

Using quantum annealing to design lattice proteins

Anders Irbäck^{1,*}, Lucas Knuthson¹, Sandipan Mohanty², and Carsten Peterson¹¹Computational Biology & Biological Physics, Centre for Environmental and Climate Science (CEC), Lund University, 223 62 Lund, Sweden²Institute for Advanced Simulation, Jülich Supercomputing Centre, Forschungszentrum Jülich, D-52425 Jülich, Germany

(Received 3 November 2023; accepted 16 January 2024; published 13 February 2024)

Quantum annealing has shown promise for finding solutions to difficult optimization problems, including protein folding. Recently, we used the D-Wave Advantage quantum annealer to explore the folding problem in a coarse-grained lattice model, the HP model, in which amino acids are classified into two broad groups: hydrophobic (H) and polar (P). Using a set of 22 HP sequences with up to 64 amino acids, we demonstrated the fast and consistent identification of the correct HP model ground states using the D-Wave hybrid quantum-classical solver. An equally relevant biophysical challenge, called the protein design problem, is the inverse of the above, where the task is to predict protein sequences that fold to a given structure. Here, we approach the design problem by a two-step procedure implemented and executed on a D-Wave machine. In the first step, we perform a pure sequence-space search by varying the type of amino acid at each sequence position, and seek sequences which minimize the HP-model energy of the target structure. After mapping this task onto an Ising spin-glass representation, we employ a hybrid quantum-classical solver to deliver energy-optimal sequences for structures with 30–64 amino acids, with a 100% success rate. In the second step, we filter the optimized sequences from the first step according to their ability to fold to the intended structure. In addition, we try solving the sequence optimization problem using only the quantum processing unit (QPU), which confines us to sizes ≤ 20 , due to exponentially decreasing success rates. To shed light on the pure QPU results, we investigate the effects of control errors caused by an imperfect implementation of the intended Hamiltonian on the QPU, by numerically analyzing the Schrödinger equation. We find that the simulated success rates in the presence of control noise semiquantitatively reproduce the modest pure QPU results for larger chains.

DOI: [10.1103/PhysRevResearch.6.013162](https://doi.org/10.1103/PhysRevResearch.6.013162)

I. INTRODUCTION

Quantum annealing (QA) [1–4] is a promising method for finding good solutions to difficult optimization problems. In this method, one aims to find the solution to an optimization problem by encoding it into the ground state of a spin Hamiltonian. The approach of mapping optimization problems to spin systems is not new. It was used already in the 1980s in the context of neural networks [5,6]. By exploiting quantum fluctuations and quantum tunneling, QA offers a potentially much faster method for minimizing spin system energies. Technological advances like the D-Wave Advantage quantum annealer, with over 5000 quantum bits (qubits) and an average connectivity of 15 [7], permit exploration of the QA approach for a wide range of scientifically interesting problems, as illustrated by recent polymer and spin-glass studies [8,9]. While most of the problems studied concern optimization, attempts are also being made to use quantum annealers for sampling [10–12].

We have recently employed this machine for protein folding using the lattice-based HP (H: hydrophobic, P: polar) model [13] as a test bed [14]. Earlier attempts to use quantum computing for similar folding problems were based on chain growth, or turn-based, algorithms [15,16], with which nonlocal interactions like chain self-avoidance are challenging to implement unless the chain length N is very short ($N \lesssim 10$). In Ref. [14], we developed a scalable field-like representation, with qubits at all lattice sites, which made it possible to tackle chain lengths up to $N = 64$ using the D-Wave hybrid quantum-classical solver. To ensure proper chain configurations, the approach requires penalty terms with Lagrange parameters but is robust with respect to the choice of these parameters.

Another equally relevant problem from the biosector is the inverse folding problem [17–19], known as protein design, where one seeks to identify *a priori* unknown sequences which fold into a given structure. Since the structure of a protein is crucial for its function, this problem is highly relevant for drug design. It represents a computational challenge because it involves the exploration of both sequence and structure spaces.

Here, we tackle the design problem for lattice proteins using QA. We use the HP model as a test bed for which exact results are available for chain lengths $N \leq 30$ [20,21]. As is commonly done, we split the design problem into two steps. First, we generate candidate sequences that minimize the energy in the target structure, which we refer to as sequence

*anders.irback@cec.lu.se

optimization. Second, we determine whether the optimized sequences actually fold into the target structure by either folding the sequence using our aforementioned folding method, or by checking in the data bank of all solutions [20,21]. Prior work proposed both QA [22] and gate-based [23] schemes for the first of these subproblems, sequence optimization. Here, we present and test a complete QA approach to lattice protein design, comprising both of the above steps.

We find that the hybrid D-Wave annealer can efficiently handle both steps, and thus provides a fast and robust approach to the HP design problem. By contrast, relying entirely on the quantum processing unit (QPU) yields modest performance for larger chains. To understand this limitation, we develop and perform time-dependent Schrödinger equation simulations for the sequence optimization problem, using classical high-performance computing clusters. One potential cause of the decrease in success rate with problem size is control errors in the Hamiltonian, which may alter the ground state and thereby lead to incorrect solutions. Indeed, when adding control noise with its strength guided by hardware data, we obtain results that qualitatively reproduce the modest D-Wave pure QPU results.

The approach throughout this work has been evaluated using the D-Wave Advantage as the latter is, currently, the only available quantum annealer with an adequate qubit count. However, our methods should be valid for any quantum annealer.

II. METHODS

In protein design, one seeks a sequence $s = (s_1, \dots, s_N)$ that folds into a given target structure C_t . In general, multiple such sequences can exist, and any such sequence constitutes a valid solution. The probability of finding the chain with sequence s in the state C_t can be written as

$$P_\beta(s) = e^{-\beta E(C_t, s)} / \sum_C e^{-\beta E(C, s)}, \quad (1)$$

where $E(C_t, s)$ is the energy of the sequence s in state C_t (see Sec. II B), β is inverse temperature and the sum runs over all possible structures C . The design problem therefore translates to finding sequences near the maximum of $P_\beta(s)$. Methods for this task have been developed [24,25]. However, maximizing $P_\beta(s)$ involves a generally time-consuming search in both sequence and structure spaces. Therefore, a common approach is to first minimize the energy in the target structure over s , $E(C_t, s)$, followed by a filtering step to reject candidate sequences which have a higher probability for a different structure C_u . Running folding computations, which determine the most probable structure for a given sequence, is sufficient for this purpose. In this work, we address the design problem by this two-step procedure rather than directly maximizing $P_\beta(s)$.

A. HP lattice proteins

We consider the minimal two-dimensional (2D) lattice-based HP model for protein folding [13] in which the protein is represented by a self-avoiding chain of N hydrophobic (H) or polar (P) beads that interact through a pairwise contact

potential. A contact between two beads is said to occur if they are nearest neighbors on the lattice but not along the chain. The energy function can be written as $E_{HP} = -N_{HH}$, where N_{HH} is the number of HH contacts [13]. This definition renders the formation of a hydrophobic core energetically favorable. The ground state may be degenerate or unique. For a 2D square lattice, it is known from exhaustive enumerations that about 2% of all HP sequences with $N \leq 30$ have a unique ground state [20,21]. The availability of exact results for all sequences with $N \leq 30$ makes the 2D HP model a useful test bed for novel computational approaches.

Despite their simplicity, coarse-grained HP models are still relevant for the qualitative insights they provide into computationally challenging problems, such as protein folding and design (explored here), liquid-liquid phase separation of intrinsically disordered proteins [26,27], and protein evolution modeling [28,29].

B. HP sequence optimization in quadratic unconstrained binary optimization form

Given a target structure, C_t , we wish to minimize the energy $E_{HP}(C_t, s)$ over sequence, s , using a D-Wave quantum annealer. To this end, the problem must be recast in quadratic unconstrained binary optimization (QUBO) form, or, equivalently, into an Ising spin glass format. Furthermore, an auxiliary energy term needs to be included, to control the total number of H beads, N_H , in a candidate sequence of length N , since the all-H homopolymer sequence constitutes a trivial solution for unbiased $E_{HP}(C_t, s)$ minimization.

The only information needed about the target structure in order to compute $E_{HP}(C_t, s)$ is its connectivity matrix w_{ij} , which indicates whether two arbitrary beads i and j are in contact ($w_{ij} = 1$) or not ($w_{ij} = 0$). This holds for any model with pairwise contact interactions, irrespective of the dimensionality and the size of the amino acid alphabet. When using the HP model, a suitable choice of total energy $E(s)$ to minimize is given by

$$E(s) = - \sum_{1 \leq i < j \leq N} w_{ij} s_i s_j + \lambda \left(\sum_{i=1}^N s_i - N_H \right)^2, \quad (2)$$

where s_i describes whether bead i is of type P ($s_i = 0$) or H ($s_i = 1$). In Eq. (2), the first term represents $E_{HP}(C_t, s)$, whereas the second term biases the total number of H beads toward a preset value, N_H . The balance between the two terms is set by the parameter λ . The 0,1 spins of Eq. (2) can be easily transformed into Ising ± 1 spins without losing the desired quadratic structure of E . This energy function has a much simpler structure than the corresponding one for the folding problem in Ref. [14], requiring only one Lagrange parameter λ instead of the three in the folding study.

This Lagrange parameter λ must be sufficiently large for the generated sequences to acquire the desired composition, as set by N_H . On the other hand, if λ is too large, the energy landscape becomes rugged. Examples of how the efficiency of the hybrid quantum-classical solver varies with λ can be found in Sec. III A. All hybrid production runs were carried out using $\lambda = 2.5$. For the pure QPU computations, which are




FIG. 1. The 30-, 50-, and 64-bead target structures T_{30} , T_{50} , and T_{64} used in the sequence design computations with the hybrid solver.

limited to smaller systems, it was possible to use a smaller value, set to $\lambda = 1.1$ (Sec. III B).

Minimizing $E(s)$ in Eq. (2) can be seen as a graph bisection problem. Unlike a spin glass with nearest-neighbor interactions, it is a fully connected system. Note also that, compared with the HP-folding problem [14], much fewer spin variables are needed, since only the H or P identity of the beads (whose locations and contacts are fixed for the target structure) needs to be encoded.

For all instances studied in the present paper, it is possible to infer the minimum E_{HP} for a given N_H , by inspection of the bead-bead contacts present in the target structure (see Appendixes A and B). Hence, it is possible to decide whether a proposed solution is correct, even without additional calculations.

C. Hybrid quantum-classical computations

As an alternative to pure QPU computation, the D-Wave Advantage system also offers access to a hybrid quantum-classical solver [30]. The hybrid approach uses classical solvers while sending suitable subproblems as queries to the QPU, to speed up the execution and improve the solutions of challenging QUBO problems. Given the limited connectivity (15) within the D-Wave Advantage architecture, the hybrid approach is particularly relevant. With the hybrid solver, it is possible to tackle problem sizes much larger than with pure QPU computation.

Using the hybrid solver, we performed sequence optimization for three target structures with $N = 30$, $N = 50$, and $N = 64$ (see Fig. 1). For each target structure C_t , we searched for minimum- $E_{HP}(C_t, s)$ sequences s , for several fixed compositions, N_H . For each combination of C_t and N_H , we conducted a set of 10 runs, thus generating a set of up to 10 optimized sequences.

In the hybrid approach, the runtime needs to be chosen with some care. Indeed, in our previous HP folding study [14], the success rate of the hybrid solver for $N > 30$ was poor for short runtimes, while rapidly increasing to values close to 100% once the runtime passed a system size dependent threshold. To determine runtimes for the sequence optimization problem, we performed a set of preliminary runs for our largest target structure ($N = 64$), using $N_H = 42$. The hybrid solver consistently returned sequences with the known minimum E_{HP} (Appendix A) for runtimes ranging from 15 s down to the shortest possible time of 3 s, which is also the default runtime

for the D-Wave Advantage hybrid solver. Therefore, all the production runs were carried out using this default runtime.

To test whether the generated sequences actually fold to the desired target structures, we need to perform folding calculations. To this end, we also employ the D-Wave hybrid solver given its demonstrated power for the folding problem [14]. Here, for a given optimized sequence s_o , the energy $E_{HP}(C, s_o)$ was minimized over chain structure C , using the methods and parameters in Ref. [14] and a 10×10 grid. Based on the findings in Ref. [14], the runtime was set to 4, 120, and 300 s for $N = 30$, $N = 50$, and $N = 64$, respectively. These runtimes are larger than the threshold times, above which the success rate was shown to be high [14].

D. Pure quantum processor unit computations

The Pegasus topology of the D-Wave Advantage QPU connects each of its qubits to 15 others [7]. Problems requiring higher connectivity have to be embedded into the Pegasus graph. This embedding is done by forming “chains” of qubits which act as single qubits. The strength of the coupling between the qubits within a chain is a tunable parameter, called the chain strength. This parameter is typically chosen slightly larger than the minimum chain strength needed to avoid having too many chain breaks (see Sec. III C).

D-Wave offers several so-called samplers for finding embeddings into the QPU topology and performing the QPU computation. We used the DWAVECLIQUESAMPLER, designed for dense binary quadratic models [31]. It has the property that the chains representing logical qubits share a common length, which facilitates the analysis in Sec. III C. With this method, the number of physical qubits was three times the number of logical qubits in almost all instances studied. For the smallest system (with $N = 10$), this ratio was two instead of three. All the computations used a chain strength between 2.25 and 4.25 (Appendix B). The annealing time was set to $t_f = 2000 \mu s$, its maximum allowed value. The number of output reads per run (annealing cycles), which must be $< 10^6 \mu s t_f^{-1}$, was set to 100.

E. Time-dependent Schrödinger equation simulations

As will be seen in Sec. III B, the pure QPU performance deteriorates rapidly with system size. In an attempt to understand this phenomenon, we perform quantum-mechanical simulations with the time-dependent Schrödinger equation.

Similarly, the pure QPU performance was also meager in the folding case [14]. With its simple form [Eq. (2)], the sequence design problem is better suited for analyzing the shortcomings of the pure QPU performance as compared with the folding case.

We consider an N -qubit system governed by a time-dependent Hamiltonian

$$H(t) = a(t)H_D + b(t)H_P, \quad (3)$$

where H_D and H_P are the driver and problem Hamiltonians, respectively. On a D-Wave annealer, these two terms take the forms

$$H_D = \sum_i \sigma_i^x, \quad H_P = \sum_i h_i \sigma_i^z + \sum_{i<j} J_{ij} \sigma_i^z \sigma_j^z, \quad (4)$$

where σ_i^x and σ_i^z denote Pauli matrices. For specificity, we assume a linear annealing schedule given by $a(t) = 1 - t/t_f$ and $b(t) = t/t_f$, where t_f is the annealing time. We integrate the Schrödinger equation with this Hamiltonian in time t (with $\hbar = 1$) using the formalism and algorithm described in Appendix C. The units for time and energy are arbitrary but related through $\{\text{unit of time}\} \times \{\text{unit of energy}\} = \hbar$.

In addition to the driver Hamiltonian H_D in Eq. (4), we also consider the so-called XY -mixer [32], given by

$$H_D^{XY} = \frac{1}{2} \sum_{i<j} (\sigma_i^x \sigma_j^x + \sigma_i^y \sigma_j^y), \quad (5)$$

which is currently not available on D-Wave's annealers. Transitions generated by this mixer have the property of leaving the constrained sum in Eq. (2) unchanged. Hence, if the H_D^{XY} mixer is used, and the initial state is a uniform superposition of all states satisfying the constraint, it would be possible to minimize E_{HP} at a fixed N_H without including the constraint term.

F. Test bed: HP target structures

We seek HP sequences that fold to given target structures with 10–30, 50, and 64 beads. For $N \leq 30$, all HP sequences with unique ground states and the corresponding structures are known from exhaustive enumerations [20,21]. This data can be used to decide whether a generated sequence actually folds to the desired structure. To evaluate 50- and 64-bead sequences, for which no such data are available, we determine minimum-energy structures by using the D-Wave hybrid solver as described in Ref. [14].

III. RESULTS

Given a target structure C_t and a composition N_H , we wish to find minimum- E_{HP} sequences by minimizing the energy E in Eq. (2) on a quantum annealer. In all instances studied, the minimum E_{HP} is known (Appendixes A and B), so it is possible to decide whether an obtained sequence is a correct solution to the sequence optimization problem.

A sequence that minimizes E_{HP} in C_t may, however, have the same or even lower energy in other structures $C_u \neq C_t$. Whether or not this is the case can be checked against existing exact results if $N \leq 30$ [20,21], or by performing an energy

minimization in the structure space using the hybrid quantum-classical computations, as described in Ref. [14].

D-Wave offers solvers based entirely on quantum annealing, as well as a hybrid quantum-classical scheme. In this section, we first try out the hybrid quantum-classical approach with success, even for large chains (Sec. III A). After that, in Sec. III B, we examine the effectiveness of pure QPU calculations, without the classical preprocessing involved in the hybrid approach, using smaller target structures. We observe a rapid decrease in the pure QPU hit rate when increasing the size of the structures. In Sec. III C, we attempt to explain this observation by numerically solving the Schrödinger equation on classical computers.

A. Hybrid quantum-classical computations

Using the hybrid solver, we conducted sequence design for the three target structures shown in Fig. 1 with $N = 30$, 50, and 64, which will be referred to as T_{30} , T_{50} , and T_{64} , respectively. For each target structure, we used a few different compositions, N_H (Appendix A).

With no exceptions, the hybrid solver generated sequences with the known minimum E_{HP} (Appendix A) in the target structure, with a 100% success rate.

Of note, there are sequences that minimize E_{HP} in the target structure without folding to this structure. For such a sequence, the target structure may be one of multiple structures in a degenerate ground state. Alternatively, there exists at least one other structure in which E_{HP} is lower than it is in the target structure, so that a subsequent energy minimization in the conformation space yields a different structure. In such cases, those sequences are not solutions of the design problem for the target structure, and are discarded, even when they are valid solutions for the first step of our two-step approach. This is the price we pay for foregoing an expensive simultaneous search in sequence and structure spaces in favor of a sequence space search as the first step. The sequences emerging from sequence space minimization must be filtered by their ability to fold to the target structure.

Note also that our search for candidate sequences below is not exhaustive; further minimum-energy sequences may exist. Our goal is to find some sequence that folds into the target structure, not all such sequences.

1. Target structure T_{30}

Our first target structure, T_{30} (Fig. 1, left panel), is known from exact results [21] to be the unique ground state of >800 HP sequences. Using the hybrid solver, we minimized E in Eq. (2), with $\lambda = 2.5$, for this structure for several compositions, $12 \leq N_H \leq 17$. For every N_H , 10 hybrid runs all successfully returned sequences with the known minimum E_{HP} .

While most of the thus generated sequences had T_{30} as their unique ground state [21], some of them (all with $N_H = 12, 13$, or 16) did not. For each of the latter, a search for possible structures with lower energy was performed, using the hybrid solver (Sec. II C). No such structure was found, which suggests that the ground states for those sequences are degenerate, and T_{30} is one of the structures having the lowest energy.

2. Target structure T_{50}

The second target structure, T_{50} (Fig. 1, middle panel), comes from a study of a Monte Carlo-based sequence design algorithm [24], which actually optimizes the target population, Eq. (1), rather than the energy in this structure. The best sequence found in that study contained 31 H beads [24].

Here, we searched for sequences minimizing the energy E_{HP} of the T_{50} structure for $N_H = 29, 30$, and 31, using Eq. (2) with $\lambda = 2.5$ and the hybrid solver. For $N_H = 31$, the solution to the E_{HP} -minimization problem is unique (Appendix A) and given by the sequence identified in Ref. [24].

As in the T_{30} case, for every N_H , all 10 hybrid runs successfully gave sequences with the known minimum E_{HP} (Appendix A). For $N_H = 29$ and 30, where the minimum- E_{HP} level is degenerate (Appendix A), the number of distinct sequences obtained from the 10 runs were six and four, respectively.

For each of the 11 distinct optimized sequences, we subsequently minimized E_{HP} over structure by a set of 10 hybrid runs (Sec. II C). Six of the sequences turned out to have T_{50} as one of the structures at the lowest lying E_{HP} minimum. For the remaining five sequences, T_{50} was the only structure at the global minimum of E_{HP} . Among them, was the sequence with $N_H = 31$ from Ref. [24]. The results obtained here support the conclusion that T_{50} is the unique ground state of the sequence found in Ref. [24], while at the same time finding several new solutions to the design problem for T_{50} .

3. Target structure T_{64}

The target structure T_{64} (Fig. 1, right panel) has the lowest known energy for an HP sequence with 42 H beads that has been extensively studied [14,33,34]. Although not a unique energy minimum, all known structures sharing the same energy ($E_{HP} = -42$) have a similar shape. The structures differ only within the entirely hydrophobic core, where the chain can be rearranged without altering E_{HP} .

Using Eq. (2) with $\lambda = 2.5$ and the hybrid solver, we minimized E_{HP} in the target structure T_{64} for $N_H = 36, 40$, and 42. As for the previous two target structures, the hybrid solver consistently found sequences with the known minimum E_{HP} (Appendix A) for all N_H values. In the $N_H = 36$ and 42 cases, where multiple solutions exist (Appendix A), the returned sequence varied from run to run.

As in the case of T_{30} and T_{50} , we searched for possible structures with lower E_{HP} using a set of 10 hybrid runs (Sec. II C) per optimized sequence. For every sequence, all runs returned structures with the same E_{HP} as the target structure. The results thus suggest that T_{64} is a minimum- E_{HP} structure for all these sequences. However, the minimum is not unique for any of the optimized sequences with $N_H = 40$ or 42. For all these sequences, the core of the T_{64} structure is entirely hydrophobic, which makes it possible to rearrange the chain without changing E_{HP} .

Interestingly, the situation appears to be different for one of the optimized $N_H = 36$ sequences, shown in Fig. 2, which has P beads at four core positions. For this sequence, all 10 hybrid folding runs returned the target structure. Changing these four beads to P in the otherwise hydrophobic core appears to lift the degeneracy of the minimum- E_{HP} level.




FIG. 2. An optimized HP sequence that appears to have the target structure T_{64} as its unique ground state. The sequence is composed of 36 H (filled) and 28 P (open) beads. For the sequence with 42 H beads studied in Refs. [33,34], this structure is one of several with minimum energy. The degeneracy arises because, for that sequence, the chain can be rearranged in the core of the structure without altering the energy.

The above results show that the hybrid quantum-classical method efficiently solves the sequence optimization problem for all the systems studied. All these computations were done with the Lagrange parameter in Eq. (2) set to $\lambda = 2.5$. To gauge the sensitivity of the success rate to changes in λ , we conducted additional sets of hybrid runs for the three systems (T_{30} , T_{50} , and T_{64} with $N_H = 15, 31$, and 36, respectively), using the default runtime. We found that $\lambda > 0.25$ was necessary for any correct solutions to be found in all three cases. The lower limit on λ is problem-dependent but not larger than 2.0 for any of the systems studied in this paper. As shown in Fig. 3(a), large values for λ also lead to performance degradation, although there is a wide window of λ values having a 100% hit rate, showing that no excessive fine tuning of λ is required.

The low hit rates at large λ in Fig. 3(a) can be improved at the cost of increasing the runtime. Figure 3(b) shows the runtime required to attain a 50% success rate $\tau_{1/2}$ plotted against λ . For the two larger problems with $N = 50$ and $N = 64$, respectively, at $\lambda = 7$, the hit rate is tiny when using a runtime of 3 s [Fig. 3(a)], but can be improved to 50% by increasing the runtime to about 100 s [Fig. 3(b)]. Note that $\tau_{1/2}$ grows faster with λ for the two larger systems than it does for the $N = 30$ system. Not unexpectedly, Figs. 3(a) and 3(b) both show that the $N = 30$ problem is significantly easier than the other two.

To summarize, here we have designed sequences for three target structures using a two-step procedure involving energy minimization in the sequence and structure spaces. With this widely used approach, a pure sequence space search is performed first, where the types of the amino acids at each position in the target structure are treated as the optimization parameters. An optimized sequence from the first stage is accepted as a solution to the design problem only if a subsequent minimization in the conformation space finds the target structure to be the global minimum for that sequence. We stress that, for both tasks, the hybrid solver gave reliably




FIG. 3. Dependence of the hit rate on the Lagrange parameter λ [Eq. (2)] when solving the sequence optimization problem by hybrid quantum-classical computations for the target structures T_{30} , T_{50} , and T_{64} (Fig. 1) with $N_H = 15, 31$, and 36 , respectively. (a) Hit rate as a function of λ when using the default runtime, which was 3 s for all systems. Each data point represents an average over 100 runs. (b) The time required to attain 50% hit rate, $\tau_{1/2}$, plotted on a log scale against λ . The horizontal dotted line indicates the default runtime (3 s).

good results, with robustness with respect to the Lagrange parameter settings.

B. Pure quantum computations

In this section, rather than using the D-Wave hybrid solvers as in the previous section, we explore the ability of pure QPU computations to solve the sequence optimization problem for $10 \leq N \leq 20$ target structures. We pick, as our target, the most designable structure for a given N denoted by T_N , where the designability of a structure is defined as the number of sequences sharing it as their unique ground state. This number is known for all structures with $N \leq 30$ from exhaustive enumerations [21]. The same data bank [21] also provides exact answers to whether the generated sequences actually fold to the target structures.

For each of the target structures T_{10} – T_{20} , we performed pure QPU computations for one or a few choices of N_H (Appendix B), using the DWAVECLIQUESAMPLER. As in the case of the structures in Sec. III A, the minimum E_{HP} , given N_H , can be inferred from the contacts present in the target structure, and the solution may be unique or degenerate (Appendix B).

The pure QPU computations recovered all possible solutions to the sequence optimization problem for every (T_N, N_H) pair. To quantify the success rate of the pure QPU computations, 10 000 annealing cycles were generated for each combination of T_N and N_H . The fraction of these yielding a correct solution is referred to as the hit rate as was done in Sec. III A. For this purpose, a correct solution is a solution sequence whose energy matches the previously known minimal E_{HP} for the (T_N, N_H) pair. The hit rates obtained in the pure QPU runs can be found in Fig. 4(a).

As noted previously, a sequence that minimizes the energy in the target structure, for a given N_H , may not fold to that structure. However, for every (T_N, N_H) combination studied here, there is at least one solution to the sequence optimization problem that has the target structure T_N as its unique ground state. The precise number of such solutions for different (T_N, N_H) pairs can be found in Appendix B.

For the generated sequences that did not have the target structure as its unique ground state, we minimized E_{HP} over

structure using the hybrid solver (Sec. II C). In these runs, for almost all the sequences, we found other structures with the same energy as the target structure but none with lower energy, suggesting that the target structure is part of a degenerate ground state. The only exceptions occurred for the target structure T_{13} and $N_H = 6$. In this case, five of 18 minimum- E_{HP} sequences attained a lower energy in other structures. Nothing in our procedure precludes the existence of lower-energy structures for a sequence obtained by minimizing the target structure energy. However, had such situations been more common, the recipe used here (sequence space optimization followed by filtering based on folding runs) would not be effective. Note that, for T_{13} and $N_H = 6$, the existence of lower-energy alternatives is intuitively unsurprising, considering that the target structure has a relatively high energy for the given N_H (Appendix B).

Although the pure QPU computations recovered all possible solutions to the sequence optimization problem, the hit rate was strongly problem-dependent [Fig. 4(a)]. While also depending on N_H , the hit rate shows a clear decreasing trend with the problem size N , which limits the range of N which is meaningful to study.

There are several factors that may contribute to the rapid decay of the pure QPU hit rate with problem size as seen in Fig. 4(a). These include (i) finite- t_f effects, (ii) chain breaks, (iii) thermal noise, and (iv) control errors. Here, we briefly comment on factors (i)–(iii), whereas factor (iv), which relates to the implementation of the couplings J_{ij} and h_i in Eq. (2), will be discussed in Sec. III C below.

(i) *Finite t_f .* An obvious potential source of error is the use of a finite annealing time t_f (≤ 2000 μ s). The t_f dependence of the pure QPU hit rate is illustrated in Fig. 5(a) by data obtained for the target structure T_{12} and two values of N_H (4 and 6). In both cases, the hit rate does increase with t_f for small t_f . However, it levels off around $t_f = 400$ μ s for $N_H = 4$ and already before $t_f = 100$ μ s for $N_H = 6$, at values well below unity. This behavior suggests that there must be other, more important error sources than the upper limit on t_f . This conclusion is further supported by results from numerically integrating the Schrödinger equation (Sec. II E). Here, we computed the ground-state probabilities P_g for the different systems at $t = t_f$, for a fixed t_f [Fig. 5(b)]. With our choice




FIG. 4. (a) Hit rate when minimizing $E(s)$ in Eq. (2) by pure QPU computation for systems with $10 \leq N \leq 20$ and different N_H (Appendix B). The value of N_H is indicated by the plot symbol, and the Lagrange parameter was set to $\lambda = 1.1$. For each system, a set of 10 000 annealing cycles was generated using the DWAVECLIQUESAMPLER, and the hit rate is the fraction of these that gave a correct solution. The annealing time was set to its maximum value, $t_f = 2000 \mu s$. The chain strength was chosen individually for each system among the values 2.25, 2.50, ..., 4.25 for best performance. A circle around the plot symbol indicates that the system has a large energy gap, $\Delta E \geq 1.0$ (Appendix B). The dashed line is a least-square fit to the encircled data points. (b) Same as panel (a) after restricting the analysis to a filtered dataset without chain breaks (77% of the full dataset) and removing systems with a small energy gap ($\Delta E < 1.0$).

of $t_f = 20$ a.u., the P_g data roughly match the measured pure QPU hit rates for small N [Fig. 4(a)] but not the rapid decay with N seen in the latter case.

(ii) *Chain breaks.* When embedding the problems into the Pegasus graph, the pure QPU solver creates chains of strongly coupled physical qubits collectively behaving as a single qubit. In computations, it may happen that such chains, representing logical qubits, break. To check how our results are affected by such chain breaks, we recomputed the pure QPU hit rates using data only from annealing cycles in which no chain break occurred (77% of the full dataset). A scatter plot comparing the original and recomputed pure QPU hit rates can be found in Appendix D. In many cases, the removal of chain breaks leads to a statistically significant change of the measured hit rate, although the overall agreement between the two datasets is quite good (Pearson correlation coefficient 0.94),

(iii) *Thermal noise.* The effects of thermal noise are likely to be more severe if the energy gap ΔE between the ground

state and the first-excited state of the problem Hamiltonian is small. In our systems, it turns out that ΔE takes on one of three possible values, namely, 0.1, 1.0, and 1.1 (Appendix B). In Fig. 4(a), there are three systems with markedly lower hit rates than the others, all of which have a small energy gap $\Delta E = 0.1$. It is conceivable that the low hit rates for these systems, at least in part, is due to thermal noise. However, modeling thermal effects is difficult without extensive details of the underlying physics of the D-Wave processors.

In what follows, we remove from the analysis systems where thermal effects are potentially much stronger than in the others by focusing on systems with $\Delta E \geq 1.0$. Furthermore, we use the filtered dataset without chain breaks, for a cleaner comparison with results from Schrödinger simulations. Redrawing Fig. 4(a) after making these two restrictions, we obtain Fig. 4(b), where the pure QPU hit rate falls off roughly exponentially with N , albeit still with some scatter.




FIG. 5. Effects of using a finite annealing time t_f in solving the sequence optimization problem [Eq. (2) with $\lambda = 1.1$]. (a) Hit rate against t_f in pure QPU computations for two of the problems in Fig. 4 (target structure T_{12} , $N_H = 4$ and 6). (b) Ground-state probability P_g at $t = t_f$ when numerically integrating the Schrödinger equation (Sec. II E) for systems with $10 \leq N \leq 20$ (Appendix B) using a fixed finite t_f ($t_f = 20$ a.u.), plotted against system size, N . The value of N_H is indicated by the plot symbol. The Hamiltonian is $H(t) = a(t)H_D + b(t)H_P$ [Eqs. (3) and (4)].

C. Probing the effects of control noise on the success rate of a pure quantum processor unit

One potentially limiting factor in the pure QPU computations is analog control errors in the fields h_i and couplers J_{ij} [35,36] of the Ising Hamiltonian H_P [Eq. (4)], which are referred to as integrated control errors in D-Wave's documentation [31]. The presence of control errors δh_i and δJ_{ij} leads to a perturbed Hamiltonian

$$\tilde{H}_P = \sum_i (h_i + \delta h_i) \sigma_i^z + \sum_{i < j} (J_{ij} + \delta J_{ij}) \sigma_i^z \sigma_j^z, \quad (6)$$

whose ground state may not coincide with that of the intended Hamiltonian H_P [Eq. (4)]. Previous work showed that small errors in individual parameters collectively can cause an exponential decay of the success rate with problem size [35,36].

To be able to explore the effects of control errors in our pure QPU computations, we have to make some simplifying assumptions. First, following Refs. [35,36], we assume that all errors δh_i and δJ_{ij} are statistically independent and normally distributed, with zero mean and standard deviations σ_h and σ_J for all δh_i and δJ_{ij} , respectively. Second, for computational reasons, we consider only logical qubits, thus essentially ignoring the auxiliary qubits needed when embedding the problems into the QPU topology. However, we take the QPU embedding into account in setting the values of σ_h and σ_J (see below).

In D-Wave QPU computations, all couplers and fields are rescaled to $\hat{h}_i = h_i/r$ and $\hat{J}_{ij} = J_{ij}/r$, where r is the smallest number such that all rescaled parameters fall in given intervals, $|\hat{h}_i| \leq h_{\max}$ and $|\hat{J}_{ij}| \leq J_{\max}$. In all systems studied here, the rescaling factor r is set by the chain strength J_{cs} (see Sec. II D) and given by $r = J_{cs}/J_{\max}$. In particular, this implies that the energy gap of the rescaled problem Hamiltonian scales as $1/J_{cs}$, which should lead to a decrease in success rate with increasing J_{cs} .

The strength of the control noise on D-Wave's systems has been investigated [31]. D-Wave Support suggests using $\sigma_h = x \max |h_i|$ and $\sigma_J = x \max |J_{ij}|$ with $x = 0.015$, where the maxima are taken over all fields and couplers of the Hamiltonian, including those associated with auxiliary qubits. As indicated above, in our systems, the largest $|J_{ij}|$ is the chain strength J_{cs} . Following the suggestion above, we therefore set $\sigma_J = x J_{cs}$. Each logical qubit is represented by a chain of k physical qubits (Sec. II D), where $k = 2$ for the $N = 10$ system and $k = 3$ for all other systems studied. As the physical qubits representing a logical qubit with field h_i have fields h_i/k , we set $\sigma_h = \sqrt{k} x \max |h_i|/k$, where the square root comes from summing over k physical qubits. Summarizing, we then have

$$\sigma_h = \frac{x \max |h_i|}{\sqrt{k}} \text{ and } \sigma_J = x J_{cs}. \quad (7)$$

Using Eq. (7) with k and J_{cs} as in the pure QPU computations and $x = 0.015$, we generated 10 000 perturbed Hamiltonians $\tilde{H}_P$ [Eq. (6)] for each of the systems in Fig. 4(b), and determined the fraction of these, P_g , that shared ground state with the intended Hamiltonian H_P . For the two systems with $(N, N_H) = (10, 4)$ and $(N, N_H) = (20, 11)$, we




FIG. 6. The dependence of the hit rate on the chain strength, J_{cs} , in pure QPU computations and in Schrödinger simulations with control noise of strength $x = 0.015$, for two systems with $(N, N_H) = (10, 4)$ (black) and $(N, N_H) = (20, 11)$ (red), respectively (Appendix B). Full lines represent data from the Schrödinger simulations. Plot symbols show pure QPU hit rates, calculated over the full dataset (plus) or the filtered dataset without chain breaks (circle). Dotted lines are drawn to guide the eye. Data presented elsewhere in the paper for these systems were obtained using $J_{cs} = 2.25$ for $(N, N_H) = (10, 4)$ and $J_{cs} = 4.25$ for $(N, N_H) = (20, 11)$.

performed additional sets of pure QPU computations to elucidate how the hit rate depends on J_{cs} .

Figure 6 shows the J_{cs} dependence of the pure QPU hit rate for these two systems, along with the simulated ground-state probability P_g . Clearly, J_{cs} must be sufficiently large to ensure chain stability. On the other hand, we expect the hit rate to drop if J_{cs} gets too large, due to the $1/J_{cs}$ scaling of the energy gap (see above). The data confirm that J_{cs} must be neither too small nor too large for good performance (Fig. 6), and therefore needs to be chosen with some care. It can also be seen that the hit rate depends only weakly on whether the full dataset or the filtered one without chain breaks is used. As chain breaks do not occur in the simulated systems with only logical qubits, P_g does not drop at small J_{cs} . At large J_{cs} , P_g decreases at a rate similar to what is observed for the pure QPU hit rate.

Figure 7(a) shows the simulated hit rates P_g , for the systems studied in Fig. 4(b), plotted against N . As in the pure QPU computations [Fig. 4(b)], the hit rate falls off roughly exponentially with N . The scatter plot in Fig. 7(b) compares simulated hit rates P_g [Fig. 7(a)] with pure QPU hit rates [Fig. 4(b)]. The pure QPU hit rate appears to decrease somewhat more slowly than P_g with N . Nevertheless, at a semiquantitative level, the pure QPU hit rates agree quite well with the P_g values obtained using the suggested noise strength $x = 0.015$. While refraining from attempts to fine-tune x , we note that using $x = 0.003$ or $x = 0.030$ leads to, respectively, too high or too low P_g values, compared with the QPU hit rates.

In conclusion, the results presented here are consistent with the hypothesis that control errors are an important factor behind the strong N dependence of the pure QPU hit rates. To fully account for the measured hit rates, and in particular




FIG. 7. (a) N dependence of hit rates from Schrödinger simulations with control noise of strength $x = 0.015$. The function optimized is $E(s)$ in Eq. (2) with $\lambda = 1.1$ and different N_H , as indicated by the plot symbol. The systems studied are the same as in Fig. 4(b). For each system, 10 000 perturbed Hamiltonians $\tilde{H}_p$ [Eq. (6)] were generated, and the hit rate is the fraction of these that share ground state with the noise-free Hamiltonian H_p . The dashed line is a least-square fit. (b) Scatter plot comparing hit rates from the Schrödinger simulations in panel (a) with the pure QPU hit rates in Fig. 4(b).

their N_H dependence, additional factors need to be considered, including thermal noise.

There are complementary methods, not explored here, to improve on modest pure QPU hit rates. One is to add a postprocessing step, in which the QPU output state is subject to a local optimization with a greedy classical algorithm [37], to drive approximately correct solutions to the desired minimum-energy level. However, in most systems studied here (Fig. 7), all first-excited states are local minima left unchanged by this method, which suggests that it is of limited use for our systems. Another approach, called “shimming,” is to refine the calibration of the Hamiltonian by using symmetries of the system [38]. These symmetries typically assume zero fields, and coupler values related by a sign flip ($\pm J$). These conditions are not met in our systems, which makes useful symmetries hard to find.

Let us finally mention that we also investigated whether the XY-mixer offers an advantage in our problem, by comparing the time evolution of the Schrödinger system when using respectively H_D and H_D^{XY} in Eqs. (4) and (5) as drivers. In the cases studied, we found that the improvement is at most marginal.

IV. SUMMARY AND OUTLOOK

Protein design, determining sequences corresponding to a given structure, is a highly relevant biophysical problem. Since both sequence and structure space have to be explored, the problem is very challenging, especially for large chains. We have approached this problem for the HP lattice protein model using quantum annealing by first minimizing the target structure energy to generate sequences and then checking if the generated sequences do in fact fold to the target structure.

The approach was evaluated by using the D-Wave Advantage hybrid quantum-classical solver for three structures with chain lengths $N = 30$, $N = 50$, and $N = 64$. Without exceptions, the D-Wave hybrid solver swiftly solves the sequence optimization problem, for which the ground-state energy can be deduced (Appendix A). These solutions were then tested for their ability to fold to the target structure, a problem

that can also be efficiently tackled using the hybrid solver [14]. The two-step procedure was successfully applied to all three target structures. In particular, we identified a previously unknown sequence that appears to have the $N = 64$ structure as its unique ground state.

In addition, we tested using the D-Wave pure QPU for sequence optimization problems with $10 \leq N \leq 20$, for which solutions exist in the data bank. For all structures, sequences that had the desired structure as their unique ground state were found. However, in line with previous results for the folding problem [14], when using only the QPU for the sequence optimization problem, the performance deteriorates with system size. To understand this behavior, which very likely is due to hardware-induced phenomena (noise), we solved the time-dependent Schrödinger equation numerically for different scenarios.

Two possible sources of error are inadequate annealing time and control errors in the couplers and fields of the problem Hamiltonian. Whereas inadequate annealing time turned out not to be the problem, our results suggest that control errors have a significant impact on the success rate. Here, we computed the fraction of perturbed Hamiltonians sharing ground state with the original Hamiltonian and compared with the pure QPU results. Employing the same rescaling on the D-Wave machine and using standard deviations of the noise supplied by D-Wave, we found a semiquantitative agreement with the observed pure QPU hit rates [Fig. 7(b)]. A more detailed error model should also incorporate thermal noise. The latter would have to be based on deeper knowledge about the inner workings of the D-Wave architecture, which is, at present, not accessible.

These noise investigations were exclusively probing the D-Wave QPU properties for the simple reason that D-Wave is the only available annealer for realistic computations. However, our approach to elucidate the problem by comparing with the time-dependent Schrödinger equation is of wider relevance.

Overall, it is clear that the quantum annealer naturally lends itself to protein design. Using the hybrid quantum-classical annealer, both generating sequences, and subsequently filtering them based on their folding ability, work very well. When

using just the QPU, without leveraging D-Wave's hybrid solver, we found that generating sequences with the lowest HP energy for the target structure becomes difficult for larger structures. The results presented lend support to the notion that, at least in part, this difficulty originates from imperfect implementation of the problem Hamiltonian.

Replacing the X driver with an XY driver in the Schrödinger simulations, which has been suggested for quantum annealing in general [32], did not improve the results for the sequence optimization problem. This is somewhat surprising since in our case it would remove the only constraint term in Eq. (2).

Toward more realistic protein models, one could replace the binary HP encoding spins s_i by discrete multistate spins S_i encoding both amino acid type, e.g., in the canonical 20-letter alphabet, and the corresponding side-chain conformations, or rotamers [22]. Assuming an energy with the quadratic structure $E = \sum_{i=1}^N A(S_i) + \sum_{i<j} B(S_i, S_j)$ and a given target backbone conformation, one could then determine the amino acid sequence and rotamers by minimizing E over the S_i variables. Given all possible values of all one- and two-body terms $A(S_i)$ and $B(S_i, S_j)$, which would have to be predetermined, this minimization could in principle be carried out using QA and a one-hot encoding of the spins S_i . However, checking whether the optimized sequences fold to the intended and real, rather than lattice-based, backbone conformation would require classical computing. One possibility would be to use the AlphaFold structure prediction method [39,40].

ACKNOWLEDGMENTS

This work was in part supported by the Swedish Research Council (Grant no. 621-2018-04976). We gratefully acknowledge the Jülich Supercomputing Centre [41] for supporting this project by providing computing time on the D-Wave Advantage™ System JUPSI through the Jülich UNified Infrastructure for Quantum computing (JUNIQ; budget ID PDQA2023). We acknowledge helpful conversations with Ken Robbins from D-Wave, and have also benefited from discussions with Göran Johansson, Hanna Linn, and Laura García-Álvarez at the Wallenberg Centre for Quantum Technology at Chalmers University of Technology.

APPENDIX A: MINIMUM E_{HP} FOR THE TARGET STRUCTURES T_{30} , T_{50} , AND T_{64}

Given a target structure C_t , we determine sequences s by minimizing $E_{\text{HP}}(C_t, s)$ at different fixed compositions N_H , with E_{HP} defined as minus the number of HH contacts. In all instances studied in this paper, the solution to this problem can be inferred from a contact map showing all contacts present in the target structure. Figure 8 shows contact maps for the three largest target structures studied: T_{30} , T_{50} , and T_{64} (Sec. III A).

To illustrate how the minimum E_{HP} can be found, consider the contact map for T_{64} (Fig. 8, right panel) for the three N_H values used (Sec. III A), namely, 36, 40, and 42. In this structure, 10 of the beads do not form any contact, and can therefore be ignored. The remaining beads can be divided into 12 distinct clusters: seven two-bead clusters, three four-bead clusters and two larger clusters with 12 and 16 beads, respectively. The five clusters with more than two beads contain 40 beads in total. It is easy to see that taking these 40 beads as H represents a unique E_{HP} minimum for $N_H = 40$ ($E_{\text{HP}} = -41$). For $N_H = 42$, the minimum E_{HP} is obtained by adding any of the seven two-bead clusters to the set of H beads ($E_{\text{HP}} = -42$). Finally, by instead removing one of the three four-bead clusters from the $N_H = 40$ solution, one finds the minimum E_{HP} for $N_H = 36$ ($E_{\text{HP}} = -37$).

The T_{30} and T_{50} problems can be analyzed in a similar way. In Table I, we summarize minimum E_{HP} values and the degeneracy of the solutions for all problems considered in our hybrid computations.

APPENDIX B: MINIMUM E_{HP} FOR THE TARGET STRUCTURES WITH $N \leq 20$

Table II shows the minimum E_{HP} and the degeneracy of the solution to the sequence optimization problem for the systems studied using pure QPU computations and Schrödinger simulations, all with $N \leq 20$. How many of the solutions that actually have the target structure as their unique ground state is indicated within parentheses. This number is known from exact enumerations for these system sizes [20]. Finally, Table II also shows the energy gap ΔE between the ground state and the first-excited state of the problem Hamiltonian. All target structures T_N with $N \leq 20$ can be found in Fig. 9.




FIG. 8. Contact maps showing all contacts present in the 30-, 50-, and 64-bead target structures T_{30} , T_{50} , and T_{64} (Fig. 1). Two beads are said to be in contact if they are nearest neighbors on the lattice but not along the chain.

TABLE I. Minimum E_{HP} and the degeneracy of the solution for all sequence optimization instances studied in Sec. III A, using the target structures T_{30} , T_{50} , and T_{64} and different compositions N_{H} .

Target structure	N_{H}	Minimum E_{HP}	Degeneracy
T_{30}	12	-11	14
	13	-12	14
	14	-14	1
	15	-15	1
	16	-15	18
T_{50}	17	-16	4
	29	-28	10
	30	-29	4
	31	-30	1
T_{64}	36	-37	3
	40	-41	1
	42	-42	7

APPENDIX C: NUMERICAL INTEGRATION OF THE TIME-DEPENDENT SCHRÖDINGER EQUATION

To numerically integrate the Schrödinger equation, we first split the time evolution into M steps with length ϵ ($t_f = M\epsilon$), yielding


$$\psi(t_f) = U(t_M, t_{M-1}) \cdots U(t_1, t_0) \psi(0), \quad (\text{C1})$$

where $\psi(t)$ is the wave function corresponding to the Hamiltonian in Eq. (3), $t_m = m\epsilon$ ($m = 0, \dots, M$), and

$$U(t_{m+1}, t_m) = \mathcal{T} \exp \left[-i \int_{t_m}^{t_{m+1}} dt H(t) \right] \quad (\text{C2})$$

 TABLE II. Minimum E_{HP} , the degeneracy of the solution, the energy gap ΔE , and the chain strength J_{cs} used, for all sequence optimization instances studied in Sec. III B, using the target structures T_{10} – T_{20} and different compositions N_{H} . Also included is the system used for verification in Appendix C (T_8). The number of the solutions that have the target structure as its unique ground state is given within parentheses. ΔE is the gap between the two lowest values of $E(s)$ [Eq. (2)], when using $\lambda = 1.1$.

Target structure	N_{H}	Minimum E_{HP}	Degeneracy	ΔE	J_{cs}
T_8	4	-3	1 (1)	1.0	
T_{10}	4	-4	1 (1)	1.1	2.25
T_{11}	5	-4	1 (1)	1.0	2.25
T_{12}	4	-4	1 (1)	1.1	2.25
T_{12}	6	-5	1 (1)	1.0	2.75
T_{13}	6	-4	18 (1)	0.1	2.75
T_{13}	8	-6	1 (1)	1.0	2.75
T_{14}	6	-5	5 (1)	0.1	3.00
T_{14}	8	-7	1 (1)	1.0	3.00
T_{15}	5	-4	6 (1)	0.1	3.25
T_{15}	6	-5	5 (1)	0.1	3.00
T_{16}	6	-6	1 (1)	0.1	3.00
T_{16}	7	-7	1 (1)	1.0	3.00
T_{16}	8	-7	10 (5)	0.1	3.25
T_{17}	6	-6	1 (1)	1.0	3.50
T_{17}	7	-6	14 (5)	0.1	3.50
T_{18}	8	-8	1 (1)	0.1	3.50
T_{18}	9	-9	1 (1)	1.0	3.75
T_{19}	8	-8	1 (1)	0.1	3.75
T_{19}	9	-9	1 (1)	1.0	4.00
T_{20}	8	-8	1 (1)	0.1	4.25
T_{20}	9	-9	1 (1)	1.0	4.00
T_{20}	11	-10	1 (1)	1.0	4.25


 FIG. 9. The target structures used for the sequence optimization in Sec. III B and the structure used for verification in Appendix C (T_8). The different compositions used can be found in Table II.




FIG. 10. Step size dependencies when using Eq. (C4) to integrate the Schrödinger equation of an eight-qubit system [T_8 (Fig. 9) with $N_H = 4$ (Table II)] for two choices of the annealing time, $t_f = 10$ and $t_f = 50$. (a) The probability P_g of finding the final system in the ground state of the problem Hamiltonian H_P plotted against the step size ϵ . The lines are drawn to guide the eye. (b) The ϵ dependence of the finite-difference ratio $\chi(\epsilon) = [P_g(\epsilon) - P_g(\epsilon/2)]/[P_g(2\epsilon) - P_g(\epsilon)]$, which approaches 2^{-k} for small ϵ if the integrator is k th-order accurate. The data for $\chi(\epsilon)$ fall close to 0.25 for small ϵ , as expected for the second-order integrator in Eq. (C4).

(with $\hbar = 1$). Assuming a linear t dependence of $a(t)$ and $b(t)$ (Sec. II E), a leading-order Magnus expansion [42] yields $\ln U(t_{m+1}, t_m) \approx -i\epsilon H_m + O(\epsilon^3)$, where $H_m = H(t_m + \epsilon/2)$. It follows that

$$\tilde{U}(t_{m+1}, t_m) = \exp(-i\epsilon H_m) \quad (C3)$$

provides a unitary, second-order accurate integrator (cubic local error). With this approximation, the evolution from time t_m to time t_{m+1} is governed by the constant Hamiltonian H_m . Still, due to the Hilbert-space dimensionality (2^N), the numerical implementation requires care. To this end, we replace $\tilde{U}(t_{m+1}, t_m)$ [Eq. (C3)] by the Crank-Nicolson integrator [43]

$$\tilde{U}(t_{m+1}, t_m) = (\mathbf{T}^\dagger)^{-1} \mathbf{T} \text{ where } \mathbf{T} = 1 - \frac{i\epsilon}{2} H_m, \quad (C4)$$

which is implicit but still feasible, thanks to the sparseness of H_m . Like $\tilde{U}(t_{m+1}, t_m)$, $\tilde{U}(t_{m+1}, t_m)$ is unitary and second-order accurate. The sparseness of H_m can be easily exploited by rewriting the relation $\psi(t_{m+1}) = \tilde{U}(t_{m+1}, t_m)\psi(t_m)$ as

$$\mathbf{A}\mathbf{v} = \mathbf{u} \text{ with } \mathbf{A} = \mathbf{T}\mathbf{T}^\dagger \text{ and } \mathbf{u} = \mathbf{T}^2\psi(t_m), \quad (C5)$$

and solving this linear system of equations for $\mathbf{v} = \psi(t_{m+1})$ by the conjugate gradient method [44]. In Eq. (C5), both sides of the equation were multiplied by $\mathbf{T}$, in order to have a Hermitian and positive-definite matrix $\mathbf{A}$, as required by the conjugate gradient method. We implemented this algorithm, based on the Crank-Nicolson and conjugate-gradient methods, into a C++ program. Figure 10 shows results from a test of the program on an eight-qubit system, using two choices of annealing time, $t_f = 10$ and $t_f = 50$, and different step sizes ϵ . From Fig. 10(a) it can be seen that the value $t_f = 50$ is sufficiently large for the ground-state probability P_g at $t = t_f$ to be close to one, which is not case for $t_f = 10$. Figure 10(b) shows data for a finite-difference ratio, $\chi(\epsilon)$, which confirm that, for

fixed t_f , P_g displays the expected quadratic dependence on ϵ for small ϵ .

APPENDIX D: PURE QUANTUM PROCESSING UNIT HIT RATES WITH AND WITHOUT CHAIN BREAKS

In our pure QPU computations, using the DWAVECLIQUE-SAMPLER, each logical qubit is represented by a chain of two or three physical qubits. It may happen that such chains break. The pure QPU hit rates shown in Fig. 4(a) represent averages over all generated annealing cycles, irrespective of whether all chains were intact. Figure 11 compares these hit rates with those obtained when omitting from the analysis all annealing cycles in which any chain break occurred (23% of the data). The results obtained with and without this filter are similar (Pearson correlation coefficient 0.94).




FIG. 11. Scatter plot comparing pure QPU hit rates computed using all data [Fig. 4(a)] to those obtained after filtering out all annealing cycles in which any chain break occurred.

- [1] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse Ising model, *Phys. Rev. E* **58**, 5355 (1998).
- [2] J. Brooke, D. Bitko, T. F. Rosenbaum, and G. Aeppli, Quantum annealing of a disordered magnet, *Science* **284**, 779 (1999).
- [3] 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).
- [4] 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).
- [5] J. J. Hopfield and D. W. Tank, Neural computation of decisions in optimization problems, *Biol. Cybern.* **52**, 141 (1985).
- [6] C. Peterson and J. R. Anderson, Neural networks and NP-complete optimization problems; A performance study on the graph bisection problem, *Complex Syst.* **2**, 59 (1988).
- [7] C. McGeoch and P. Farré, The D-Wave Advantage System: An Overview, Tech. Rep. 14-1049A-A (D-Wave Systems Inc., 2020).
- [8] F. Slongo, P. Hauke, P. Faccioli, and C. Micheletti, Quantum-inspired encoding enhances stochastic sampling of soft matter systems, *Sci. Adv.* **9**, eadi0204 (2023).
- [9] A. D. King, J. Raymond, T. Lanting, R. Harris, A. Zucca, F. Altomare, A. J. Berkley, K. Boothby, S. Ejtemaee, C. Enderud, E. Hoskinson, S. Huang, E. Ladizinsky, A. J. R. MacDonald, G. Marsden, R. Molavi, T. Oh, G. Poulin-Lamarre, M. Reis, C. Rich, Y. Sato, N. Tsai, M. Volkmann, J. D. Whittaker, J. Yao, A. W. Sandvik, and M. H. Amin, Quantum critical dynamics in a 5,000-qubit programmable spin glass, *Nature (London)* **617**, 61 (2023).
- [10] M. Vuffray, C. Coffrin, Y. A. Kharkov, and A. Y. Lokhov, Programmable quantum annealers as noisy Gibbs samplers, *PRX Quantum* **3**, 020317 (2022).
- [11] J. Nelson, M. Vuffray, A. Y. Lokhov, T. Albash, and C. Coffrin, High-quality thermal Gibbs sampling with quantum annealing hardware, *Phys. Rev. Appl.* **17**, 044046 (2022).
- [12] R. Sandt and R. Spatschek, Efficient low temperature Monte Carlo sampling using quantum annealing, *Sci. Rep.* **13**, 6754 (2023).
- [13] K. F. Lau and K. A. Dill, A lattice statistical mechanics model of the conformational and sequence spaces of proteins, *Macromolecules (Washington, DC, U.S.)* **22**, 3986 (1989).
- [14] A. Irbäck, L. Knuthson, S. Mohanty, and C. Peterson, Folding lattice proteins with quantum annealing, *Phys. Rev. Res.* **4**, 043013 (2022).
- [15] A. Perdomo-Ortiz, N. Dickson, M. Drew-Brook, G. Rose, and A. Aspuru-Guzik, Finding low-energy conformations of lattice protein models by quantum annealing, *Sci. Rep.* **2**, 571 (2012).
- [16] A. Robert, P. K. Barkoutsos, S. Woerner, and I. Tavernelli, Resource-efficient quantum algorithm for protein folding, *npj Quantum Inf.* **7**, 38 (2021).
- [17] B. Kuhlman, G. Dantas, G. C. Ireton, G. Varani, B. L. Stoddard, and D. Baker, Design of a novel globular protein fold with atomic-level accuracy, *Science* **302**, 1364 (2003).
- [18] G. Bhardwaj, V. K. Mulligan, C. D. Bahl, J. M. Gilmore, P. J. Harvey, O. Cheneval, G. W. Buchko, S. V. S. R. K. Pulavarti, Q. Kaas, A. Eletsky, P.-S. Huang, W. A. Johnsen, P. J. Greisen, G. J. Rocklin, Y. Song, T. W. Linsky, A. Watkins, S. A. Rettie, X. Xu, L. P. Carter, R. Bonneau, J. M. Olson, E. Coutsiyas, C. E. Correnti, T. Szyperski, D. J. Craik, and D. Baker, Accurate *de novo* design of hyperstable constrained peptides, *Nature (London)* **538**, 329 (2016).
- [19] K. K. Yang, Z. Wu, and F. H. Arnold, Machine-learning-guided directed evolution for protein engineering, *Nat. Methods* **16**, 687 (2019).
- [20] A. Irbäck and C. Troein, Enumerating designing sequences in the HP model, *J. Biol. Phys.* **28**, 1 (2002).
- [21] C. Holzgräfe, A. Irbäck, and C. Troein, Mutation-induced fold switching among lattice proteins, *J. Chem. Phys.* **135**, 195101 (2011).
- [22] V. K. Mulligan, H. Melo, H. I. Merritt, S. Slocum, B. D. Weitzner, A. M. Watkins, P. D. Renfrew, C. Pleasure, P. S. Arora, and R. Bonneau, Designing peptides on a quantum computer, *bioRxiv*, 752485 (2020).
- [23] M. H. Khatami, U. C. Mendes, N. Wiebe, and P. M. Kim, Gate-based quantum computing for protein design, *PLoS Comput. Biol.* **19**, e1011033 (2023).
- [24] A. Irbäck, C. Peterson, F. Potthast, and E. Sandelin, Design of sequences with good folding properties in coarse-grained protein models, *Structure (Oxford, U.K.)* **7**, 347 (1999).
- [25] A. Aina and S. Wallin, Multisequence algorithm for coarse-grained biomolecular simulations: Exploring the sequence-structure relationship of proteins, *J. Chem. Phys.* **147**, 095102 (2017).
- [26] D. Nilsson and A. Irbäck, Finite-size scaling analysis of protein droplet formation, *Phys. Rev. E* **101**, 022413 (2020).
- [27] A. Statt, H. Casademunt, C. P. Brangwynne, and A. Z. Panagiotopoulos, Model for disordered proteins with strongly sequence-dependent liquid phase behavior, *J. Chem. Phys.* **152**, 075101 (2020).
- [28] E. Bornberg-Bauer and H. S. Chan, Modeling evolutionary landscapes: Mutational stability, topology, and superfunnels in sequence space, *Proc. Natl. Acad. Sci. USA* **96**, 10689 (1999).
- [29] J. Aguirre, P. Catalán, J. A. Cuesta, and S. Manrubia, On the networked architecture of genotype spaces and its critical effects on molecular evolution, *Open Biol.* **8**, 180069 (2018).
- [30] C. McGeoch, P. Farré, and W. Bernoudy, D-Wave Hybrid Solver Service + Advantage: Technology Update, Tech. Rep. 14-1048A-A (D-Wave Systems Inc., 2020).
- [31] D-Wave system documentation, <https://docs.dwavesys.com/docs/latest/> (2023).
- [32] I. Hen and F. M. Spedalieri, Quantum annealing for constrained optimization, *Phys. Rev. Appl.* **5**, 034007 (2016).
- [33] R. Unger and J. Moult, Genetic algorithms for protein folding simulations, *J. Mol. Biol.* **231**, 75 (1993).
- [34] U. Bastolla, H. Frauenkron, E. Gerstner, P. Grassberger, and W. Nadler, Testing a new Monte Carlo algorithm for protein folding, *Proteins: Struct., Funct., Genet.* **32**, 52 (1998).
- [35] T. Albash, V. Martin-Mayor, and I. Hen, Analog errors in Ising machines, *Quantum Sci. Technol.* **4**, 02LT03 (2019).
- [36] A. Pearson, A. Mishra, I. Hen, and D. A. Lidar, Analog errors in quantum annealing: Doom and hope, *npj Quantum Inf.* **5**, 107 (2019).
- [37] D-Wave system documentation, https://docs.ocean.dwavesys.com/en/latest/examples/pp_greedy.html.
- [38] K. Chern, K. Boothby, J. Raymond, P. Farré, and A. D. King, Tutorial: Calibration refinement in quantum annealing, *Front. Comput. Sci.* **5**, 1238988 (2023).

- [39] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko, A. Bridgland, C. Meyer, S. A. A. Kohl, A. J. Ballard, A. Cowie, B. Romera-Paredes, S. Nikolov, R. Jain, J. Adler, T. Back, S. Petersen, D. Reiman, E. Clancy, M. Zielinski, M. Steinegger, M. Pacholska, T. Berghammer, S. Bodenstein, D. Silver, O. Vinyals, A. W. Senior, K. Kavukcuoglu, P. Kohli, and D. Hassabis, Highly accurate protein structure prediction with AlphaFold, *Nature (London)* **596**, 583 (2021).
- [40] M. Varadi, S. Anyango, M. Deshpande, S. Nair, C. Natassia, G. Yordanova, D. Yuan, O. Stroe, G. Wood, A. Laydon, A. Žídek, T. Green, K. Tunyasuvunakool, S. Petersen, J. Jumper, E. Clancy, R. Green, A. Vora, M. Lutfi, M. Figurnov, A. Cowie, N. Hobbs, P. Kohli, G. Kleywegt, E. Birney, D. Hassabis, and S. Velankar, AlphaFold protein structure database: Massively expanding the structural coverage of protein-sequence space with high-accuracy models, *Nucleic Acids Res.* **50**, D439 (2021).
- [41] <https://www.fz-juelich.de/ias/jsc>.
- [42] W. Magnus, On the exponential solution of differential equations for a linear operator, *Commun. Pure Appl. Math.* **7**, 649 (1954).
- [43] J. Crank and P. Nicolson, A practical method for numerical evaluation of solutions of partial differential equations of the heat-conduction type, *Math. Proc. Cambridge Philos. Soc.* **43**, 50 (1947).
- [44] M. R. Hestenes and E. Stiefel, Methods of conjugate gradients for solving linear systems, *J. Res. Natl. Bur. Stand.* **49**, 409 (1952).