

# A QUBO model of the RNA folding problem optimized by variational hybrid quantum annealing

Tristan Zaborniak\*

Department of Computer Science  
University of Victoria  
Victoria, Canada  
tristanz@uvic.ca

Juan Giraldo\*

Department of Computer Science  
University of Victoria  
Victoria, Canada  
juangiraldo@uvic.ca

Hausi Müller

Department of Computer Science  
University of Victoria  
Victoria, Canada  
hausi@uvic.ca

Hosna Jabbari

Department of Computer Science  
University of Victoria  
Victoria, Canada  
jabbari@uvic.ca

Ulrike Stege

Department of Computer Science  
University of Victoria  
Victoria, Canada  
ustege@uvic.ca

**Abstract**—RNAs self-interact through hydrogen-bond base-pairing between nucleotides and fold into specific, stable structures that substantially govern their biochemical behaviour. Experimental characterization of these structures remains difficult, hence the desire to predict them computationally from sequence information. However, correctly predicting even the base pairs involved in the folded structure of an RNA, **known as secondary structure, from its sequence using minimum free energy models is understood to be NP-hard**. Classical approaches rely on heuristics or avoid considering pseudoknots in order to render this problem more tractable, with the cost of inexactness or excluding an entire class of important RNA structures. Given their prospective and demonstrable advantages in certain domains, including combinatorial optimization, quantum computing approaches by contrast have the potential to compute the full RNA folding problem while remaining more feasible and exact. Herein, we present a physically-motivated QUBO model of the RNA folding problem amenable to both quantum annealers and circuit-model quantum computers and compare the performance of this formulation versus current RNA folding QUBOs after tuning the parameters of all against known RNA structures using an approach we call “variational hybrid quantum annealing”.

**Index Terms**—RNA folding, quantum computing, machine learning, computational complexity, variational hybrid quantum annealing

## I. INTRODUCTION

RNAs are crucial to biological function, assisting in the (de)coding and regulation of genes, protein construction, and catalysis [1], [2]. Indeed, RNA forms the genetic material in some organisms (e.g., RNA viruses) and may have preceded DNA and proteins in evolutionary history [3], [4]. Integral to these roles is the ability and tendency of RNA to fold into specific structures through hydrogen-bond base-pairing with itself [5]. These structures dictate in combination with the primary sequence the behaviour of any particular RNA and its interactions with other biomolecules [6]–[8].

While recent advances in sequencing technology have made possible rapid accessioning of RNA primary sequences, experimental determination of RNA structure remains slow and expensive by comparison [9]–[13]. In addition to the desire to efficiently design structured RNAs, this provides incentive toward the end of predicting RNA structure computationally from **primary sequence information** [14]–[17].

However, the computational approach is not without its challenges. Given that the most stable structure of a molecule is with **minimum free energy** (MFE) at equilibrium, many RNA structure prediction algorithms attempt to determine the MFE structure of an RNA from primary sequence information [18]. This problem is understood to be NP-hard when allowing for all possible observed intricacies of RNA secondary structure (e.g., pseudoknots) [19]. Various classical approaches achieve speedups through either heuristic methods or introducing simplifications to the energy model for RNA folding (e.g., ignoring pseudoknots, or focusing on only certain classes thereof) [20]–[28]. Other studies to optimize the balance between RNA folding speed and realism are ongoing. [29]–[31].

Quantum computing approaches have the potential to contribute considerably to this objective, given their proven and anticipated abilities to dramatically reduce the time-complexity of certain problem cases compared to classical counterparts, including optimization [32]–[37]. This suggests that pseudoknot-inclusive RNA folding in the MFE approximation could see improvements if framed in a manner amenable to solution by quantum computational techniques.

In particular, both quantum annealing machines and circuit-model quantum computers are capable of solving quadratic unconstrained binary optimization (QUBO) problems, while offering advantage over classical approaches [38]–[42]. Such problems have the following cost function:

\*These authors contributed equally to this work.

$$\mathcal{H}_{\text{QUBO}}(\mathbf{q}) = \sum_i h_i q_i + \sum_{i>j} J_{ij} q_i q_j \quad (1)$$

The parameter  $h_i$  is called the bias on the  $i$ -th binary variable,  $q_i \in \{0, 1\}$ , and  $J_{ij}$  is called the coupling between binary variables  $i$  and  $j$ . The  $h_i$  and  $J_{ij}$  therefore define the problem to be solved. In the case of quantum annealing, the adiabatic theorem is exploited to continuously evolve the known, easy-to-prepare minimum-energy solution (ground-state) of an initial Hamiltonian to the ground-state of a final, non-commuting Hamiltonian encoding the problem cost-function [43]. Circuit-model quantum computers solve QUBOs through digitized versions of adiabatic quantum evolution protocols and quantum sampling methods [44]–[47].

Given the diversity of quantum approaches suitable to solve QUBOs, a sensible challenge then for RNA folding is to design a **robust energy function in QUBO form**. Currently, two such QUBO functions have been presented, but each remains with various shortcomings [48]–[50]. In this paper, we first review these existing solutions, highlighting their respective benefits and drawbacks, before advancing a novel formulation which seeks to combine their principal virtues with additional sophistications and commenting on complexity results. We then tune the parameters of each QUBO using a training set of known RNA structures using what we term “variational hybrid quantum annealing”, and compare their performance over an additional testing set of known RNA structures. Finally, we speculate on future directions to quantum RNA folding.

## II. METHODS

Our approach and contributions are organized into four important parts: (II-A) literature review, where we discuss each current RNA QUBO turn to motivate our own work, (II-B) RNA QUBO design, where we outline the features of our RNA QUBO model, (II-C) RNA QUBO training, where we describe our procedure for tuning the parameters of our QUBO model against real RNA structures, and (II-D), RNA QUBO testing, where we delineate our method of testing our QUBO model against real structures using the D-Wave hybrid quantum-classical annealer.

### A. Literature review

One existing RNA QUBO model (hereafter referred to as model 1) seeks to identify the set of potential stems that maximizes both the number of **consecutive base pairs and the average length of stems** [48]. This amounts to a minimization of the following QUBO equation:

$$\begin{aligned} \mathcal{H}_{\text{RNA}_1}(\mathbf{q}) = & \sum_i \{c_L(k_i - \mu)^2 - c_B k_i^2\} q_i \\ & - \sum_{i>j} \{2c_B k_i k_j \delta_{ij} + \delta'_{ij}\} q_i q_j \end{aligned} \quad (2)$$

Here,  $q_i$  corresponds to the binary variable for the  $i$ -th potential stem, the parameters  $c_L$  and  $c_B$  are tunable constants,  $k_i$  is the length of potential stem  $i$  measured in base-pairs, and  $\mu$

is the longest stem length in the potential set, all determined by classical pre-processing. A stem is a set of consecutive base-pairs, as shown in Figure 1.  $\delta_{ij}$  and  $\delta'_{ij}$  are delta-function penalties applied when  $i$  and  $j$  are pseudoknots and overlapping, respectively. They are defined as:

$$\delta_{ij} = \begin{cases} 1 & \text{if } i, j \text{ are not pseudoknotted} \\ X & \text{if } i, j \text{ are pseudoknotted} \end{cases} \quad (3a)$$

$$\delta'_{ij} = \begin{cases} 1 & \text{if } i, j \text{ are not overlapping} \\ -\infty & \text{if } i, j \text{ are overlapping} \end{cases} \quad (3b)$$

Note that  $\delta_{ij}$  increasingly penalizes pseudoknots with decreasing  $X$ . The authors of [48] chose  $X = 0.5$  in equation (3a).

The linear terms of the QUBO model consist of two parts, one (associated with  $c_L$ ) comparing the length of the  $i$ -th stem to the longest stem of the potential set, thereby enforcing the notion that average stem length should be maximized, and the second (associated with  $c_B$ ) offering quadratic reward to longer stems. The quadratic terms penalize pseudoknotted stems in proportion to the product of their lengths ( $\delta_{ij}$ ), and prevent overlapping stems from appearing in the solution based on an overwhelming constant addition to the energy function if this is the case ( $\delta'_{ij}$ ).

Though the essence of this QUBO model reflects certain known features of RNA folding (that longer stems are more energetically favorable, and that pseudoknotted stems are less stable than the same stems not in pseudoknot), these reflections are fairly coarse-grained. For one, the length of a stem is not directly proportional to its stability. The various typical base pairs ( $G, C$ ), ( $A, U$ ), and ( $G, U$ ) have differential stabilities when considered independently [51] that combine non-trivially in nearest-neighbor stacking interactions [52], [53]. Furthermore, the approximation that pseudoknots should be penalized in proportion to the product of their lengths does not match previous, experimentally-informed attempts at representing these structures and remains without strong physical reasoning [54]–[56]. Finally, given that this QUBO is heuristic in its construction, optimal performance must require parameter tuning and testing against known RNA structures; to the best of our knowledge, this has not been carried out so far.

It should be noted that each term of equation 2 must be pre-computed classically before the Hamiltonian can be embedded within quantum hardware for solving. It is therefore important to understand the computational complexity involved for reference when comparing to existing, classical approaches. The appendix offers proof that the time-complexity of the pre-computation step for this Hamiltonian is  $\mathcal{O}(N^3 \log N)$ , where  $N$  is RNA length, and the length of the binary vector  $\mathbf{q} = S(N, m)$ , where  $m$  is the minimum stem length, is  $\mathcal{O}(N^3)$  in the worst case. Native RNAs, however, typically demonstrate a  $q$ -length on the order of  $\mathcal{O}(N^2)$  (computed based on a fit of  $S(N, m) = aN^b$  to all structures of the Archive II RNA benchmarking suite [57]).

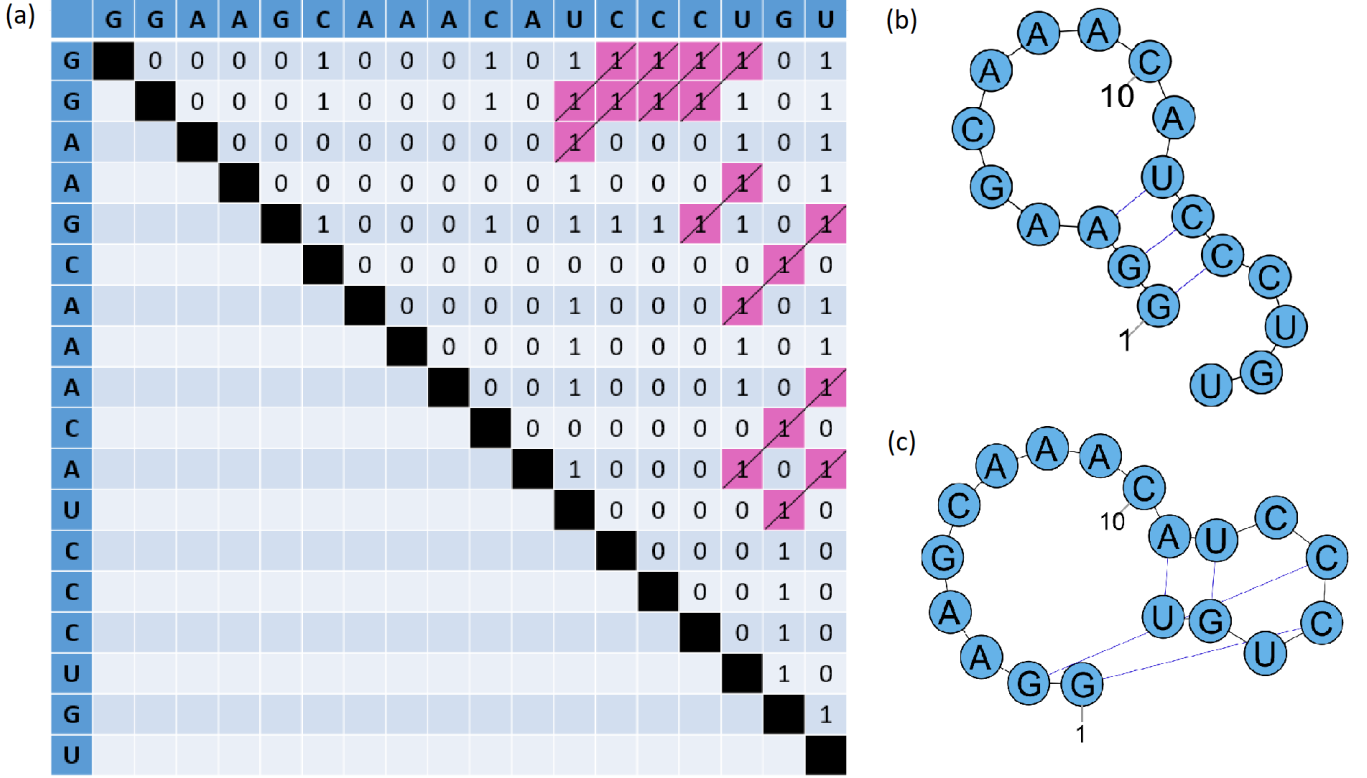


Fig. 1: (a) Matrix representation of potential base-pairs and stems for an example sequence. Potential base-pairs are marked with a 1, and stems (of at least two successive base-pairs) are marked in pink with diagonal lines to show the base-pairs involved in the stem. Note that any stem of length 3 or greater consists of overlapping stacked quartets (stems of length 2). For model 2, stacked quartets are mapped to individual qubits; for models 1 and 3, stems are mapped to individual qubits. (b) A possible structure for the example sequence comprised of one stem of length 3. In this case, two stacked quartets combine to form this stem. (c) A possible structure for the example sequence comprised of two stems, each of length 2. Notice that these stems are pseudoknotted. This figure is adapted from ref. [48].

The other existing RNA QUBO model (hereafter referred to as model 2) seeks to identify the set of potential stems of length two (stacked quartets) that maximizes successive stacked quartets while guarding against overlaps and penalizing pseudoknots [49], [50]. This amounts to a minimization of the following QUBO equation:

$$\mathcal{H}_{\text{RNA}_2}(\mathbf{q}) = - \sum_i N_i q_i - \sum_{i>j} M_{ij} q_i q_j \quad (4)$$

Here,  $q_i$  is the binary variable corresponding to the  $i$ -th potential stacked quartet,  $N_i$  is the nearest-neighbor stacking energy of the  $i$ -th quartet determined from optical melting experiments at 37 °C [58], and  $M_{ij}$  is of the form:

$$M_{ij} = \begin{cases} M^+ & \text{if } i, j \text{ are stacked} \\ M^- & \text{if } i, j \text{ are pseudoknotted} \\ M^P = -\infty & \text{if } i, j \text{ are overlapping} \end{cases} \quad (5)$$

$M^+$  rewards consecutive stacked quartets to incentivize stem formation,  $M^-$  penalizes pseudoknot formation between stacked quartets, and  $M^P$  is a large penalty that prevents the same base from being assigned to a different matching pair.  $M^+$  and  $M^-$  are positive, tunable values, with  $M^- < M^+$

representing a soft constraint, where violation of the constraint might allow for the formation of a longer stem that further minimizes the cost-function.  $M^P$ , by contrast, is negative.

An advantage of this approach versus the former is its direct use of experimental data to inform the model, where the real free energy contributions of stacked quartets are considered. However, due to the computational variables being defined at this level, it is not possible to reward the formation of larger stems containing some number of base pairs versus several smaller stems totalling the same number of base pairs. Further, as with the previous model, the pseudoknot penalty is again proportional to the product of the stem lengths forming the pseudoknot, which deviates significantly from previous work [54]–[56]. Finally, the  $M_{ij}$  terms are not experimentally-informed, and so must be pre-selected in advance, or tuned against known RNA structures. It is our understanding that this tuning has not been carried out previously.

As with the previous model, each parameter must be pre-computed classically before quantum embedding. The computational complexity of this stage is shown in the appendix to be on the order of  $\mathcal{O}(N^2 \log N)$  in time, and generates a binary vector  $\mathbf{q}$  of  $\mathcal{O}(N^2)$  in length. Therefore, there is a polynomial advantage to this approach versus the former

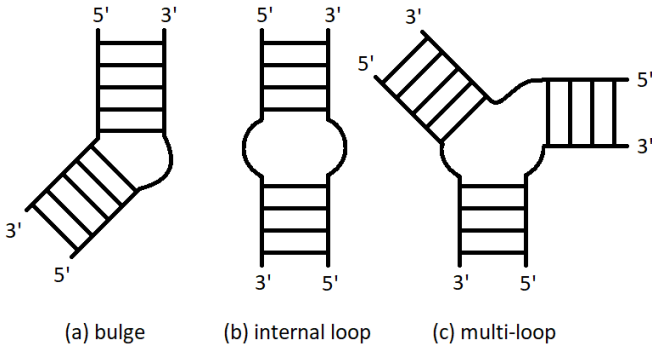


Fig. 2: RNA structural motifs that are difficult to penalize directly within an RNA folding QUBO model. (a) bulges consist in unpaired stretches of base located within one strand of an otherwise continuous “stem”, (b) internal loops consist in unpaired stretches of base separating two stems, and (c) multi-loops are internal loops with three or more participant stems.

in pre-processing time and quantum space-complexity, which becomes significant at larger sequence lengths.

### B. RNA QUBO design

As outlined in section II-A, we identify several features desirable at this stage in an RNA QUBO model, some of which are implemented variously by the two existing solutions covered. These features include: physically-informed representations of base-pairing energies at the stem level and physically-informed pseudoknot penalties. Further, we consider rewarding the formation of larger stems over a number of smaller stems that are the same in summed energy when considering only nearest-neighbor interactions important within a QUBO model context. Specifically, this originates from the fact that **bulges, internal loops, multi-loops**, and other such structural motifs are known to be destabilizing [59], features not feasibly pre-computable given that we must enumerate every possibility in advance prior to submission to a QUBO. See Figure 2 for a schematic representation of these structures. This heuristic therefore favors the formation of continuous base-pairs, and may be constructed as to be tunable, as per ref. [48].

With these considerations in mind, a sketch of our proposed QUBO model (hereafter called model 3) is as follows: we first pre-compute and optimize over stems as opposed to stacked quartets, even though this comes with a slight polynomial cost. This permits us to appraise in advance the lengths of two stems that may be involved in a pseudoknot, thereby allowing for pseudoknots to be penalized in a more sophisticated manner. Moreover, computing stems allows for an understanding of possible hairpin loop lengths, and the inclusion of an associated penalty. Secondly, we opted to maintain the term of ref. [48] which compares the base-pair length of the  $i$ -th potential stem to **the maximum potential stem length in order to maximize the average stem length**. Here, we replace base-pair stem length with **total nearest-neighbor-based stem energy**. Finally, we include a term that is linear in nearest-neighbor stem energy, and incorporate a penalty which prevents against

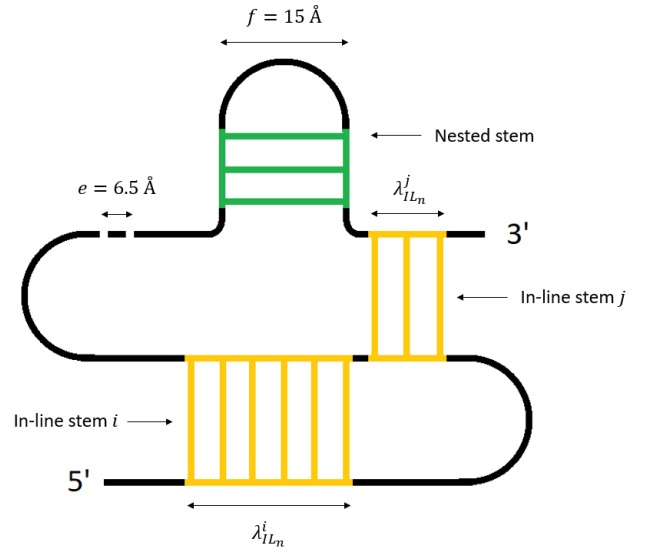


Fig. 3: Overview of pseudoknot structure model and entropic penalty terms. Energy penalties for single-stranded nucleotides and nested stems are based on ref. [60], and the penalty for in-line stems was developed in Hajdin et al., 2013 [56]. This figure is adapted from ref. [56].

bases from being forced to pair un-physically with more than one base.

The pseudoknot penalty we abstract from ref. [56], which is based on a polymer-physics model [60]. The form of this penalty is based on the insight that pseudoknots typically exhibit short single-stranded sections between their constituent in-line stems, and avoid the inclusion of nested stems. Essentially, the energetic penalties associated with single-stranded nucleotides and nested stems between in-line stems are in proportion to the linear dimension of each of these features along the strand, as shown in Figure 3. Furthermore, in-line stems are penalized based on their length with reference to the distribution of in-line stem lengths of known structures (see supplementary material of ref. [56]). In full form, this penalty ( $P_{PK}$ ) appears as:

$$P_{PK} = P_1 \ln(e^2 N_{SS} + f^2 N_{NH}) + P_2 \ln \left( \sum_k \lambda_{IL_n}^k \right) \quad (6)$$

Where  $P_1$  and  $P_2$  are adjustable parameters,  $e$  and  $f$  are the linear dimensions of nucleotides and nested stems in-strand, respectively,  $N_{SS}$  and  $N_{NH}$  are the numbers of single-stranded nucleotides and nested stems involved in the pseudoknot, respectively, and  $\lambda_{IL_n}^k$  is the penalty constant for the  $k$ -th in-line stem of length  $n$ .

Now, as before with internal loops and bulges, it is difficult with the QUBO approach to feasibly pre-compute this penalty term per pair of stems participating in a pseudoknot, given that many potential nested stems may be involved in the actual pseudoknot. As a compromise, we forgo the  $N_{NH}$  term

and simply consider the number of nucleotides between the constituent stems, assuming them to be single-stranded, such that we operate with this modified pseudoknot penalty:

$$P'_{PK} = P_1 \ln(e^2 N_{SS}) + P_2 \ln \left( \sum_k \lambda_{IL_n}^k \right) \quad (7)$$

Assuming the previous model to be more accurate, all else equal, this compromise has us over-penalizing pseudoknots with large loops to some degree. Nevertheless, the essence of the model is captured in that larger loop structures are penalized more within the context of a pseudoknot.

We also incorporate a hairpin loop penalty ( $l_i$ ) when considering a single stem, as per refs. [58], [61]. With this penalty, loops of size 1 and 2 are forbidden, while loops of size 3 or greater are with decreasing penalties until loops of size 7, beyond which all loops are assigned the same penalty. It is possible that the nucleotides of certain hairpin loops penalized in this way might participate in pseudoknots, thereby negating the need for a loop penalty (given that single-stranded sections to pseudoknots are penalized according to equation 7), but the probability of pseudoknots forming with short hairpin loops, where the loop penalty is highest in magnitude and rate of change, is small. Finally, we include a linear term proportional to the nearest-neighbor energy of the  $i$ -th stem ( $k_i$ ) and the delta-function of equation 3b, yielding this QUBO equation:

$$H = \sum_i \{ \alpha(k_i - \mu)^2 - \beta(k_i - l_i) \} q_i + \sum_{i,j>i} (P'_{PK} + \delta'_{ij}) q_i q_j \quad (8)$$

$\alpha$  and  $\beta$  of equation 8 are tunable parameters, as are  $P_1$  and  $P_2$  in equation 7. In section II-C, we describe our tuning of these parameters, as well as those which appear in the existing solutions of section II-A.

### C. RNA QUBO training

To enable the evaluation of our parameterized RNA folding QUBO model against the models of previous work and to assess general performance, we tuned the parameters of each model by training over known RNA secondary structures with and without pseudoknots. To accomplish this, we collected 110 sequences and associated connectivity table structure files with less than 90% sequence identity from the bpRNA-1m database [62] ranging between 20 and 50 bases in length. We divided the RNAs into training and testing groups, containing 70 and 40 sequences, respectively. Within each of these groups, half contained pseudoknots. The purpose of the with/without pseudoknots division is in following the RNA structure prediction benchmarking practices outlined in ref. [63], where it is recommended that structures with and without pseudoknots be tested against models separately and their individual performances noted. Training was carried out using the D-Wave Advantage system together with a classical parameter optimizer as instantiated in the Amazon Braket Hybrid Jobs platform. Below we present our approach in

more detail, which is summarized in Figure 4. We propose to call this approach “variational hybrid quantum annealing”, in parallel to gate-based variational hybrid quantum algorithms.

**Stem and stacked-quartet identification:** We implemented two versions of an algorithm to extract lists of potential stems and stacked quartets per sequence, respectively. A potential base-pairing matrix populated only in the upper diagonal as shown in Figure 1 was constructed per RNA sequence, where 0s indicate the impossibility of a base-pair, and 1s base-pair possibilities. We considered  $(G, C)$ ,  $(A, U)$ , and  $(G, U)$  as legitimate base-pairs. We proceeded to extract stems (of at least two base-pairs, given the basis of model 3 being in nearest-neighbor stacking interaction energies) and stacked-quartets by scanning matrices for consecutive possible base-pairs in upper-right diagonals in  $\mathcal{O}(N^2)$  time. Following ref. [48], we represented stems and stacked quartets by listing the indices of the first and last base involved in pairing, and either the length of the stem (in the case of model 1) or the total nearest-neighbor stacking interaction energy (in the case of model 2 and model 3), keeping track of the largest potential stem,  $\mu$ , in the case of model 1 and model 3. The stem identification process for an example sequence 5'-GGAAGCAAACAUCCCUGU-3' is depicted in Figure 1.

**Pseudoknot and overlap penalties:** For each model and RNA considered, each pairwise combination of stems or stacked quartets was checked for pseudoknotting and overlap using a line-sweep interval overlap algorithm. For pseudoknots, we created lists of 3-tuples that identified a combination of stems  $(i, j)$  and their penalty score. If stems were found to be in pseudoknot in the case of models 1 and 2, this penalty score assumes the form of a tunable constant. With model 3, each potential pseudoknot was evaluated according to equation 7, requiring an evaluation of the assumed single-stranded nucleotides within the pseudoknot and a constant lookup of  $\lambda_{IL_n}$  with respect to the lengths of the participant stems. In the case of overlapping stems, we used the same 3-tuple notation, but with the constant penalty  $-\infty \approx -1000$ .

**Parsing training structure data:** As mentioned at the beginning of this section, we retrieved the known secondary structures for each sequence as connectivity table files from bpRNA-1m. We implemented an auxiliary method to take the contents of a connectivity table file and output the stems in a format matching the potential stems on a per-model basis, as well as the length of the longest actual stem,  $\mu$ , in model 1 and model 3. This allowed for easy comparison of the known-structure base-pairs of a sequence to the base-pairs predicted by the models, as well as energy evaluation of known structures under the assumption of the models.

**Model construction:** With lists of all possible stems or stacked quartets and pairwise combination penalties for a sequence, we employed equations 2, 4, and 8 to instantiate each QUBO model programmatically in Python dictionary form, one dictionary corresponding to linear terms and the other dictionary to quadratic terms. The keys of the former we set to individual stems or stacked quartets, and the latter we set to pairs of stems or stacked quartets. To start, we initialized

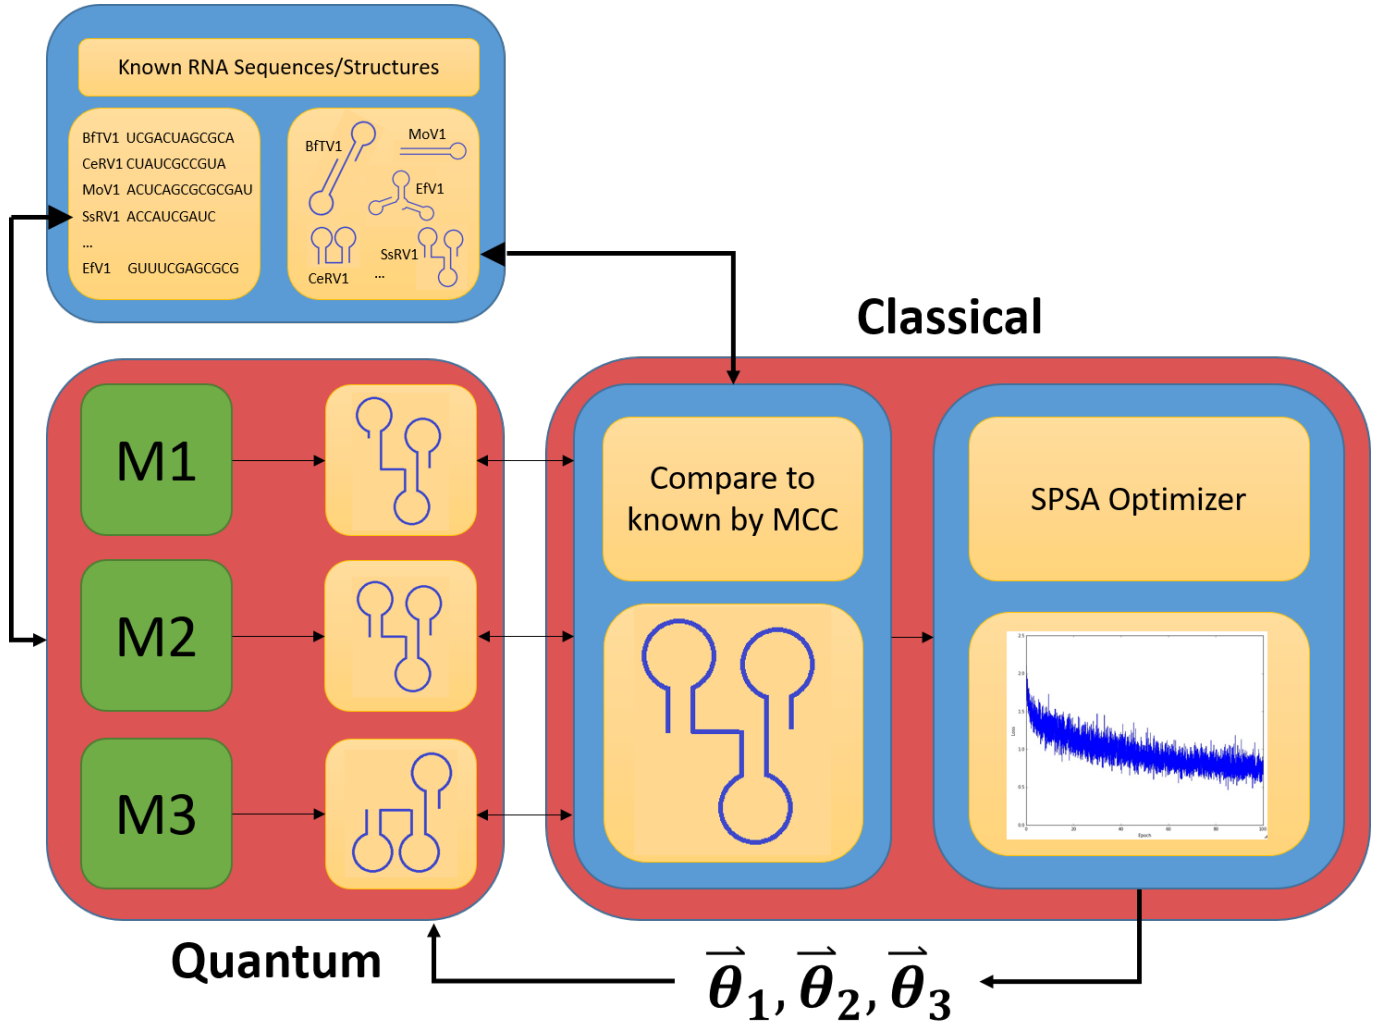


Fig. 4: Overview of RNA QUBO training protocol. RNA sequences with known structures are passed to parameterized QUBO models, which are executed by the D-Wave hybrid solver. Predicted structures of each model are compared to known structures using the MCC scoring function [64]. An SPSA optimizer iteratively updates the parameters of the QUBO models. The arithmetic mean of the MCC evaluation metric over all test structures is optimized.

all tunable parameters of each model to 1.

**Hybrid quantum annealing:** Robust quantum error correcting codes are yet to be implemented on quantum annealing machines, such that purely quantum processing unit-based (QPU-based) outputs are noise-prone. For this reason, ref. [48] executed their experiment on the D-Wave hybrid solver, which breaks a QUBO into smaller sub-QUBOs and iterates between running these sub-QUBOs on the D-Wave QPU and through a classical tabu search after recombination [65]. We opted for this same hybrid approach using the 5000-qubit D-Wave Advantage machine with Pegasus-graph qubit connectivity, keeping default settings to the annealing schedule. As detailed above, we formulated each problem corresponding to an RNA sequence, model, and pseudoknot penalty in a dictionary form acceptable by the D-Wave QPU interface. Problems are mapped such that single stems or stacked quartets correspond to individual qubits, and the couplings between any pair  $(i, j)$  map to chains of interconnected qubits between  $(i, j)$ . For each

run, the solution was returned as a binary vector with a length corresponding to the number of potential stems or stacked quartet of the tested RNA, entries of 1 recording predicted stems or stacked quartets, and an entries of 0 recording non-predicted stems or stacked quartets. Additionally, the energy of the predicted solution and minimum estimated time-to-solution were returned per run.

**Model evaluation:** To evaluate the predicted structures versus the known structures of an RNA, the Matthews Correlation Coefficient (MCC) score over exact base-pairs was employed [64]. This metric measures the quality of binary classifications, taking into account true positive, true negative, false positive, and false negative counts such that perfect predictions score +1, total disagreement is indicated by -1, and random predictions score 0. Versus other two-class confusion matrix scoring metrics such as  $F_1$ , accuracy, markedness, bookmaker informedness, and balanced accuracy, the MCC in general proves more informative and reliable [66]. Our requirement

that base-pairs must be exact excludes similar solutions from scoring well [63]. For example, if we consider an RNA with the single known stem (4, 13, 3), where the last tuple entry records the base-pair length of the stem, the predicted stem (5, 14, 3) will score poorly, even though four of the six bases involved in pairing in the known structure are paired in the predicted. However, strict training is enforced through use of this scoring protocol.

**Model training:** We trained the parameters of each model over the 30 training structures using the Qiskit implementation of the simultaneous perturbation stochastic approximation (SPSA) optimization algorithm while carrying out QUBO function evaluations using the D-Wave hybrid solver as described [67], [68]. This method is based on a gradient approximation using two measurements of the loss function, where we took the arithmetic mean of the MCC score across all structures of a given training class is taken to be the loss, and perturbations to the current parameter set are stochastically generated from a zero-mean Bernoulli distribution. We observed convergence of this optimization at approximately 30-60 iterations of the optimizer in each model implementation.

#### D. RNA QUBO testing

Machine learning methods run the risk of over-fitting to training data [69]–[71]. It is therefore necessary to test machine-learning models against data other than that used in training to evaluate their realism and applicability to future instances. We test each fit model as described in (II-C) over 40 known RNA structures with less than 90% sequence identity versus those used in training, divided as before into two groups: one containing 20 structures with pseudoknots, and the other containing 20 structures without pseudoknots. We test for significant difference in performance between models over all test structures (with and without pseudoknots) using the Kolmogorov-Smirnov 2-sample test, and additionally test for intra-model difference in performance between structures with and without pseudoknots.

### III. RESULTS

Training each RNA folding QUBO yielded a set of parameters optimized against the training set, as well as a mean-squared loss score characterizing goodness of fit. For model 1, we found the parameters:  $c_L = 0.639$ ,  $c_B = 0.223$ , and  $\delta = 0.681$ , with a loss of 0.113. For model 2, we found:  $M^+ = 1.748$  and  $M^- = 0.386$ , with a loss of 0.263. For model 3, we found:  $\alpha = 1.604$ ,  $\beta = 2.212$ ,  $P_1 = 1.495$ , and  $P_2 = 1.338$ , with a loss of 0.116.

Given these optimized parameters, we tested each model against an additional 40 RNA structures. We present the results of this testing in Figure 5 as a boxplot with reference to MCC, where we differentiate performance over structures with pseudoknots and structures without pseudoknots. Comparing model performance over all test structures using the Kolmogorov-Smirnov 2-sample test, we did not find any significant difference between them (at significance levels of

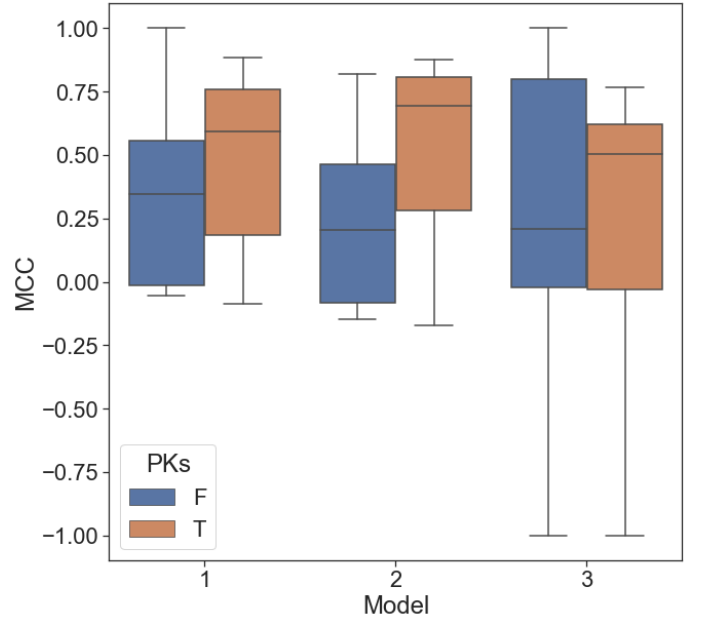


Fig. 5: Performance of RNA folding QUBOs after SPSA training as measured by the Matthews Correlation Coefficient. The same 40 structures (20 structures with pseudoknots, labelled ‘T’ and 20 structures without pseudoknots, labelled ‘F’) were evaluated per model. Whiskers extend to the lowest (highest) data point within 1.5 times the interquartile range as measured from the lower (upper) quartile.

0.01 and 0.05). However, it should be pointed out that models 1 and 2 demonstrate a tighter variance than model 3.

Additionally, it may be seen that models 1 and 2 appear to differ appreciably in their performance over structures with pseudoknots versus structures without. A Kolmogorov-Smirnov 2-sample test confirms this difference at both 0.01 and 0.05 levels of significance, while reporting model 3 as consistent in its performance across structures with and without pseudoknots.

### IV. DISCUSSION

We approach this section by first discussing our results and contributions in further detail, before proceeding to the limitations of our methodology and suggestions about how it might be improved.

#### A. Performance of models over pseudoknotted versus unpseudoknotted structures

We begin by considering the significant difference in performance seen in models 1 and 2 over structures with pseudoknots versus structures without pseudoknots. We suspect that this is due to our training and test structures being limited to sequences of less than 50 bases. With short sequences, the number of potential stems is quite small, and the number of potential stems in potential pseudoknots of course even smaller. Therefore, it is likely that if a potential pseudoknot (or multiple) is present for a given sequence, the actual structure will contain a pseudoknot. This means that during training, it is

possible for the pseudoknot penalty to offer a relatively high reward for pseudoknot formation without compromising too much model performance over structures without pseudoknots. In model 3, despite no statistically significant difference, the median MCC score of structures with pseudoknots versus those without is greater by an appreciable degree.

These non-pseudoknotted structures are at somewhat a “disadvantage” in that they are with one less term in helping to determine which of a set of possible structures is correct. Consequently, it is more difficult to select the proper structure, as the energy landscape for such structures is of a higher degeneracy than structures with pseudoknots in the small regime we have tested. We expect that the difference in performance between structures with and without pseudoknots should lessen in each model if trained over larger RNAs, but this remains to be tested. We suggest this step as necessary, as it should be the case that the trained parameters reflect a training set representative of the full array of RNA structures. In this work, we limited ourselves to small structures to expedite computations.

### *B. Comparison of overall model performance*

We now comment on the negligible difference in overall performance between the models reviewed, as determined by the Kolmogorov-Smirnov 2-sample test. This result indicates that with respect to the training and testing structures of this work, model 3 does not demonstrate any advantage versus models 1 and 2, representing existing RNA folding QUBO models. This is despite deliberate construction of model 3 to account for terms excluded from consideration by models 1 and 2 that are commonly included within the majority of other RNA folding models. Specifically, model 3 includes hairpin loop penalties and a pseudoknot penalty which takes into account intervening bases involved in a pseudoknot of two stems (either single-stranded or in nested stems).

### *C. Limitations of approach*

Before speculating further, we discuss limitations that constrain the presented results. One, already noted, concerns the number and length of RNA structures used in both training and testing. Due to a limited amount of computational resources and the larger amount of time it takes to compute longer RNA sequences using the D-Wave hybrid solver, we had to restrict our training data set to only 30 small structures (from 20 to 50 bases). This might affect the ability of the models to capture the effect of structure size on their performance. Furthermore, the SPSA optimizer used during the training phase is based on gradient descent, which can converge on local minima for certain starting points. We acknowledge this being a possibility within the training of our models and suggest to try different starting points for the optimizer in further iterations of this work, or a Monte Carlo optimization approach to avoid local minima trapping. We note as well that the D-Wave hybrid solver is not guaranteed to find the global optimum of a given problem instance. Typically, many executions of the same problem instance are computed to determine the optimum with

greater likelihood, but due to resource limitations we opted for a greedy approach in which only one execution of each problem instance was run, both in the training and testing phases of our methods.

Turning our attention once again to specific shortcomings of model 3, we first remark that its pseudoknot penalty modeled on that of Hajdin et al. [56] ignores the possibility of nested stems, assuming and penalizing instead the intervening bases as if they are all single-stranded. Additionally, our hairpin loop penalty assumes the stem involved is in fact a hairpin loop, with the loop not being involved in another stem. This is clearly impossible to satisfy in every case, especially as structure size increases. It is true as well that we do not consider direct penalizations of internal loops, multi-loops, bulges, and other such structural motifs, except through the heuristic that tends to avoid them in favour of longer stems if possible. We note that some of these penalties may prove beyond the capacity of the QUBO representation: multi-loops, for instance, are formed more than by pairwise interactions between stems.

Indeed, in general it might prove exceedingly challenging to design a sufficiently accurate RNA folding QUBO model for the reason that the method essentially demands the encoding of all possible structures at once. To achieve this while penalizing common RNA structural motifs that do not lend themselves easily to pairwise interaction terms in advance of solving the QUBO will likely remain possible only through clever heuristics. For example, consider an internal loop. If modelling at the stem level, we can determine whether two stems form an internal loop together. However, it is not possible to penalize this internal loop directly, as it might be that a portion of the internal loop will actually prove involved in another stem. However, it might be possible to penalize internal loops directly if they are smaller than a certain size, as beyond this size, it becomes impossible (or at least very improbable) for the internal loop to participate in other structures. A similar approach might be used for hairpin loops and pseudoknots, perhaps informed by empirically-derived statistical models of the probabilities of their being of a certain size, but complex motifs such as kissing hairpins and multi-loops escape simple treatment.

Despite these difficulties, we suggest the pursuit of a QUBO model of the RNA folding problem as one of salience, as this method lends itself well to the two leading quantum computing architectures (annealing and circuit-model), both of which may result in quantum advantage versus classical machines in solving QUBOs. Though in our work we execute RNA folding QUBOs by way of annealing hardware only, we emphasize that these models are easily transferable to universal quantum computers, where algorithms such as the Variational Quantum Eigensolver (VQE) and the Quantum Approximate Optimization Algorithm (QAOA) provide a feasible way of computing solutions to QUBOs.

#### D. Variational hybrid quantum annealing

In this section, we highlight the uniqueness of our method. We present a practical example of how quantum annealing can be used in the context of a supervised machine learning problem, where our RNA folding QUBO models are the analogs to ML models that then were trained using a hybrid quantum-classical solver in coupling with a classical gradient descent optimizer. To the best of our knowledge, this is the first experimental paper that demonstrates the implementation of variational hybrid quantum annealing. Other research presents different approaches to exploiting quantum annealing in machine learning, such as using it to optimize the training pipeline of a classical ML model for classification problems [72] or to construct diabatic Ising-Born machines that act as the analog of an ansatz for parameterized quantum circuits in generative learning problems [73]. Given the similarity in structure of variational quantum algorithms on circuit-model hardware, we suggest our method as the quantum annealing extension to these methods. We also suggest that though the classical optimization of this work took place at a large physical distance from the QPU, speedup might be seen if classical resources dedicated to optimization were to be integrated in close proximity to the QPU, as with Qiskit Runtime [74].

#### E. Future work

Moving forward, we expect to include larger and more complex RNA structures in the training process of the models to identify possible size dependence of the optimized model parameters. As well, we plan to increase the number of optimizer iterations and initial parameter points used to train the models to ensure that the final set of parameters represents the lowest loss possible. Another important consideration is to thoroughly review the structures which performed particularly poorly in model 3, and seek to improve the corresponding Hamiltonian based on our findings. We also plan to investigate penalizations of other structural motifs besides hairpin loops and pseudoknots as mentioned. One final point is that MCC does not provide a full measure of fitness for an RNA folding model. For instance, the relative difference in energy between the predicted and actual structures of a large suite of test sequences under a given model suggests the closeness of this model to reality if the model is reasonably complex, where this metric tends to zero as the model converges to accurate replication of reality. We aim to consider this means and others in assessing model performance.

#### V. CONCLUSION

We present herein a novel, physically-motivated RNA folding QUBO model amenable to quantum computers, and train the free parameters of this and other existing RNA folding QUBOs using hybrid quantum-classical annealing. We find that after training, all models perform similarly, but stress that both training and testing were limited to RNA sequences of less than 50 bases, and that working with larger sequences will

be necessary before full conclusions of relative performance between models can be drawn.

We would like to highlight our approach of combining a quantum annealing solver with classical parameter optimization as an additional important contribution of this paper. While existing quantum variational algorithms employ this principle with circuit-model solvers, to date it appears that its application to quantum annealers has not been carried out. We suggest that our variational hybrid quantum annealing approach might be of especial benefit to training the free parameters of Ising problem models.

Finally, despite the limitations of RNA folding QUBO models (i.e., that certain structural motifs cannot be directly penalized), we advocate that the motivation for continuing research in their development for solution by quantum computers is strong. We firmly believe that quantum computing is with the potential to increase the tractability of the full RNA folding problem including pseudoknots and suggest this pursuit in earnest.

#### ACKNOWLEDGMENT

We thank the NSERC CREATE in Quantum Computing program and the University of Victoria for providing funding for this project. Additionally, we thank the organizers and sponsors of QHack 2022 for offering a platform through which to showcase this work. We especially thank Amazon Braket for awarding us the compute credits to carry out our experiments.

#### REFERENCES

- [1] P. G. Higgs, "RNA secondary structure: physical and computational aspects," *Quarterly Reviews of Biophysics*, vol. 33, no. 3, pp. 199–253, Aug. 2000.
- [2] I. Brierley, S. Pennell, and R. J. C. Gilbert, "Viral RNA pseudoknots: versatile motifs in gene expression and replication," *Nature Reviews Microbiology*, vol. 5, no. 8, pp. 598–610, Aug. 2007.
- [3] S. R. Eddy, "Non-coding RNA genes and the modern RNA world," *Nature Reviews Genetics*, vol. 2, no. 12, pp. 919–929, Dec. 2001.
- [4] G. F. Joyce and J. W. Szostak, "Protocells and RNA self-replication," *Cold Spring Harbor Perspectives in Biology*, vol. 10, no. 9, p. a034801, Sep. 2018.
- [5] S. Halder and D. Bhattacharyya, "RNA structure and dynamics: a base pairing perspective," *Progress in Biophysics and Molecular Biology*, vol. 113, no. 2, pp. 264–283, 2013.
- [6] G. L. Conn and D. E. Draper, "RNA structure," *Current opinion in structural biology*, vol. 8, no. 3, pp. 278–285, 1998.
- [7] S. R. Holbrook, "RNA structure: the long and the short of it," *Current Opinion in Structural Biology*, vol. 15, no. 3, pp. 302–308, Jun. 2005.
- [8] T. Pan and T. Sosnick, "RNA folding during transcription," *Annu. Rev. Biophys. Biomol. Struct.*, vol. 35, pp. 161–175, 2006.
- [9] Y. Wan, M. Kertesz, R. C. Spitale, E. Segal, and H. Y. Chang, "Understanding the transcriptome through RNA structure," *Nature Reviews Genetics*, vol. 12, no. 9, pp. 641–655, Aug. 2011.
- [10] E. J. Strobel, A. M. Yu, and J. B. Lucks, "High-throughput determination of RNA structures," *Nature Reviews Genetics*, vol. 19, no. 10, pp. 615–634, Jul. 2018.
- [11] S. Bottaro, F. D. Palma, and G. Bussi, "The role of nucleobase interactions in RNA structure and dynamics," *Nucleic Acids Research*, vol. 42, no. 21, pp. 13 306–13 314, Oct. 2014.
- [12] B. A. Shapiro, Y. G. Yingling, W. Kasprzak, and E. Bindewald, "Bridging the gap in RNA structure prediction," *Current Opinion in Structural Biology*, vol. 17, no. 2, pp. 157–165, Apr. 2007.
- [13] R. Giegerich, "Abstract shapes of RNA," *Nucleic Acids Research*, vol. 32, no. 16, pp. 4843–4851, Sep. 2004.
- [14] E. Westhof, B. Masquida, and L. Jaeger, "RNA tectonics: towards RNA design," *Folding and Design*, vol. 1, no. 4, pp. R78–R88, Aug. 1996.

- [15] M. Andronescu, A. P. Fejes, F. Hutter, H. H. Hoos, and A. Condon, "A new algorithm for RNA secondary structure design," *Journal of Molecular Biology*, vol. 336, no. 3, pp. 607–624, Feb. 2004.
- [16] M. McKeague, R. S. Wong, and C. D. Smolke, "Opportunities in the design and application of RNA for gene expression control," *Nucleic Acids Research*, vol. 44, no. 7, pp. 2987–2999, Mar. 2016.
- [17] É. Bonnet, P. Rzazewski, and F. Sikora, "Designing RNA secondary structures is hard," *Journal of Computational Biology*, vol. 27, no. 3, pp. 302–316, Mar. 2020.
- [18] I. Tinoco and C. Bustamante, "How RNA folds," *Journal of Molecular Biology*, vol. 293, no. 2, pp. 271–281, Oct. 1999.
- [19] R. B. Lyngsø and C. N. S. Pedersen, "RNA pseudoknot prediction in energy-based models," *Journal of Computational Biology*, vol. 7, no. 3–4, pp. 409–427, Aug. 2000.
- [20] M. Zuker and P. Stiegler, "Optimal computer folding of large RNA sequences using thermodynamics and auxiliary information," *Nucleic acids research*, vol. 9, no. 1, pp. 133–148, 1981.
- [21] J. Ruan, G. D. Stormo, and W. Zhang, "An iterated loop matching approach to the prediction of RNA secondary structures with pseudoknots," *Bioinformatics*, vol. 20, no. 1, pp. 58–66, Dec. 2003.
- [22] J. Ren, B. Rastegari, A. Condon, and H. H. Hoos, "HotKnots: Heuristic prediction of RNA secondary structures including pseudoknots," *RNA*, vol. 11, no. 10, pp. 1494–1504, Sep. 2005.
- [23] S. Janssen and R. Giegerich, "Faster computation of exact RNA shape probabilities," *Bioinformatics*, vol. 26, no. 5, pp. 632–639, Jan. 2010.
- [24] S. Bellaousov and D. H. Mathews, "ProbKnot: Fast prediction of RNA secondary structure including pseudoknots," *RNA*, vol. 16, no. 10, pp. 1870–1880, Aug. 2010.
- [25] C. Waldsich, *RNA folding: methods and protocols*. New York: Humana Press, 2014.
- [26] S. Takitou and A. Taneda, "Ant colony optimization for predicting RNA folding pathways," *Computational Biology and Chemistry*, vol. 83, p. 107118, Dec. 2019.
- [27] L. Huang, H. Zhang, D. Deng, K. Zhao, K. Liu, D. A. Hendrix, and D. H. Mathews, "LinearFold: linear-time approximate RNA folding by 5'-to-3' dynamic programming and beam search," *Bioinformatics*, vol. 35, no. 14, pp. i295–i304, Jul. 2019.
- [28] K. Sato, M. Akiyama, and Y. Sakakibara, "RNA secondary structure prediction using deep learning with thermodynamic integration," *Nature Communications*, vol. 12, no. 1, Feb. 2021.
- [29] M. Gray, S. Chester, and H. Jabbari, "KnotAli: informed energy minimization through the use of evolutionary information," *BMC Bioinformatics*, vol. 23, no. 1, may 2022.
- [30] M. Gruzewski and M. Palkowski, "Implementation of nussinov's RNA folding using the kokkos library," in *Progress in Image Processing, Pattern Recognition and Communication Systems*. Springer International Publishing, Aug. 2021, pp. 310–317.
- [31] M. W. Lewis, A. Verma, and R. Hennig, "Predicting 3d RNA folding patterns via quadratic binary optimization," 2021.
- [32] R. Jozsa and N. Linden, "On the role of entanglement in quantum-computational speed-up," *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, vol. 459, no. 2036, pp. 2011–2032, Aug. 2003.
- [33] T. Zanca and G. E. Santoro, "Quantum annealing speedup over simulated annealing on random Ising chains," *Physical Review B*, vol. 93, no. 22, Jun. 2016.
- [34] A. Gilyén, S. Arunachalam, and N. Wiebe, "Optimizing quantum optimization algorithms via faster quantum gradient computation," in *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, Jan. 2019, pp. 1425–1444.
- [35] P. K. Barkoutsos, G. Nannicini, A. Robert, I. Tavernelli, and S. Woerner, "Improving variational quantum optimization using CVaR," *Quantum*, vol. 4, p. 256, Apr. 2020.
- [36] P. Hauke, H. G. Katzgraber, W. Lechner, H. Nishimori, and W. D. Oliver, "Perspectives of quantum annealing: Methods and implementations," *Reports on Progress in Physics*, vol. 83, no. 5, p. 054401, 2020.
- [37] A. Gilliam, S. Woerner, and C. Gondiulea, "Grover adaptive search for constrained polynomial binary optimization," *Quantum*, vol. 5, p. 428, Apr. 2021.
- [38] J. Golden, A. Bärttschi, S. Eidenbenz, and D. O'Malley, "Evidence for super-polynomial advantage of QAOA over unstructured search," 2022.
- [39] E. Zahedinejad and A. Zaribafian, "Combinatorial optimization on gate model quantum computers: A survey," 2017.
- [40] "Quantum annealing with manufactured spins," *Nature*, vol. 473, p. 194, 2011.
- [41] 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.*, vol. 10, p. 218, 2014.
- [42] S. Mukherjee and B. Chakrabarti, "Multivariable optimization: Quantum annealing and computation," *The European Physical Journal Special Topics*, pp. 17–24, Feb. 2015.
- [43] T. Albash and D. A. Lidar, "Adiabatic quantum computation," *Reviews of Modern Physics*, vol. 90, no. 1, Jan. 2018.
- [44] D. Poulin and P. Wocjan, "Sampling from the thermal quantum Gibbs state and evaluating partition functions with a quantum computer," *Physical review letters*, vol. 103, no. 22, p. 220502, 2009.
- [45] K. Temme, T. J. Osborne, K. G. Vollbrecht, D. Poulin, and F. Verstraete, "Quantum metropolis sampling," *Nature*, vol. 471, no. 7336, pp. 87–90, 2011.
- [46] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," *arXiv preprint arXiv:1411.4028*, 2014.
- [47] M. H. Amin, E. Andriyash, J. Rolfe, B. Kulchitsky, and R. Melko, "Quantum boltzmann machine," *Physical Review X*, vol. 8, no. 2, p. 021050, 2018.
- [48] D. M. Fox, C. M. MacDermaid, A. M. Schreij, M. Zwierzyzna, and R. C. Walker, "RNA folding using quantum computers," May 2021.
- [49] M. W. Lewis, A. Verma, and T. T. Eckdahl, "Qfold: a new modeling paradigm for the RNA folding problem," *Journal of Heuristics*, vol. 27, no. 4, pp. 695–717, Mar. 2021.
- [50] M. W. Lewis, A. Verma, and R. Hennig, "Predicting 3d RNA folding patterns via quadratic binary optimization," 2021.
- [51] G. Varani and W. H. McClain, "The g-u wobble base pair," *EMBO reports*, vol. 1, no. 1, pp. 18–23, Jul. 2000.
- [52] T. Xia, J. SantaLucia Jr, M. E. Burkard, R. Kierzek, S. J. Schroeder, X. Jiao, C. Cox, and D. H. Turner, "Thermodynamic parameters for an expanded nearest-neighbor model for formation of RNA duplexes with Watson-Crick base pairs," *Biochemistry*, vol. 37, no. 42, pp. 14719–14735, 1998.
- [53] M. Andronescu, A. Condon, D. H. Turner, and D. H. Mathews, "The determination of RNA folding nearest neighbor parameters," *RNA Sequence, Structure, and Function: Computational and Bioinformatic Methods*, pp. 45–70, 2014.
- [54] R. M. Dirks and N. A. Pierce, "An algorithm for computing nucleic acid base-pairing probabilities including pseudoknots," *Journal of computational chemistry*, vol. 25, no. 10, pp. 1295–1304, 2004.
- [55] M. S. Andronescu, C. Pop, and A. E. Condon, "Improved free energy parameters for RNA pseudoknotted secondary structure prediction," *RNA*, vol. 16, no. 1, pp. 26–42, 2010.
- [56] C. E. Hajdin, S. Bellaousov, W. Huggins, C. W. Leonard, D. H. Mathews, and K. M. Weeks, "Accurate SHAPE-directed RNA secondary structure modeling, including pseudoknots," *Proceedings of the National Academy of Sciences*, vol. 110, no. 14, pp. 5498–5503, Mar. 2013.
- [57] M. F. Sloma and D. H. Mathews, "Exact calculation of loop formation probability identifies folding motifs in RNA secondary structures," *RNA*, vol. 22, no. 12, pp. 1808–1818, Oct. 2016.
- [58] D. H. Turner and D. H. Mathews, "NNDDB: the nearest neighbor parameter database for predicting stability of nucleic acid secondary structure," *Nucleic Acids Research*, vol. 38, no. suppl\_1, pp. D280–D282, Oct. 2009.
- [59] J. Nowakowski and I. Tinoco Jr, "RNA structure and stability," in *Seminars in virology*, vol. 8, no. 3. Elsevier, 1997, pp. 153–165.
- [60] D. P. Aalberts and N. Nandagopal, "A two-length-scale polymer theory for RNA loop free energies and helix stacking," *RNA*, vol. 16, no. 7, pp. 1350–1355, May 2010.
- [61] S. T. Kelley and D. Didulo, *Computational Biology*. American Society for Microbiology, Feb. 2018.
- [62] P. Danaee, M. Rouches, M. Wiley, D. Deng, L. Huang, and D. Hendrix, "bpRNA: large-scale automated annotation and analysis of RNA secondary structure," *Nucleic Acids Research*, vol. 46, no. 11, pp. 5381–5394, May 2018.
- [63] D. H. Mathews, "How to benchmark RNA secondary structure prediction accuracy," *Methods*, vol. 162–163, pp. 60–67, Jun. 2019.
- [64] B. Matthews, "Comparison of the predicted and observed secondary structure of t4 phage lysozyme," *Biochimica et Biophysica Acta (BBA) - Protein Structure*, vol. 405, no. 2, pp. 442–451, Oct. 1975.

- [65] M. Booth, S. P. Reinhardt, and A. Roy, "Partitioning optimization problems for hybrid classical/quantum execution," D-Wave Systems, Inc., Tech. Rep., 2017.
- [66] D. Chicco, N. Tötsch, and G. Jurman, "The Matthews correlation coefficient (MCC) is more reliable than balanced accuracy, bookmaker informedness, and markedness in two-class confusion matrix evaluation," *Biodata Mining*, vol. 14, no. 1, Feb. 2021.
- [67] J. Spall, "Multivariate stochastic approximation using a simultaneous perturbation gradient approximation," *IEEE Transactions on Automatic Control*, vol. 37, no. 3, pp. 332–341, Mar. 1992.
- [68] M. Treinish, J. Gambetta, P. Nation, Qiskit-Bot, P. Kassebaum, D. M. Rodríguez, S. De La Puente González, Shaohan Hu, K. Krsulich, J. Garrison, L. Zdanski, J. Yu, J. Gacon, D. McKay, J. Gomez, L. Capelluto, Travis-S-IBM, M. Marques, A. Panigrahi, J. Lishman, Lerongil, Rafey Iqbal Rahman, S. Wood, L. Bello, Toshinari Itoko, Divyanshu Singh, , Drew, E. Arbel, J. Schwarm, and J. Daniel, "Qiskit/qiskit: Qiskit 0.36.0," 2022.
- [69] T. Dietterich, "Overfitting and undercomputing in machine learning," *ACM computing surveys (CSUR)*, vol. 27, no. 3, pp. 326–327, 1995.
- [70] R. Roelofs, V. Shankar, B. Recht, S. Fridovich-Keil, M. Hardt, J. Miller, and L. Schmidt, "A meta-analysis of overfitting in machine learning," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [71] X. Ying, "An overview of overfitting and its solutions," in *Journal of Physics: Conference Series*, vol. 1168, no. 2. IOP Publishing, 2019, p. 022022.
- [72] R. K. Nath, H. Thapliyal, and T. S. Humble, "A review of machine learning classification using quantum annealing for real-world applications," *SN Computer Science*, vol. 2, no. 5, p. 365, Jul 2021.
- [73] E. J. Crosson and D. A. Lidar, "Prospects for quantum enhancement with diabatic quantum annealing," *Nature Reviews Physics*, vol. 3, no. 7, pp. 466–489, Jul 2021.
- [74] B. Johnson, "Qiskit runtime, a quantum-classical execution platform for cloud-accessible quantum computers," *Bulletin of the American Physical Society*, 2022.

The QUBO formulations for RNA folding presented herein as well as that ref. [48] requires that each candidate stem of a query RNA be enumerated and mapped to a qubit. Further, the pairwise interactions between these candidate stems must also be enumerated. It is therefore useful to consider the upper limits to the number of stems and their pairwise interactions to estimate bounds for both classical pre-processing time-complexity and quantum spatial resource requirements.

We work with the extreme case of an RNA consisting in alternating bases of a canonical base pair along its sequence. For example: GCGCGC... or AUAUAU... We denote the sequence length of the query RNA  $N$ , the minimum stem size  $m$ , the function over these two variables which yields the number of candidate stems  $S(N, m)$ , and the number of pairwise interactions  $P(N, m)$ .

Consider a GCGCGC... sequence in the matrix formation of Figure 1, and suppose odd  $N$ . Then, there are  $N - 1$  anti-diagonals containing all 1s and  $N - 2$  diagonals containing all 0s. We refer to the 1-containing anti-diagonals as “anti-diagonals” from here.

By symmetry,  $(N - 1)/2$  of these anti-diagonals are with unique lengths, these lengths following the sequence: 1, 2, 3, ...,  $(N - 1)/2$ . If stems are limited to being of size  $m$  or larger, replace  $(N - 1)/2$  with  $(N - 1)/2 - m + 1$ . To find the total number of stems, we count over each anti-diagonal the number of sub-anti-diagonals of length  $\geq m$ . Allowing  $i$  to be the anti-diagonal length and  $j$  to be the stem length, both with lower bounds of  $m$ :

$$S(N, m) = 2 \sum_{i=m}^{\frac{N-1}{2}} \sum_{j=m}^i (i - j + 1) \quad (\text{A.1})$$

The summand  $i - j + 1$  tells simply how many stems of a length  $j$  fit within a anti-diagonal of length  $i$ .

Now, evaluating the inner sum, we find:

$$\sum_{j=m}^i (i - j + 1) = \frac{1}{2}(i - m + 1)(i - m + 2) \quad (\text{A.2})$$

Therefore, A.1 becomes with this substitution:

$$S(N, m) = \sum_{i=m}^{\frac{N-1}{2}} (i - m + 1)(i - m + 2) \quad (\text{A.3})$$

This then evaluates to:

$$S(N, m) = \sum_{k=1}^n k(k + 1) = \frac{1}{3}n(n + 1)(n + 2) \quad (\text{A.4})$$

Where  $k = i - m + 1$  and  $n = \frac{N-1}{2} - m + 1$ . Resubstituting, we find the analytic expression:

$$S(N, m) = \frac{1}{24} \left( N^3 + N^2(9 - 6m) + N(12m^2 - 36m + 23) + (-8m^3 + 36m^2 - 46m) + 15 \right) \quad (\text{A.5})$$

For  $N \gg m$ , this is  $\mathcal{O}(N^3)$ , which is to say that the number of stems mapping directly to qubits is cubic in sequence length, in the worst case.

Now, given that every stem may form a pseudoknot/overlap with any other stem, we enumerate the total number of possible pairwise interactions between stems we must check as:

$$P(N, m) = \binom{S(N, m)}{2} = \frac{S(N, m)(S(N, m) - 1)}{2} \quad (\text{A.6})$$

In the  $N \gg m$  approximation, this is  $\mathcal{O}(N^9)$ , such that the number of pairwise interaction grows to the ninth power of sequence length, in the worst case.

Having discussed the space-complexity of our pre-processing, we now consider the time-complexity, which again divides into two stages: enumerating over potential stems and enumerating pairwise interactions between stems.

To find stems, we iterate over all upper anti-diagonals of the RNA potential stems matrix. Straightforwardly, by considering all 1s as beginnings to potential stems and maintaining these in a stack until proven otherwise, we achieve a full enumeration in  $\mathcal{O}(N^2)$  time, again where  $N \gg m$ .

Once we are with a list of stems, finding the overlaps and pseudoknots of a set of stems is precisely an interval overlap problem, which may be solved in  $\mathcal{O}(S(N, m) \log S(N, m)) = \mathcal{O}(N^3 \log N)$  time using interval tree methods.

In sum, with the worst-case RNA, the number of potential stems is  $\mathcal{O}(N^3)$ , requiring  $\mathcal{O}(N^2)$  time to find, and the number of pairwise interactions between potential stems is  $\mathcal{O}(N^9)$ , requiring  $\mathcal{O}(N^3 \log N)$  time to find. This grants us a final pre-processing time-complexity of  $\mathcal{O}(N^3 \log N)$  and a quantum space-complexity of  $\mathcal{O}(N^3)$  (under the assumption of full connectivity of the qubit graph; in reality, extra qubits are required to achieve proper coupling, at least in the case of quantum annealing).

For the refs. [49], [50] QUBO, only stems of length 2 ( $m = 2$ ) are considered. Therefore, equation A.1 reduces to  $\mathcal{O}(N^2)$  in the worst case, and equation A.6 reduces to  $\mathcal{O}(N^4)$ . The pairwise interaction terms can therefore be computed in  $\mathcal{O}(N^2 \log N)$  time, such that the pre-processing stage of this algorithm is of  $\mathcal{O}(N^2 \log N)$  time-complexity, and a quantum space-complexity of  $\mathcal{O}(N^2)$  is required in the worst case.