

Folding lattice proteins with quantum annealing

Anders Irbäck^{1,*}, Lucas Knuthson¹, Sandipan Mohanty², and Carsten Peterson¹¹*Computational Biology & Biological Physics, Department of Astronomy, and Theoretical Physics, Lund University, Box 43, SE-221 00 Lund, Sweden*²*Institute for Advanced Simulation, Jülich Supercomputing Centre, Forschungszentrum Jülich, D-52425 Jülich, Germany*

(Received 6 May 2022; accepted 29 August 2022; published 10 October 2022)

Quantum annealing is a promising approach for obtaining good approximate solutions to difficult optimization problems. Folding a protein sequence into its minimum-energy structure represents such a problem. For testing new algorithms and technologies for this task, the minimal lattice-based [hydrophobic (H) or polar (P) beads] HP model is well suited, as it represents a considerable challenge despite its simplicity. The HP model has favorable interactions between adjacent, not directly bound hydrophobic residues. Here, we develop a novel spin representation for lattice protein folding tailored for quantum annealing. With a distributed encoding onto the lattice, it differs from earlier attempts to fold lattice proteins on quantum annealers, which were based upon chain growth techniques. With our encoding, the Hamiltonian by design has the quadratic structure required for calculations on an Ising-type annealer, without having to introduce any auxiliary spin variables. This property greatly facilitates the study of long chains. The approach is robust to changes in the parameters required to constrain the spin system to chainlike configurations, and performs very well in terms of solution quality. The results are evaluated against existing exact results for HP chains with up to $N = 30$ beads with 100% hit rate, thereby also outperforming classical simulated annealing. In addition, the method allows us to recover the lowest known energies for $N = 48$ and $N = 64$ HP chains, with similar hit rates. These results are obtained by the commonly used hybrid quantum-classical approach. For pure quantum annealing, our method successfully folds an $N = 14$ HP chain. The calculations were performed on a D-Wave Advantage quantum annealer.

DOI: [10.1103/PhysRevResearch.4.043013](https://doi.org/10.1103/PhysRevResearch.4.043013)

I. INTRODUCTION

Quantum computers with their interacting qubits as basic units appear very promising for optimization problems with binary variables such as those present in spin systems. These technologies are being developed along two main tracks: quantum annealers [1] and gate-based systems [2]. In quantum annealing (QA) [3–5], the idea is to encode the solution to a given optimization problem in the ground state of a Hamiltonian and efficiently locate the energy minimum by exploiting quantum fluctuations and tunneling, in analogy with the role played by thermal fluctuations in classical simulated annealing (SA) [6]. Mapping difficult binary optimization problems onto Ising spin glass systems goes back to the 1980s in the context of neural networks [7,8]. For a recent review, see Ref. [9]. In the context of QA, this approach is called quadratic unconstrained binary optimization (QUBO).

Protein folding, going from sequence to structure by minimizing an energy function, represents a difficult optimization problem. Simplified lattice-based models for this problem can often provide qualitatively relevant results, but remain

computationally challenging and are therefore ideal testbeds for novel algorithms. A pioneering QUBO formulation of the folding problem for lattice proteins was given by Perdomo and coworkers [10]. They considered the HP model [11], where proteins are represented by linear chains of N hydrophobic (H) or polar (P) beads, residing on a lattice. Their model used binary variables encoding bead coordinates on the lattice, and an additional set of auxiliary binary variables had to be added in order to obtain a quadratic Hamiltonian.

An early attempt to fold a short lattice protein ($N = 6$) by QA was carried out on a D-Wave (D-Wave Systems Inc.) machine [12]. This implementation relied on a growth algorithm, where turns along the chain were mapped onto qubits. Recent work implemented on an IBM gate-based quantum computer used a similar encoding [13], again for a short protein chain ($N = 7$). The growth algorithms offer a compact, resource-efficient representation of the structure of a chain, but creating a quadratic Hamiltonian requires additional spin variables and implementing interactions such as self-avoidance becomes tedious for long chains with this type of encoding. For a recent review of lattice protein folding on quantum computers, see Ref. [14].

In contrast to Refs. [12,13], here we propose a different binary encoding for lattice proteins. In our model, all sites of the lattice host qubits. The approach was inspired by recent D-Wave applications for homopolymers [15], and shares some similarities with a QUBO formulation for lattice heteropolymers [16] which, to our knowledge, has not yet been implemented. The energy E is such that its minimization

*anders.irback@thep.lu.se

makes the values of the qubits coalesce to a finite set of active qubits defining the desired folded structure. Importantly from the viewpoint of QA, the entire function E , including a self-avoidance term, is manifestly quadratic, or two-local, in the binary spin variables, without having to add any auxiliary spins. As a result, E retains a convenient form for long chains. This fully distributed dynamical encoding method can be considered as a clustering approach driven by requiring a legal chain on the lattice. It is also somewhat reminiscent of a molecular field theory [17] where the fields reside on a lattice and give rise to structures through their dynamics.

We evaluate the performance of our approach using the two-dimensional HP model [11] as a testbed. Although this model at first sight looks simple, finding its lowest energy has been shown to belong to the class of NP-complete problems [18]. In particular, we consider a set of sequences with up to 30 beads, for which the exact solutions are known from exhaustive enumerations of all structures [19,20]. In addition, we present results obtained for two longer sequences with 48 and 64 beads, respectively, which have been studied by various classical methods [21–23].

In what follows, we first briefly describe the HP model and then map the energy minimization problem for HP sequences to a QUBO problem. We then evaluate our encoding using both classical SA and explorations performed on the D-Wave Advantage quantum annealer, with over 5000 qubits and 15-way qubit connectivity [24].

II. METHODS

A. HP lattice proteins

We consider the minimal lattice-based HP model for protein folding [11], in which the protein is represented by a self-avoiding chain of hydrophobic (H) or polar (P) beads on a lattice. Two beads are said to be in contact if they are nearest neighbors on the lattice, but not along the chain. A given chain configuration is assigned an energy defined as $E_{\text{HP}} = -N_{\text{HH}}$, where N_{HH} is the number of HH contacts [11]. With this choice of energy, low-energy configurations tend to exhibit a hydrophobic core of H beads. Despite the simplicity of the model, there are HP sequences with a unique ground state, and therefore a well-defined structure. In particular, on a two-dimensional square lattice, it is known from exhaustive enumerations that about 2% of all HP sequences with ≤ 30 beads have a unique ground state [19,20].

B. Binary quadratic model for HP lattice proteins – QUBO encoding

Given an HP sequence $(h_1, \dots, h_N)$, $h_i \in \{P, H\}$, we wish to determine its ground state using QA. To this end, in this section, we present a binary encoding for HP lattice proteins, assuming a square grid with L^2 sites.

Inspired by the binary representation of homopolymers of Ref. [15], rather than directly encoding chain configurations, we introduce fields of binary variables along with penalty terms. These terms serve to ensure that the final binary field configurations correspond to proper chain configurations. To reduce the number of binary variables, we make a checkerboard division of the lattice into even and odd sites, and use the fact that in a valid chain configuration all even (odd) beads

share the same lattice site parity (see Fig. 1). As a result, we may assume that even (odd) beads reside on even (odd) lattice sites. Thus, we introduce one set of binary fields, σ_s^f , to describe the location of even beads, and another set for odd beads $\sigma_{s'}^{f'}$. Here, the indices f and s run over even beads and sites, respectively, while f' and s' run over odd beads and sites. We set $\sigma_s^f = 1$ if bead f is located on site s , and $\sigma_s^f = 0$ otherwise. The odd fields $\sigma_{s'}^{f'}$ are defined in the same way. The division into even and odd sites reduces the number of variables required from $N \times L^2$ to $\approx N \times L^2/2$.

Having defined the degrees of freedom, we now describe the energy function. In our QUBO model, the total energy E has the form

$$E = E_{\text{HP}} + \sum_{i=1}^3 \lambda_i E_i, \quad (1)$$

where E_{HP} is the energy of the HP model (see above) and the remaining three terms E_1 , E_2 and E_3 are constraint energies. The strengths of the constraints are set by the parameters λ_i .

Specifically, in terms of the binary fields, the four energies can be expressed as follows:

(i) The HP energy $E_{\text{HP}} = -N_{\text{HH}}$ can be rewritten as

$$E_{\text{HP}} = - \sum_{|f-f'|>1} C(h_f, h_{f'}) \sum_{\langle s, s' \rangle} \sigma_s^f \sigma_{s'}^{f'}, \quad (2)$$

where the interaction strength $C(h_f, h_{f'}) = 1$ if $h_f = h_{f'} = H$ and $C(h_f, h_{f'}) = 0$ otherwise. In Eq. (2), the second sum runs over all nearest-neighbor pairs of sites, $\langle s, s' \rangle$. Such a pair always consist of one even and one odd site. The beads f and f' must both be of type H for a nonzero energy contribution, and must not, with our definition of a contact, be adjacent along the chain.

(ii) The first constraint energy, E_1 , is given by

$$E_1 = \sum_f \left(\sum_s \sigma_s^f - 1 \right)^2 + \{\text{same for odd parity}\}, \quad (3)$$

and serves to ensure that each bead is located at exactly one lattice site.

(iii) The energy E_2 makes the chain self-avoiding. It is given by

$$E_2 = \frac{1}{2} \sum_{f_1 \neq f_2} \sum_s \sigma_s^{f_1} \sigma_s^{f_2} + \{\text{same for odd parity}\}, \quad (4)$$

and provides an energy penalty whenever two beads occupy the same site.

(iv) The final energy, E_3 , has the form

$$E_3 = \sum_{1 \leq f < N} \sum_s \sigma_s^f \sum_{\|s'-s\|=1} \sigma_{s'}^{f+1} + \{\text{same with odd/even parity interchanged}\}, \quad (5)$$

and is responsible for connecting the beads to a chain. It provides an energy penalty whenever two adjacent beads along chain are not nearest neighbors on the lattice.

Our model contains three parameters; λ_1 , λ_2 , and λ_3 [Eq. (1)]. It is desirable that when executing the model it is reasonably robust with respect to these parameters. This will turn out to be the case in Sec. III when exploring the method.




FIG. 1. Illustration of a hypothetical evolution of the binary model described in Sec. II B for the 6-bead sequence HPHPPH. Circles represent beads, and numbers indicate bead positions along the sequence. By construction, odd/even beads can reside only on grey/white sites. (a) Early stage. Typically, all the three constraints are violated ($E_1, E_2, E_3 > 0$). (b) Intermediate stage. Some but not all of the constraints are satisfied (in this example: $E_1 = E_2 = 0, E_3 > 0$). (c) The final state, in this example corresponding to the desired minimum-energy structure of the given sequence ($E_{\text{HP}} = -2, E_1 = E_2 = E_3 = 0$).

As indicated in Sec. I, the above binary model shares similarities with the “diamond” encoding proposed in Ref. [16]. The latter method is able to reduce the number of binary variables required for very short chains, by fixing the position of the first bead and using the fact that odd and even beads can be assumed to belong to different “diamond” layers. For long chains, our choice of a freely moving chain on a simple odd/even checkerboard is more resource-efficient, because, in general, the search for the ground state can be carried out on a smaller grid if the chain is freely moving. Our constraint energies E_i also differ from those of Ref. [16]. In our case, all three constraint energies are manifestly non-negative for both physical and unphysical spin configurations, which makes our method more robust to changes in the strength parameters λ_i . Since the encoding in Ref. [16] was never explored, a robustness analysis is not available there.

C. Scaling properties

The binary model introduced above uses $\approx NL^2/2$ spins to describe the structure of an N -bead HP protein on an L^2 grid. In order that the (in general *a priori* unknown) minimum-energy structure fits inside the simulation grid, the lattice size L has to be chosen according to the protein sequence. A safe choice is to take $L = N$, in which case the number of spins scales cubically with N . However, a typical minimum-energy structure is compact, since it contains many attractive H-H contacts, and fits onto a grid with L not much larger than $\sqrt{N}$. Such a structure may therefore be found using N^2 rather than N^3 spins. Note that the assumption of a square grid is unnecessary and was only made for simplicity. For a general grid shape, the number of spins scales as N times the number of grid points.

Compared to our encoding, turn-based ones [12,16] provide a more economic description of the chain geometry, requiring only $\sim N$ spins. However, additional spins are needed in order to formulate the Hamiltonian. A resource-efficient turn-based binary model with a quadratic Hamiltonian was proposed in Ref. [16], which in total uses $\sim N^2$ spins. Here, the number of spins is thus comparable to that of our model when using a reduced grid size. However, in the turn-based model [16], the Hamiltonian is somewhat complicated already for $N = 6$, and it would be challenging to implement it for

the chain lengths studied here. In contrast, the Hamiltonian in Sec. II B retains its simple and convenient structure as the chain length is increased.

D. Simulated annealing

Before turning to QA, we tested this QUBO model using SA, with the system defined by the partition function $Z = \sum_{\{\sigma_s^f, \sigma_{s'}^{f'}\}} e^{-\beta E}$, where β denotes inverse temperature and E is given by Eq. (1). All runs spanned the same set of 25 temperatures, given by $\beta_0 = 1$ and $\beta_{i+1} = 1.05\beta_i$. At each temperature, 10^4 sweeps were performed, where one sweep comprises, on average, one attempted update per spin variable. The updates were single-spin flips, controlled by a Metropolis acceptance criterion. All runs were started from random initial spin configurations, and used a 10^2 grid.

For comparison, we also conducted SA runs based on the conventional explicit-chain representation of the HP model. Here, the energy was given by E_{HP} , without the constraint terms. The set of temperatures was the same as in the QUBO SA runs. The simulations used three standard Metropolis-type elementary moves: local one- and two-bead updates, and a nonlocal pivot update. At each temperature, 10^5 sweeps were performed, with one sweep consisting of $N - 1$ one-bead moves, $N - 2$ two-bead moves and one pivot move. The chains were not confined to a finite-size grid.

All SA simulations were run on a standard desktop computer. For $N = 30$, each QUBO SA run required 21 CPU-core seconds, whereas each explicit-chain SA run required 5 CPU-core seconds. To gather statistics, for each sequence, we performed 1000 runs with each method, using different random number seeds.

E. Hybrid quantum-classical computations

D-Wave offers access to solvers based solely on QA as well as hybrid quantum-classical solvers. The hybrid approach uses classical solvers while sending suitable subproblems as queries to the quantum processing unit (QPU). The solutions to the subproblems serve to guide the classical solvers [25]. The goal is to speed up the solution of challenging QUBO problems by queries to the QPU. With the hybrid approach, it is possible to tackle much larger problems, with many

thousands of fully connected variables, than what can be dealt with using QA alone.

We conducted hybrid quantum-classical computations for HP chains with up to 64 beads, using D-Wave’s hybrid solver services and a D-Wave Advantage quantum computer. All sequences with $N \leq 30$ were folded on a 10^2 grid using the default run time set by the hybrid solver, which depends on problem size and was 4 s for $N = 30$. To gather statistics, a set of 100 runs were performed for each sequence. Two additional sequences were studied, with $N = 48$ and $N = 64$, respectively, using both 10^2 and 15^2 grids. For these sequences, the run time for the hybrid solver had to be taken larger than the default run time, to ensure satisfactory hit rates. For a given sequence and grid, computations were performed for a number of different run times, to investigate the run time dependence of the hit rate. For each choice of sequence, grid and run time, 10 runs were performed.

F. Pure QPU computations

The 5000-qubit Advantage machine uses a Pegasus topology with a connectivity of 15 [24]. Thus, in order to solve a problem with higher connectivity, it has to be embedded into the Pegasus graph. This embedding is done by forming “chains” of qubits that 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 chain breaks.

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. All the computations used a chain strength between 1 of 7.5 and the annealing time was set to $\tau = 2000 \mu\text{s}$, its maximum allowed value. The number of output reads (annealing cycles) per run, which must be $< 10^6/(\tau/\mu\text{s})$, was set to 490. For each system studied, 100 runs were executed, each with this number of annealing cycles.

G. Testbed – HP sequences

As a testbed, we use a selected set of HP sequences with 4–30 beads, all of which are known from exhaustive enumerations to have a unique minimum-energy structure [19,20]. The sequences are labeled S_N , where N indicates the number of beads.

A sequence having a unique minimum-energy structure is said to design that structure. The number of different sequences designing a given structure is called the designability of the structure. Structures with high designability thus show robustness to mutation.

For a given $N \leq 30$, the selected sequence S_N is one of those that fold to the structure with highest designability for that N .

In addition, we included in our study two sequences with $N = 48$ and $N = 64$, respectively, for which low-energy structures have been explored by various classical methods.

The sequences S_N along with illustrations of low-energy structures can be found in Appendix A.




FIG. 2. Run-time evolution of the HP energy E_{HP} and the constraint energies E_1 , E_2 , and E_3 in a QUBO SA folding simulation for the 30-bead sequence S_{30} (Appendix A) on 10^2 grid, with $\lambda = (2.1, 2.4, 3.0)$.

III. RESULTS

Using the spin representation of Sec. II B, we wish to find minimum-energy structures of given HP sequences by minimizing the total energy $E = E_{\text{HP}} + \sum_i \lambda_i E_i$ [Eq. (1)] on a quantum annealer. As a first step toward this goal, we investigate the power of the QUBO approach under classical SA, and how it depends on the Lagrange parameters λ_i . We next do the same using the hybrid quantum-classical solver. Finally, we compare the results by using the QPU annealer only. We find that the hybrid quantum-classical solver outperforms all the other approaches listed above for our application, in terms of ease of achieving 100% hit rate and consumed computer time.

A. Simulated annealing with QUBO encoding

In our binary model, the E_{HP} energy can become substantially lower than it is in any proper chain conformations. For this QUBO approach to work, it is therefore essential that the λ_i parameters that force the solutions to be “legal” are sufficiently large. On the other hand, by choosing large λ_i values, one risks making the energy landscape rugged, and therefore the dynamics potentially slow. Hence, the λ_i parameters should be neither too large nor too small.

To gain insight into the behavior of the binary model and its dependence on the λ_i parameters, we conducted a set of classical Monte-Carlo-based SA runs (Sec. II D), using the HP sequences S_{18} – S_{30} in Appendix A. The runs had a fixed length and were deemed successful if the final state corresponded to the known minimum-energy structure of the given HP sequence. As expected, in order to have an acceptable hit rate, it was necessary to choose the λ_i parameters with some care. Nevertheless, without excessive fine-tuning, it was possible to find a single set of parameters, $\lambda = (2.1, 2.4, 3.0)$, that gave a hit rate $\gtrsim 0.1$ for all the sequences S_{18} – S_{30} (see below). We refrained from attempting any further optimization of the parameters, as the optimal values need not be the same on a quantum annealer. The optimal parameters would, of course, also depend on both HP sequence and grid size.

Figure 2 shows the run-time evolution of the four different energy terms in one of 1000 QUBO SA runs for the sequence S_{30} . At the end of the run, all the three constraint energies




FIG. 3. Parameter dependence of the fraction of correct solutions (hit rate) in the vicinity of a reference point λ^* , when using QUBO SA and hybrid quantum-classical computation to search for the ground state of the S_{30} sequence (Appendix A) on a 10^2 grid. The hit rate is plotted against $\Delta\lambda_i = \lambda_i - \lambda_i^*$, keeping $\lambda_j = \lambda_j^*$ for $j \neq i$. Lines are drawn to guide the eye. (a) QUBO SA with $\lambda = (2.1, 2.4, 3.0)$. (b) Hybrid quantum-classical computations with $\lambda = (2.0, 3.0, 3.0)$. Note the difference in scale between the two panels, reflecting the difference in performance as shown in Fig. 4.

E_i vanish, while E_{HP} takes its known minimum value for an HP chain with this sequence ($E_{\min} = -15$). Hence, the final spin configuration corresponds to the S_{30} ground state in the HP model. The hit rate, defined as the fraction of runs ending in the ground state, was 0.226 ± 0.013 . The remaining runs ended in spin configurations that either did not correspond to a proper chain, or corresponded to a structure with $E_{HP} > E_{\min}$. In the beginning of the runs, the spin system undergoes a rapid relaxation, which brings the energies from initial values $E_{HP} \sim -10^3$ and $\lambda_1 E_1 + \lambda_2 E_2 + \lambda_3 E_3 \sim 10^5$ to the plotted range before the first measurement is taken (after 10^3 sweeps). We note that among the three constraint energies, E_1 and E_3 tend to relax much more slowly than E_2 , as is the case in Fig. 2. Note also that the HP energy takes values $E_{HP} < E_{\min}$ many times during the course of the run. Such values can occur only when at least one constraint is broken.

To elucidate the λ dependence of the hit rate, we also performed QUBO SA calculations for a set of additional λ near $\lambda = \lambda^* = (2.1, 2.4, 3.0)$ using the sequence S_{30} . Figure 3(a) shows the observed hit rates when changing one λ_i at a time. In all three λ_i , the hit rate stays tiny until a threshold is passed, followed by a steep increase to the maximum observed hit rate, for $\lambda_i = \lambda_i^*$. When further increasing λ_i beyond λ_i^* , the hit rate decays, probably due to an increasingly rugged energy landscape. This decay leads to an upper limit on the parameter λ_1 , beyond which the hit rate is impractically small. By contrast, the hit rate stays significant even for λ_2 and λ_3 values much larger than those in Fig. 3(a). In fact, setting $\lambda_2 = 100$ or $\lambda_3 = 100$, we still obtained hit rates of 0.132 ± 0.011 and 0.074 ± 0.008 , respectively. Hence, overall the parameter sensitivity is low, although λ_1 must be chosen with some care.

The fact that the λ_1 dependence has a different shape than the dependences on λ_2 and λ_3 can be, at least in part, understood. With the single-spin updates employed, the system cannot move from one chainlike configuration to another, both with $E_1 = E_2 = E_3 = 0$, without visiting intermediate nonchain configurations with $E_1 > 0$. By contrast, E_2 and E_3 may stay zero during such a move. These observations suggest that the energy landscape indeed is rugged for large λ_1 , but not necessarily so for large λ_2 or λ_3 .

To explore how the hit rate of the QUBO SA approach depends on chain length, we conducted calculations for all the HP sequences S_{18} – S_{30} in Appendix A, using $\lambda = \lambda^*$. As expected, the measured hit rates show a decreasing trend with increasing N (Fig. 4). However, the decrease is not monotonous, indicating that the hit rate is sequence-dependent and not a simple function of N .

For comparison, we also carried out a set of direct SA minimizations of E_{HP} based on conventional explicit-chain Monte Carlo methods (Fig. 4). Despite being faster, the hit rate is higher in these runs than it is with QUBO-based SA. However, the difference in hit rate is modest given that state space for explicit chains is comparatively tiny. Note the similarities in shape between the hit rates obtained from these two unrelated sets of SA calculations. These similarities suggest




FIG. 4. Fraction of correct solutions (hit rate) when searching for the ground state of HP chains with $18 \leq N \leq 30$ beads by using the D-Wave hybrid solver with QUBO encoding (blue), SA with explicit chains (black), and SA with QUBO encoding (red). The HP sequences studied can be found in Appendix A. The parameters λ were set to $(2.1, 2.4, 3.0)$ in the QUBO SA runs, and to $(2.0, 3.0, 3.0)$ when using the hybrid solver. All QUBO-based results were obtained using a 10^2 grid. The hybrid computations used the default run time for the solver (4 s for $N = 30$).




FIG. 5. Hybrid quantum-classical computations for the sequences S_{48} and S_{64} (Appendix A), using $\lambda = (2.0, 3.5, 3.0)$ for S_{48} and $\lambda = (3.0, 4.0, 4.0)$ for S_{64} . (a) Fraction of runs that found the lowest known energy, $E_{\min}$, in hybrid computations on 10^2 and 15^2 grids for the sequences S_{48} ($E_{\min} = -23$ [23]) and S_{64} ($E_{\min} = -42$ [22]), plotted against run time for the hybrid solver (logarithmic scale). The horizontal line indicates a hit rate of 0.5. (b) Example of an S_{64} structure with the lowest known energy ($E_{\min} = -42$) that was obtained in the hybrid computations. Filled and open symbols indicate H and P beads, respectively.

that some target structures are relatively easy or difficult to find, independent of the method employed.

B. Hybrid quantum-classical computations

A promising alternative to pure QA is provided by hybrid quantum-classical methods, by which larger systems can be studied. To assess the power of this approach, we conducted hybrid computations for all the HP sequences studied in Sec. III A, S_{18} – S_{30} (Appendix A), using the default run time for the hybrid solver. We additionally included a two longer sequences [21], which have been extensively used as testbeds for various (classical) methods. For these sequences, the dependence of the hit rate on run time was explored.

As in the SA case, with the hybrid solver, a rough search was sufficient in order to find a single λ , $\lambda^* = (2.0, 3.0, 3.0)$, for which all the sequences S_{18} – S_{30} could be correctly folded on a 10^2 grid. Figure 3(b) shows the parameter dependence of the hit rate near λ^* when using the hybrid solver. Compared to QUBO SA [Fig. 3(a)], the measured hit rates are markedly higher with the hybrid solver (Fig. 4). At the same time, the λ_i dependences share a similar shape in both cases. In particular, in both cases, the hit rate is more sensitive to changes in λ_1 than to changes in λ_2 or λ_3 . As in the SA case, λ_2 or λ_3 can be chosen far above the plotted range in Fig. 3, without any major loss in hit rate. In fact, when setting $\lambda_2 = 100$ or $\lambda_3 = 100$, we obtained hit rates of 0.980 ± 0.014 and 0.54 ± 0.05 , respectively. Overall, the parameter sensitivity is lower with the hybrid solver than with QUBO SA.

When comparing hit rates from our hybrid and QUBO SA runs for the sequences S_{18} – S_{30} , we find that it is consistently highest in the hybrid case (Fig. 4). In fact, the hit rate is one across this entire set of sequences for the hybrid solver, despite using the shortest run time that can be set for the solver.

It is important to note that when folding the sequences S_{18} – S_{30} , the hybrid solver did not always make use of the QPU. The fraction of runs that used the QPU increased with N and was above one half for $N > 21$. Still, the precise contribution of the QPU to the final results is hard to judge since the details of the hybrid solver are not publicly available information. Nevertheless, the ease with which these sequences could be

folded motivated us to also test the hybrid solver on two significantly longer sequences, namely S_{48} and S_{64} (Appendix A) with 48 and 64 beads, respectively.

For these two sequences exact results are not available, but both belong to a set of HP sequences that have been widely used to test novel (classical) algorithms [21]. The lowest known energies are $E_{\text{HP}} = -23$ for S_{48} [23] and $E_{\text{HP}} = -42$ for S_{64} [22]. In order to obtain good results, the λ_i parameters had to be adjusted for these larger chains to $\lambda = (2.0, 3.5, 3.0)$ for S_{48} and $\lambda = (3.0, 4.0, 4.0)$ for S_{64} . In addition, using the default run time for the hybrid solver (6 s for S_{48} and 8 s for S_{64}), as was done for S_{18} – S_{30} , turned out to give unsatisfactorily low hit rates for S_{48} and S_{64} . Therefore, we investigated the run time dependence of the hit rate. To this end, we carried out computations with several different run times for both S_{48} and S_{64} , using both 10^2 and 15^2 grids.

Figure 5(a) summarizes the run time dependence of the hit rate as observed in these runs, for both sequences and grids. For a given sequence and grid, a steep increase in hit rate can be seen once the run time passes a threshold. While a hit rate of unity is reachable for the systems in Fig. 5(a), a natural measure of performance is provided by the run time required to obtain a hit rate of 0.5, denoted by $t_{1/2}$. For a given grid, the value of $t_{1/2}$ is roughly 10 times larger for S_{64} than for S_{48} . An increase in $t_{1/2}$ by roughly a factor 10 is also observed when increasing the grid size from 10^2 to 15^2 for a given sequence. It would be interesting to further explore the dependence of $t_{1/2}$ on chain length and grid size. However, a systematic study of this problem is beyond the scope of the present paper, especially since such a study should also address the sequence dependence of $t_{1/2}$. For the sequences S_{18} – S_{30} studied earlier, the default run time is apparently above the threshold, as the hit rate is 1.0.

An example of an S_{64} structure with the lowest known energy that was obtained in our hybrid computations can be found in Fig. 5(b).

C. Pure QPU computations

The QUBO problem that we wish to solve for finding minimum-energy HP structures contains $\approx NL^2/2$ logical




FIG. 6. Pure QPU computations. For every choice of sequence and grid size studied, we conducted 100 runs with 490 annealing cycles each. The sequences can be found in Appendix A. (a) The hit rate on a logarithmic scale against the number of physical qubits used for the sequences S_4 , S_6 , S_7 , S_8 , S_9 , and S_{10} for various grid sizes, using $\lambda = (1.0, 2.0, 1.5)$. The grids used can be found in Appendix B. Statistical errors are in many cases comparable with or smaller than the symbol sizes. (b) The minimum-energy structure for the sequence S_{14} , which was successfully recovered on a 4^2 grid, using $\lambda = (2.0, 7.0, 4.0)$. Filled and open symbols indicate H and P beads, respectively.

qubits. Moreover, the system is almost fully connected, implying that its embedding into the QPU topology requires a significant amount of additional qubits. Therefore, pure QPU computation is effectively limited to relatively short HP chains.

To explore how the performance of the pure QPU approach depends on system size, we conducted computations for the six sequences S_4 , S_6 , S_7 , S_8 , S_9 and S_{10} (Appendix A), with each sequence run on several different grids. The grids used, which include nonsquare rectangular ones, can be found in Appendix B, along with the numbers of logical and physical qubits involved. The number of physical qubits grows to a good approximation quadratically with the number of logical qubits, as illustrated in Appendix B.

Figure 6(a) shows the fraction of all annealing cycles that recovered the known minimum-energy structure, for these systems, plotted against the number of physical qubits employed. The parameters λ_i and the annealing time were the same for all systems, whereas the chain strength was chosen individually for each system, for best performance (among the values 1.0, 1.5, ..., 4.5, 5.0). Albeit with some scatter, the hit rate shows a roughly exponential decay with the number of physical qubits used. In part, this deterioration of performance with increasing system size may stem from integrated Hamiltonian control errors, which will reduce the probability of finding the ground state of the intended Hamiltonian [26]. Remedies to this problem are being explored [26].

In Fig. 6(a), the hit rate falls off most rapidly for the shortest chain studied, with $N = 4$. This behavior is likely an artifact due to suboptimal λ_i parameters, which, for simplicity, are kept the same for all systems. In fact, for the largest $N = 4$ system with 192 physical qubits, a significantly higher hit rate (≈ 0.02) was observed when changing λ_1 from its value 1.0 in Fig. 6(a) to 1.5.

The longest sequence that was successfully folded in our pure QPU computations was S_{14} with 14 beads (Appendix A). This sequence, whose minimum-energy structure can be seen

in Fig. 6(b), was studied using a 4^2 grid, which required 112 logical and 1214 physical qubits. The chain strength was set to 7.5. To our knowledge, this is the largest protein successfully folded using a quantum computer. However, the ground state was only recovered in one of a total of 100×490 annealing cycles. This number of cycles is larger than the number of distinct conformations available to a chain with 14 beads on a 4^2 grid, which is 416. On the other hand, it is tiny compared to the $2^{112} \approx 5.2 \times 10^{33}$ states of the binary system, the vast majority of which do not correspond to proper chain configurations. Finally, we note that our S_{14} system is similar in size to the largest system solved in a recent benchmarking study of the D-Wave Advantage machine using exact cover problems, which contained 120 logical qubits [27].

It is possible that the pure QPU results can be improved by further tuning of the simulation parameters. However, at present, we conclude that pure QPU computation cannot match classical SA or the hybrid quantum-classical approach (Secs. III A and III B).

IV. SUMMARY AND OUTLOOK

We have developed a novel mapping of the two-dimensional HP lattice protein model onto a quantum annealer, which is successfully explored on the D-Wave Advantage system. This simplified model for protein structures is known to represent a difficult optimization problem when determining structure using classical methods. As benchmarks for success, we used HP chains with sizes ranging from $N = 4$ to $N = 30$, for which the exact solutions are known from exhaustive enumerations. For larger problems, with $N = 48$ and $N = 64$, we compared with the best known solutions obtained by classical means. The approach allows us to explore the largest chains studied so far on a quantum annealer, and consistently provides high percentages of correct solution on multiple runs. However, the success of our approach relies upon using the hybrid variant provided by D-Wave. The performance of pure QA is less impressive with a drastic decrease in success rate as the system size is increased. These calculations were therefore limited to chains with at most $N = 14$.

Our approach differs from previous attempts based upon growth algorithms, as it maps the problem onto a lattice spin system where the spins, or qubits, are present throughout the lattice. In comparison to earlier work, this representation greatly facilitates the handling of interactions, including self-avoidance. In particular, with this encoding, the energy function can be written in a quadratic form without adding any auxiliary spins.

The encoding requires penalty terms for the global energy minimum of the spin system to correspond to a proper chain. The method is robust to changes in the strengths of the penalty terms, which require only a modest amount of tuning.

For convenience, we have focused on the HP model with its minimal two-letter alphabet. To extend the approach to models with larger alphabets, such as the 20-letter Miyazawa-Jernigan model [28], is straightforward, as it amounts to simply changing the interaction parameters in Eq. (2). Moreover, although the checkerboard division into even and odd sites may have to be modified or abandoned, the method can be applied to an arbitrary graph. In particular, it can be directly




FIG. 7. Ground states for all the sequences S_N in Appendix A with $N \leq 30$ [19,20] except S_{14} [whose ground state can be found in Fig. 6(b)]. Also shown is an S_{48} structure, which is one of the structures with the lowest known energy for this sequence found using the hybrid solver. A low-energy structure for the sequence S_{64} can be found in Fig. 5(b).

APPENDIX B: DETAILS OF THE PURE QPU COMPUTATIONS

This appendix provides details of the pure QPU computations discussed in Fig. 6. Table II shows the grid sizes and

the numbers of logical and physical qubits used. Figure 8 illustrates the relation between the numbers of logical and physical qubits.

-
- [1] M. W. Johnson, M. H. S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk, E. M. Chapple, C. Enderud, J. P. Hilton, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Oh, I. Perminov, C. Rich, M. C. Thom *et al.*, Quantum annealing with manufactured spins, *Nature (London)* **473**, 194 (2011).
 - [2] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature (London)* **574**, 505 (2019).
 - [3] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse Ising model, *Phys. Rev. E* **58**, 5355 (1998).
 - [4] J. Brooke, D. Bitko, T. F. Rosenbaum, and G. Aeppli, Quantum annealing of a disordered magnet, *Science* **284**, 779 (1999).
 - [5] 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).
 - [6] S. Kirkpatrick Jr., C. D. Gelatt, and M. P. Vecchi, Optimization by simulated annealing, *Science* **220**, 671 (1983).
 - [7] J. J. Hopfield and D. W. Tank, Neural computation of decisions in optimization problems, *Biol. Cybern.* **52**, 141 (1985).
 - [8] C. Peterson and J. R. Anderson, Neural networks and NP-complete problems; a performance study of the graph bisectioning problem, *Complex Syst.* **2**, 59 (1988).
 - [9] A. Lucas, Ising formulations of many NP problems, *Front. Phys.* **2**, 1 (2014).
 - [10] A. Perdomo, C. Truncik, I. Tubert-Brohman, G. Rose, and A. Aspuru-Guzik, Construction of model Hamiltonians for adiabatic quantum computation and its application to finding low-energy conformations of lattice protein models, *Phys. Rev. A* **78**, 012320 (2008).
 - [11] K. F. Lau and K. A. Dill, A lattice statistical mechanics model of the conformational and sequence spaces of proteins, *Macromolecules* **22**, 3986 (1989).
 - [12] 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).
 - [13] A. Robert, P. K. Barkoutsos, S. Woerner, and I. Tavernelli, Resource-efficient quantum algorithm for protein folding, *npj Quantum Inf.* **7**, 38 (2021).
 - [14] C. Outeiral, G. M. Morris, J. Shi, M. Strahm, S. C. Benjamin, and C. M. Deane, Investigating the potential for a limited quantum speedup on protein lattice problems, *New J. Phys.* **23**, 103030 (2021).
 - [15] C. Micheletti, P. Hauke, and P. Faccioli, Polymer Physics by Quantum Computing, *Phys. Rev. Lett.* **127**, 080501 (2021).
 - [16] R. Babbush, A. Perdomo-Ortiz, B. O’Gorman, W. Macready, and A. Aspuru-Guzik, Construction of energy functions for lattice heteropolymer models: Efficient encodings for constraint satisfaction programming and quantum annealing, *Adv. Chem. Phys.* **155**, 201 (2014).
 - [17] B. A. Reva and A. V. Finkelstein, Search for the most stable folds of protein chains: II. Computation of stable architectures of β -proteins using a self-consistent molecular field theory, *Protein Eng. Des. Sel.* **9**, 399 (1996).
 - [18] P. Crescenzi, D. Goldman, C. Papadimitriou, A. Piccolboni, and M. Yannakakis, On the complexity of protein folding, *J. Comput. Biol.* **5**, 423 (1998).
 - [19] A. Irback and C. Troein, Enumerating designing sequences in the HP model, *J. Biol. Phys.* **28**, 1 (2002).
 - [20] C. Holzgräfe, A. Irback, and C. Troein, Mutation-induced fold switching among lattice proteins, *J. Chem. Phys.* **135**, 195101 (2011).
 - [21] R. Unger and J. Moult, Genetic algorithms for protein folding simulations, *J. Mol. Biol.* **231**, 75 (1993).
 - [22] U. Bastolla, H. Frauenkron, E. Gerstner, P. Grassberger, and W. Nadler, Testing a new Monte Carlo algorithm for protein folding, *Proteins* **32**, 52 (1998).
 - [23] F. Liang and W. H. Wong, Evolutionary Monte Carlo for protein folding simulations, *J. Chem. Phys.* **115**, 3374 (2001).
 - [24] C. McGeoch and P. Farré, *The D-Wave Advantage System: An overview*, Tech. Rep. No. 14-1049A-A (D-Wave Systems Inc., 2020).
 - [25] C. McGeoch, P. Farré, and W. Bernoudy, *D-Wave Hybrid Solver Service + Advantage: Technology update*, Tech. Rep. No. 14-1048A-A (D-Wave Systems Inc., 2020).
 - [26] A. Pearson, A. Mishra, I. Hen, and D. A. Lidar, Analog errors in quantum annealing: Doom and hope, *npj Quantum Inf.* **5**, 107 (2019).
 - [27] D. Willsch, M. Willsch, C. D. G. Calaza, F. Jin, H. De Raedt, M. Svensson, and K. Michielsen, Benchmarking Advantage and D-Wave 2000Q quantum annealers with exact cover problems, *Quantum Inf. Proc.* **21**, 141 (2022).
 - [28] S. Miyazawa and R. L. Jernigan, Residue-residue potentials with a favorable contact pair term and an unfavorable high packing density term, for simulation and threading, *J. Mol. Biol.* **256**, 623 (1996).
 - [29] <https://www.fz-juelich.de/ias/jsc>.