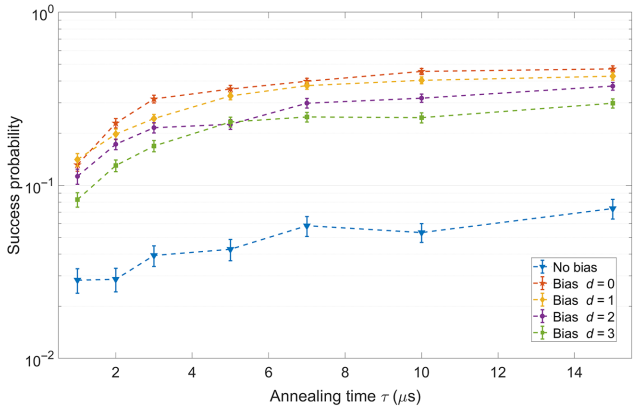

FIG. 5. We investigate the dependence of biased and unbiased annealing on the annealing time τ. We focus on $N = 20$ qubits, and plot the averaged success rates for unbiased QAS and biased QAS with different choices of the bias Hamming distance d from the correct solution. The error bars represent the standard error of the mean.

makes the biasing particularly interesting from the point of view of scalability.

In Fig. 5, we study the dependence of the success probability on the annealing time τ for $N = 20$. Increasing τ increases the success probability, both with and without bias fields, in particular for short time scales. This improvement for small τ is attributed to the fact that the unitary part of the dynamics gets closer to the adiabatic evolution. For longer time scales, the success probability improves only subtly with τ, and in this regime, the improvement is attributed to longer thermalization times. Importantly, for bias fields close to the target configuration, the ground state of H_P remains prominently present also in the Boltzmann distribution with respect to $(H_P + H_{\text{bias}})$. Hence, the presence of the bias fields does not harm the outcome of the thermalization phase. Accordingly, the improvement is seen to happen at similar rates for the different bias configurations and in the unbiased case. For instance, the success probability for the $d = 3$ bias increases by a factor 3.6 ± 0.5, when increasing τ from 1 μs to 15 μs, and by factor 2.6 ± 0.8 for unbiased QAS. From this, we conclude that the choice of τ does not seem to be crucial for benchmarking biased against unbiased QAS via average success probabilities. A more differentiated picture can be obtained when analyzing the success probability of individual instances. In Fig. 6, we present a comparison between the success probabilities of biased and unbiased annealing for individual instances. Again, we concentrate on $N = 20$, and show the data for $\tau = 1$ μs as well as for $\tau = 15$ μs, considering bias fields with $d = 0$ and $d = 3$. It is obvious that in all cases the majority of instances profit from the presence of a bias field, but the number of instances that do not take advantage from the bias is significantly reduced by increasing τ. For instance, at $\tau = 1$ μs, only 78 (65) instances (out of 100) have higher success probabilities in biased QAS with $d = 0$ ($d = 3$) than in unbiased QAS. For $\tau = 15$ μs, however, this is the case for 98 (89) instances. From this perspective, the advantage of a bias field appears to be increased if the annealing is performed at lower speed.

We also note that, if the annealing time τ was the only time scale in the annealing protocol, increasing τ would, for all N, increase the average time to solution, because the best ratios between average success probability over τ are achieved at $\tau = 1$ μs. However, since most of the computational time is actually spent by the programming of the annealer rather than by the annealing procedure itself, the best choice of τ is not obvious from these numbers alone, and it can actually be beneficial to choose a larger value of τ.

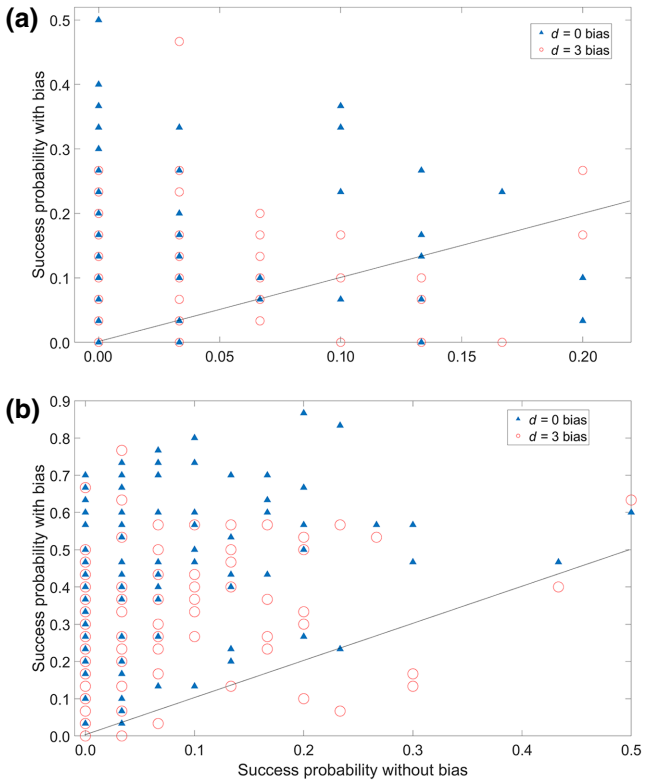

FIG. 6. We plot the success probabilities of biased QAS (with $d = 0$ and $d = 3$) versus the success probabilities of unbiased QAS, for individual instances with $N = 20$ qubits. (a) Annealing time $\tau = 1$ μs; (b) annealing time $\tau = 15$ μs. We note that, owing to the fact that the success probabilities are calculated from $N_{\text{anneal}} = 30$ annealing runs, there is a limited discrete set of possible values, and hence it occurs that several (of the total of 100) instances have identical success probabilities, that is, the corresponding data points cannot be discerned from each other.

IV. DISCUSSION AND OUTLOOK

For producing the improvements via bias fields that have been demonstrated in the previous section, strategies to choose good bias fields are needed. This choice requires some knowledge about the target state, which here we have assumed to be given *a priori*. In practice, though, this might not be the case. A strategy to find good bias configurations is proposed in Ref. [45]: the outcome of a previous annealing run can be used as a bias in the subsequent anneal, which can iteratively improve the bias, until the correct solution is found. One limitation of this strategy, discussed in the context of theoretical simulations of the algorithm in Ref. [45], is the fact that in certain cases the initially produced bias might actually attract the annealer towards an excited state rather than the ground state. In a practical implementation, however, the iterative scheme suffers from another, maybe more significant, drawback. As already mentioned above, the time τ (or even $N_{\text{anneal}}\tau$) is typically smaller than the time that is required for the programming of the `D-Wave` device. Since in the iterative procedure the annealer has to be reprogrammed after every anneal (or after some sampling period), the time efficiency of iterative QAS becomes very low. Possibly, in future hardware implementations, the fact that the iterative updates occur only in the magnetic field terms but not in the couplings might be exploited to speed up iterative schemes.

While the iterative scheme is appealing because it can be carried out entirely on a quantum device, future research should also explore the possibility of determining good bias fields through classical algorithms, e.g., through variations of the simulated annealing algorithm [72]. A final assessment as to whether biased annealing has a real advantage over unbiased annealing can only be made if the computational cost of finding a suitable bias is taken into account, but the improvements due to bias fields, seen in the present work, appear to be very promising.

ACKNOWLEDGMENTS

The project that gave rise to these results received the support of a fellowship from "La Caixa" Foundation (ID100010434). The fellowship code is LCF/BQ/PI19/11690013. I also acknowledge support from ERC AdG NOQIA; Ministerio de Ciencia y Innovation Agencia Estatal de Investigaciones (PGC2018-097027-B-I00/10.13039/501100011033, CEX2019-000910-S/10.13039/501100011033, Plan National FIDEUA PID2019-106901GB-I00, FPI, QUANTERA MAQS PCI2019-111828-2, QUANTERA DYNAMITE PCI2022-132919, Proyectos de I+D+I "Retos Colaboración" QUSPIN RTC2019-007196-7); European Union NextGenerationEU (PRTR); Fundació Cellex; Fundació Mir-Puig; Generalitat de Catalunya (European Social Fund FEDER) and CERCA program (AGAUR Grant No. 2017 SGR 134, QuantumCAT U16-011424, cofunded by ERDF Operational Program of Catalonia 2014-2020); Barcelona Supercomputing Center MareNostrum (FI-2022-1-0042); EU Horizon 2020 FET-OPEN OPTOlogic (Grant No. 899794); National Science Centre, Poland (Symfonia Grant No. 2016/20/W/ST4/00314); European Union's Horizon 2020 research and innovation program under Marie Sklodowska-Curie Grants No. 101029393 (STREDCH) and No. 847648 ("La Caixa" Junior Leaders fellowships ID100010434: LCF/BQ/PI20/11760031, LCF/BQ/PR20/11770012, LCF/BQ/PR21/11840013).

[1] T. Albash and D. A. Lidar, Adiabatic quantum computation, Rev. Mod. Phys. **90**, 015002 (2018).

[2] P. Hauke, H. G. Katzgraber, W. Lechner, H. Nishimori, and W. D. Oliver, Perspectives of quantum annealing: Methods and implementations, Rep. Prog. Phys. **83**, 054401 (2020).

[3] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse Ising model, Phys. Rev. E **58**, 5355 (1998).

[4] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, Quantum computation by adiabatic evolution, arXiv arXiv:10.48550/arxiv.quant-ph/0001106 (2000).

[5] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem, Science **292**, 472 (2001).

[6] A. Venegas-Gomez, J. Schachenmayer, A. S. Buyskikh, W. Ketterle, M. L. Chiofalo, and A. J. Daley, Adiabatic preparation of entangled, magnetically ordered states with cold bosons in optical lattices, Quantum Sci. Technol. **5**, 045013 (2020).

[7] M. Steffen, W. van Dam, T. Hogg, G. Breyta, and I. Chuang, Experimental Implementation of an Adiabatic Quantum Optimization Algorithm, Phys. Rev. Lett. **90**, 067903 (2003).

[8] K. Kim, M.-S. Chang, S. Korenblit, R. Islam, E. E. Edwards, J. K. Freericks, G.-D. Lin, L.-M. Duan, and C. Monroe, Quantum simulation of frustrated Ising spins with trapped ions, Nature **465**, 590 (2010).

[9] F. Barahona, On the computational complexity of Ising spin glass models, J. Phys. A Math. Theor. **15**, 3241 (1982).

[10] P. Hauke, L. Bonnes, M. Heyl, and W. Lechner, Probing entanglement in adiabatic quantum optimization with trapped ions, Front. Phys. **3**, 21 (2015).

[11] T. Graß, D. Raventós, B. Juliá-Díaz, C. Gogolin, and M. Lewenstein, Quantum annealing for the number-partitioning problem using a tunable spin glass of ions, Nat. Commun. **7**, 11524 (2016).

[12] D. Raventós, T. Graß, B. Juliá-Díaz, and M. Lewenstein, Semiclassical approach to finite-temperature quantum annealing with trapped ions, Phys. Rev. A **97**, 052310 (2018).

[13] T. Jörg, F. Krzakala, J. Kurchan, and A. C. Maggs, Simple Glass Models and Their Quantum Annealing, Phys. Rev. Lett. **101**, 147204 (2008).

[14] A. P. Young, S. Knysh, and V. N. Smelyanskiy, First-Order Phase Transition in the Quantum Adiabatic Algorithm, Phys. Rev. Lett. **104**, 020502 (2010).

[15] T. Jörg, F. Krzakala, G. Semerjian, and F. Zamponi, First-Order Transitions and the Performance of Quantum Algorithms in Random Optimization Problems, Phys. Rev. Lett. **104**, 207206 (2010).

[16] B. Altshuler, H. Krovi, and J. Roland, Anderson localization makes adiabatic quantum optimization fail, Proc. Nat. Acad. Sci. USA **107**, 12446 (2010).

[17] S. Knysh, Zero-temperature quantum annealing bottlenecks in the spin-glass phase, Nat. Commun. **7**, 12370 (2016).

[18] N. G. Dickson and M. H. S. Amin, Does Adiabatic Quantum Optimization Fail for NP-Complete Problems?, Phys. Rev. Lett. **106**, 050502 (2011).

[19] E. Farhi, J. Goldstone, D. Gosset, S. Gutmann, H. B. Meyer, and P. Shor, Quantum adiabatic algorithms, small gaps, and different paths, Quantum Inf. Comput. **11**, 181 (2011).

[20] N. G. Dickson and M. H. Amin, Algorithmic approach to adiabatic quantum optimization, Phys. Rev. A **85**, 032303 (2012).

[21] M. M. Rams, M. Mohseni, and A. del Campo, Inhomogeneous quasi-adiabatic driving of quantum critical dynamics in weakly disordered spin chains, New J. Phys. **18**, 123034 (2016).

[22] T. Lanting, A. D. King, B. Evert, and E. Hoskinson, Experimental demonstration of perturbative anticrossing mitigation using nonuniform driver Hamiltonians, Phys. Rev. A **96**, 042322 (2017).

[23] Y. Susa, Y. Yamashiro, M. Yamamoto, and H. Nishimori, Exponential speedup of quantum annealing by inhomogeneous driving of the transverse field, J. Phys. Soc. Jpn **87**, 023002 (2018).

[24] 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).

[25] A. Hartmann and W. Lechner, Quantum phase transition with inhomogeneous driving in the Lechner-Hauke-Zoller model, Phys. Rev. A **100**, 032110 (2019).

[26] 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).

[27] Y. Seki and H. Nishimori, Quantum annealing with antiferromagnetic fluctuations, Phys. Rev. E **85**, 051112 (2012).

[28] E. Crosson, E. Farhi, C. Y.-Y. Lin, H.-H. Lin, and P. Shor, Different strategies for optimization using the quantum adiabatic algorithm, arXiv arXiv:10.48550/arXiv.1401.7320 (2014).

[29] 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).

[30] T. Albash, Role of nonstoquastic catalysts in quantum adiabatic optimization, Phys. Rev. A **99**, 042334 (2019).

[31] I. Ozfidan, C. Deng, A. Smirnov, T. Lanting, R. Harris, L. Swenson, J. Whittaker, F. Altomare, M. Babcock, C. Baron *et al.*, Demonstration of a Nonstoquastic Hamiltonian in Coupled Superconducting Flux Qubits, Phys. Rev. Appl. **13**, 034037 (2020).

[32] T. Kiwaki, Variational optimization of annealing schedules, arXiv arXiv:10.48550/arxiv.1502.05313 (2015).

[33] L. Zeng, J. Zhang, and M. Sarovar, Schedule path optimization for adiabatic quantum computing and optimization, J. Phys. A: Math. Theor. **49**, 165305 (2016).

[34] H. Hu and B. Wu, Optimizing the quantum adiabatic algorithm, Phys. Rev. A **93**, 012345 (2016).

[35] J. Lin, Z. Y. Lai, and X. Li, Quantum adiabatic algorithm design using reinforcement learning, Phys. Rev. A **101**, 052327 (2020).

[36] Y. Susa and H. Nishimori, Variational optimization of the quantum annealing schedule for the Lechner-Hauke-Zoller scheme, Phys. Rev. A **103**, 022619 (2021).

[37] P. R. Hegde, G. Passarelli, A. Scocco, and P. Lucignano, Genetic optimization of quantum annealing, Phys. Rev. A **105**, 012612 (2022).

[38] J. Lin, Z. Zhang, J. Zhang, and X. Li, Hard-instance learning for quantum adiabatic prime factorization, Phys. Rev. A **105**, 062455 (2022).

[39] B. Irsigler and T. Grass, The quantum annealing gap and quench dynamics in the exact cover problem, Quantum **6**, 624 (2022).

[40] A. Perdomo-Ortiz, S. E. Venegas-Andraca, and A. Aspuru-Guzik, A study of heuristic guesses for adiabatic quantum computation, Quantum Inf. Process. **10**, 33 (2011).

[41] Q.-H. Duan, S. Zhang, W. Wu, and P.-X. Chen, An alternative approach to construct the initial Hamiltonian of the adiabatic quantum computation, Chin. Phys. Lett. **30**, 010302 (2013).

[42] A. Ramezanpour, Optimization by a quantum reinforcement algorithm, Phys. Rev. A **96**, 052307 (2017).

[43] M. Ohkuwa, H. Nishimori, and D. A. Lidar, Reverse annealing for the fully connected p-spin model, Phys. Rev. A **98**, 022314 (2018).

[44] C. L. Baldwin and C. R. Laumann, Quantum algorithm for energy matching in hard optimization problems, Phys. Rev. B **97**, 224201 (2018).

[45] T. Graß, Quantum Annealing with Longitudinal Bias Fields, Phys. Rev. Lett. **123**, 120501 (2019).

[46] Y. Yamashiro, M. Ohkuwa, H. Nishimori, and D. A. Lidar, Dynamics of reverse annealing for the fully connected p-spin model, Phys. Rev. A **100**, 052321 (2019).

[47] G. Passarelli, K.-W. Yip, D. A. Lidar, H. Nishimori, and P. Lucignano, Reverse quantum annealing of the p-spin model with relaxation, Phys. Rev. A **101**, 022331 (2020).

[48] A. Ramezanpour, Quantum walk in a reinforced free-energy landscape: Quantum annealing with reinforcement, Phys. Rev. A **106**, 012418 (2022).

[49] Z. Tang and E. Kapit, Unconventional quantum annealing methods for difficult trial problems, Phys. Rev. A **103**, 032612 (2021).

[50] D. Guéry-Odelin, A. Ruschhaupt, A. Kiely, E. Torrontegui, S. Martínez-Garaot, and J. G. Muga, Shortcuts to adiabaticity: Concepts, methods, and applications, Rev. Mod. Phys. **91**, 045001 (2019).

[51] A. del Campo, Shortcuts to Adiabaticity by Counterdiabatic Driving, Phys. Rev. Lett. **111**, 100502 (2013).

[52] H. Saberi, T. c. v. Opatrný, K. Mølmer, and A. del Campo, Adiabatic tracking of quantum many-body dynamics, Phys. Rev. A **90**, 060301 (2014).

[53] G. Passarelli, V. Cataudella, R. Fazio, and P. Lucignano, Counterdiabatic driving in the quantum annealing of the p-spin model: A variational approach, Phys. Rev. Res. **2**, 013283 (2020).

[54] K. Funo, N. Lambert, and F. Nori, General Bound on the Performance of Counter-diabatic Driving Acting on Dissipative Spin Systems, Phys. Rev. Lett. **127**, 150401 (2021).

[55] L. Prielinger, A. Hartmann, Y. Yamashiro, K. Nishimura, W. Lechner, and H. Nishimori, Two-parameter counterdiabatic driving in quantum annealing, Phys. Rev. Res. **3**, 013227 (2021).

[56] A. C. Santos and M. S. Sarandy, Generalized transitionless quantum driving for open quantum systems, Phys. Rev. A **104**, 062421 (2021).

[57] G. Passarelli, R. Fazio, and P. Lucignano, Optimal quantum annealing: A variational shortcut-to-adiabaticity approach, Phys. Rev. A **105**, 022618 (2022).

[58] A. Hartmann, G. B. Mbeng, and W. Lechner, Polynomial scaling enhancement in the ground-state preparation of Ising spin models via counterdiabatic driving, Phys. Rev. A **105**, 022614 (2022).

[59] M. W. Johnson, M. H. S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk *et al.*, Quantum annealing with manufactured spins, Nature **473**, 194 (2011).

[60] N. G. Dickson, M. W. Johnson, M. H. Amin, R. Harris, F. Altomare, A. J. Berkley, P. Bunyk, J. Cai, E. M. Chapple, P. Chavez *et al.*, Thermally assisted quantum annealing of a 16-qubit problem, Nat. Commun. **4**, 1903 (2013).

[61] N. Chancellor, Modernizing quantum annealing using local searches, New J. Phys. **19**, 023024 (2017).

[62] N. Chancellor and V. Kendon, Experimental test of search range in quantum annealing, Phys. Rev. A **104**, 012604 (2021).

[63] G. Passarelli, K.-W. Yip, D. A. Lidar, and P. Lucignano, Standard quantum annealing outperforms adiabatic reverse annealing with decoherence, Phys. Rev. A **105**, 032431 (2022).

[64] A. Callison, M. Festenstein, J. Chen, L. Nita, V. Kendon, and N. Chancellor, Energetic Perspective on Rapid Quenches in Quantum Annealing, Phys. Rev. X Quantum **2**, 010338 (2021).

[65] Y. Yu, C. Cao, C. Dewey, X.-B. Wang, N. Shannon, and R. Joynt, Quantum approximate optimization algorithm with adaptive bias fields, Phys. Rev. Res. **4**, 023249 (2022).

[66] V. Choi, Adiabatic quantum algorithms for the NP-complete maximum-weight independent set, exact cover and 3SAT problems, arXiv arXiv:10.48550/arxiv.1004.2226 (2010).

[67] T. Grass, Github repository: tobiasgrass/dwave_bias(2022).

[68] V. Choi, Minor-embedding in adiabatic quantum computation: I. The parameter setting problem, Quantum Inf. Process. **7**, 193 (2008).

[69] T. Boothby, A. D. King, and A. Roy, Fast clique minor generation in chimera qubit connectivity graphs, Quantum Inf. Process. **15**, 495 (2016).

[70] A. Palacios de Luis, A. Garcia-Saez, and M. P. Estarellas, Steered quantum annealing: improving time efficiency with partial information, arXiv arXiv:10.48550/arXiv.2206.07646 (2022).

[71] M. H. Amin, Searching for quantum speedup in quasistatic quantum annealers, Phys. Rev. A **92**, 052323 (2015).

[72] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, Optimization by simulated annealing, Science **220**, 671 (1983).