

An infinite class of unsaturated rooted trees corresponding to designable RNA secondary structures

Jonathan Jedwab Tara Petrie Samuel Simon

23 September 2017 (revised 8 May 2020)

Abstract

An RNA secondary structure is designable if there is an RNA sequence which can attain its maximum number of base pairs only by adopting that structure. The combinatorial RNA design problem, introduced by Haleš et al. in 2016, is to determine whether or not a given RNA secondary structure is designable. Haleš et al. identified certain classes of designable and non-designable secondary structures by reference to their corresponding rooted trees. We introduce an infinite class of rooted trees containing unpaired nucleotides at the greatest depth, and prove constructively that their corresponding secondary structures are designable. This complements previous results for the combinatorial RNA design problem.

Keywords Combinatorial RNA design problem; secondary structure; nucleotide; unsaturated tree

1 Introduction

Ribonucleic acid (RNA) is a biomolecule which performs many roles in cellular organisms, including conveying genetic information, controlling protein synthesis, and catalyzing biological reactions. A strand of RNA comprises a chain of nucleotides, each of which is one of the nitrogenous bases guanine (G), cytosine (C), adenine (A), and uracil (U). The three-dimensional spatial configuration of an RNA strand affects its biological function [4]. A first approximation to this configuration is given by the *secondary structure*, namely the two-dimensional folding of the RNA strand onto itself under only the interaction of certain pairs of nucleotides in the sequences to form *base pairs*, as illustrated in Figure 1. The three allowable base pair types are $\{G,C\}$, $\{A,U\}$, and (less frequently) $\{G,U\}$.

Given a nucleotide sequence, the stability of the RNA secondary structure increases with the number of base pairs. Nussinov and Jacobson [10] showed that, given an RNA sequence of n nucleotides, a secondary structure with the maximum number of base pairs can be predicted in

Department of Mathematics, Simon Fraser University, 8888 University Drive, Burnaby BC V5A 1S6, Canada.

J. Jedwab and T. Petrie are supported by NSERC.

Email: jed@sfu.ca, tpetrie@sfu.ca, ssimon@sfu.ca

The results of this paper form part of the Master's thesis of T. Petrie [11], who presented them in part at the CanaDAM 2017 conference in Toronto, ON.




Figure 1: An RNA secondary structure in which base pairs are represented by grey points. Reproduced from [9] under Creative Commons Attribution License.

order n^3 time. Modern algorithms for secondary structure prediction are presented in [16, 8], for example, and are widely used in current software.

The inverse problem to secondary structure prediction is the *RNA design problem*: find an RNA sequence which adopts a target secondary structure in preference to all other secondary structures, under some specified energy model. There is an extensive literature on computational approaches to the RNA design problem as an optimization problem: see, for example, [3, 1, 14, 15, 7, 5, 4, 2]. Nonetheless, there is currently no known exact polynomial-time algorithm for solving the RNA design problem, and its complexity remains open [6]; a more general problem is known to be NP-hard [13].

In view of these difficulties, Haleš et al. [6] proposed the *combinatorial RNA design problem* as an idealized version of the RNA design problem: given a target secondary structure, find an RNA sequence which can achieve its maximum number of base pairs only by adopting the specified secondary structure, or else show that no such sequence exists. This restricted problem is intended to be sufficiently tractable to allow algorithmic insights that could apply to more sophisticated models. Haleš et al. [6] considered a number of energy models under which minimizing the total energy score for an RNA sequence corresponds to maximizing a weighted sum of base pair types. In the *Nussinov-Jacobson model*, the three base pair types $\{G,C\}$, $\{A,U\}$, $\{G,U\}$ are assigned an energy score α , β , γ , respectively, where $\max(\alpha, \beta) < \gamma < 0$; all other base pair types are assigned an energy score of ∞ in order to exclude them. The *Watson-Crick model* takes $\alpha = \beta = -1$ and $\gamma = \infty$, so that minimizing the total energy score is the same as maximizing the total number of $\{G,C\}$ and $\{A,U\}$ base pairs. We shall use the Watson-Crick model throughout this paper (except that in Section 3 we extend our main result to the Nussinov-Jacobson model with $\alpha = \beta = \gamma = -1$, which corresponds to maximizing the total number of $\{G,C\}$ and $\{A,U\}$ and $\{G,U\}$ base pairs).

We point out that the immediate practical implications of our results are limited in two important respects. Firstly, the class of “P-unsaturated floral trees” that we study (see Section 3) is highly restricted: it constrains multi-loops to have exactly three branches and to have no unpaired bases; it disallows internal loops and bulges; and it requires all unpaired bases in hairpin loops




Figure 2: Labelled secondary structure (left) and corresponding labelled tree representation (right)

to occur at the same nesting depth with respect to the exterior loop. Secondly, models for which minimization of the energy score is the same as maximization of the total number of base pairs have generally been found to have poor predictive ability, particularly for longer RNA sequences. This suggests it is unlikely that an RNA sequence identified by our results as a design for a target secondary structure would in practice adopt that secondary structure in preference to all other secondary structures. In Section 6, we propose as a topic for future study to attempt to extend the methods of this paper to provide predictions of greater practical relevance.

2 Terminology and previous results

In this section we define some terminology and then summarize the main results of [6]. Consider the sequence of nucleotides

A G A C C U U C C A C C A C C U U C C C A G U.

This sequence admits a secondary structure involving 7 base pairs of types $\{G,C\}$ and $\{A,U\}$, represented on the left side of Figure 2 by arcs that are all drawn above the nucleotide sequence. The right side of Figure 2 shows the corresponding representation of this labelled secondary structure in labelled tree form. We map from the labelled tree to the labelled secondary structure by visiting tree vertices in the order given by recursively visiting the root, and then (if the root is an internal vertex) the vertices of the left subtree, the vertices of the right subtree, and the root again; and then joining by an arc each pair of nucleotides that label the same tree vertex. We map in the reverse direction by representing the outermost arc as the tree root AU , the two arcs nested directly beneath the outermost arc as a left child GC and a right child CG of the tree root AU , and so on. In this way we associate each labelled secondary structure (a set of nucleotides and base pairs) with a labelled tree.

In Figure 3, we have removed the nucleotide labels to leave the underlying secondary structure and the corresponding tree representation. Nucleotide positions belonging to a base pair are represented in the tree by a two-dot vertex, while each unpaired nucleotide position is represented in the tree by a one-dot vertex. The children of each two-dot vertex v are all the vertices corresponding to nucleotides nested directly beneath the arc corresponding to v in the secondary structure.

We consider only secondary structures for which each nucleotide in the sequence is attached to at most one arc, and which are *pseudoknot-free*: arcs drawn above the nucleotide positions as in Figure 3, representing base pairs, do not intersect (because arc crossings lead to complex




Figure 3: Secondary structure (left) and corresponding tree representation (right)

constraints that often cannot be realized in three dimensions [12]). We shall assume that the first and last nucleotide of the sequence form a base pair, so that the graph representation is a rooted tree. (Otherwise it is a forest, in which case we can introduce a virtual root in order to produce a tree.) Each tree vertex will be labelled with either one or two dots; the only vertices that can be labelled with one dot are the leaves.

Now suppose that we are given only the target secondary structure R of Figure 3. The combinatorial RNA design problem then asks whether the nucleotide positions of the secondary structure can be labelled so that the resulting nucleotide sequence admits a unique set of $\{G,C\}$ and $\{A,U\}$ arcs of maximum size and so that this arc set equals R . We shall see in Theorem 3.3 that the answer for this example is yes, and that one such labelling is given by Figure 2.

We formalize these ideas as follows.

Definition 2.1 (Secondary structure). *A (pseudoknot-free) secondary structure applied to a sequence of nucleotide positions $\{1, 2, \dots, n\}$ is a set of arcs $\{i, j\}$ satisfying $1 \leq i < j \leq n$, for which no two arcs share the same endpoint and there are no two arcs $\{i, j\}$ and $\{k, \ell\}$ satisfying $i < k < j < \ell$.*

Definition 2.2 (Maximum-size arc set). *A secondary structure applied to a nucleotide sequence is a maximum-size arc set for the sequence if it contains the largest possible number of arcs (where only arcs of types $\{G,C\}$ and $\{A,U\}$ are allowed).*

Definition 2.3 (Design, designable). *An RNA sequence is a design if it admits a unique maximum-size arc set. A secondary structure R is designable if there is a design S whose (unique) maximum-size arc set equals R ; in that case, we say that S is a design for R .*

Combinatorial RNA Design Problem: Given a target secondary structure R , find a design for R or else show that R is not designable.

We introduce some further definitions in order to summarize prior results on the combinatorial RNA design problem. Following [6], a secondary structure in which all nucleotide positions are paired is *saturated*, and otherwise is *unsaturated*. The tree corresponding to a secondary structure is saturated when every leaf has two dots, and otherwise is unsaturated.




Figure 4: Saturated (left) and unsaturated (right) secondary structures

Theorem 2.4 (Haleš et al. [6, Result R4]). *The secondary structure R corresponding to a rooted saturated tree T is designable if and only if the root of T has at most four children and every other vertex of T has at most three children.*

The proof of Theorem 2.4 given in [6] is constructive, and gives many possible nucleotide labellings for the tree.

Theorem 2.5 (Haleš et al. [6, Result R5]). *The secondary structure R corresponding to a rooted unsaturated tree T containing a vertex with at least two 2-dot children and at least one 1-dot child is not designable.*

Theorem 2.5 shows that we cannot remove the constraint “saturated” from Theorem 2.4: for example, the secondary structure corresponding to the unsaturated tree



is not designable. Nonetheless, in this paper we generalize Theorem 2.4 by showing that an infinite class of unsaturated secondary structures is designable, thereby solving the combinatorial RNA design problem for a new class of structures.

Following [6], a *colouring* of a rooted tree assigns a colour drawn from the set $\{B, W, Y\}$ (where $B = \text{black}$, $W = \text{white}$, $Y = \text{grey}$) to each paired vertex of a rooted tree except the root. The colouring is *proper* if it obeys the rules specified in Table 1. (These rules constrain the root to have at most four children, and each other vertex to have at most three children.) The *level* of a vertex in such a coloured tree is the difference between the number of black vertices and the number of white vertices on the path starting at the vertex and ending at the root. A proper colouring is *separated* if no grey vertex shares a level with an unpaired vertex.

Vertex colouring	The multiset of colours of the children must be contained in the multiset
none (root)	$\{B, W, Y, Y\}$
B	$\{B, Y, Y\}$
W	$\{W, Y, Y\}$
Y	$\{B, W, Y\}$

Table 1: Rules for a proper colouring of the vertices of a rooted tree

Theorem 2.6 (Haleš et al. [6, Result R6]). *The secondary structure R corresponding to a rooted tree T which admits a separated proper colouring is designable.*

Theorem 2.6 is proved in [6] using an algorithm for constructing a design for the secondary structure R in which all leaves are labelled with A, and then the remaining black vertices are labelled with GC, white vertices with CG, and grey vertices with one of AU and UA.

In Section 3 we shall introduce the infinite class of “P-unsaturated perfect floral trees”, and state without proof our main result (Theorem 3.3) that these trees corresponds to designable secondary structures under the “natural labelling” of the structure with nucleotides. We then show that our main result complements the previous result of Theorem 2.6, in that each can be used to establish the designability of a secondary structure that the other cannot. We then establish some extensions of the main result, as Corollaries 3.4 and 3.5. Verifying that the natural labelling gives a maximum-size arc set is straightforward, but proving that the corresponding sequence is a design for the target structure appears to be much more difficult and subtle. In Section 4 we illustrate the proof techniques for Theorem 3.3 by reference to an extended example. We then prove the main result in full generality in Section 5. We conclude in Section 6 with some open questions.

3 Main result and some extensions

We wish to identify a new class of unsaturated trees that are designable (and which must necessarily avoid the class of unsaturated trees described in Theorem 2.5). In order to do so, we introduce Definition 3.1. Note that the *depth* of a vertex v in a rooted tree is the length of the path from the root to v .

Definition 3.1 (P-unsaturated). *A rooted tree is P-unsaturated if every leaf at the maximum depth is assigned exactly one dot and all other vertices exactly two dots.*

(The letter P in the name P-unsaturated is inspired by the requirement that all one-dot vertices of the rooted tree must be Pendant vertices.)




Figure 5: P-unsaturated binary tree

A *binary tree* is a rooted tree in which every vertex has at most two children. A *perfect binary tree* is a rooted tree in which every internal vertex has exactly two children and all leaves are at the same depth. We call a rooted tree *floral* if removing all vertices at the maximum depth leaves a binary tree, and *perfect floral* if removing all vertices at the maximum depth leaves a perfect binary tree. For example, Figure 6 shows a P-unsaturated perfect floral tree of height 3.

Theorem 3.3 (whose full proof is postponed until Section 5) states that the secondary structure corresponding to a P-unsaturated perfect floral tree is designable. The statement of this result involves the following assignment of nucleotides.

Definition 3.2 (Natural labelling). *Let T be a P-unsaturated perfect floral tree of height $n + 1$. The natural labelling of T labels all vertices at depth $n + 1$ with A if $n + 1$ is even, and with C if*




Figure 6: P-unsaturated perfect floral tree

$n + 1$ is odd; and labels the remaining vertices at even depths from left to right with alternating AU and UA , and those at odd depths from left to right with alternating GC and CG .

For example, the natural labelling of the tree in Figure 6 is shown in Figure 7.




Figure 7: Natural labelling of P-unsaturated perfect floral tree

Theorem 3.3 (Main Result). *The secondary structure R corresponding to a P-unsaturated perfect floral tree T is designable, and a design for R is given by the nucleotide sequence S corresponding to the natural labelling of T .*

Theorems 2.6 and 3.3 complement each other. The unsaturated tree shown in Figure 8, for example, admits the displayed separated colouring and so is designable by Theorem 2.6, but is not floral and so cannot be shown to be designable using Theorem 3.3 (nor by its Corollary 3.4 below). Conversely, consider a P-unsaturated perfect floral tree T of height 5, each of whose vertices at depth 4 has one child that is assigned one dot. We know this tree is designable by Theorem 3.3, but we now demonstrate with reference to Figure 9 that it cannot be shown to be designable using Theorem 2.6. Suppose, for a contradiction, that we can assign a separated proper colouring to T . By the separated colouring condition, no vertex at depth 4 can be coloured Y (else it would share a level with its unpaired child), and then by the colouring rules given in Table 1 each pair of sibling vertices at depth 4 must be coloured as one B and one W. The colouring rules applied to vertices at depth 3, then depth 2, and so on, then force all vertices at depths 1 and 3 to be coloured Y, and all vertices at depth 2 to be coloured in sibling pairs as one B and one W. But then the grey

vertices at depth 1 share the level 0 with the shaded unpaired vertex contained in the subtree of T shown in Figure 9, contradicting that the proper colouring is separated.




Figure 8: Unsaturated tree admitting a separated proper colouring (levels shown beneath vertices)




Figure 9: Subtree of proper colouring of P-unsaturated perfect floral tree of height 5

We show in Corollary 3.4 that we can relax the condition “perfect floral” in Theorem 3.3 to “floral”. The example structures in Figure 10 illustrate the ideas of the proof.

Corollary 3.4. *The secondary structure R corresponding to a P-unsaturated floral tree T is designable, and a design for R is given by the nucleotide sequence corresponding to a labelled subtree of the natural labelling of a perfect floral tree.*

Proof. Let T' be the smallest perfect floral tree containing T as a subtree; T' is necessarily P-unsaturated. Let τ' be the labelled perfect floral tree obtained by assigning nucleotides to T' according to the natural labelling. By Theorem 3.3, the labelled tree τ' yields a design S' of nucleotides whose (unique) maximum-size arc set corresponds to T' . Let τ be the labelled subtree of τ' whose unlabelled version is T , and let S be the associated sequence of nucleotides.

Colour the nucleotides and arcs that belong to S' but not to S red. The red substructure of S' consists of all nucleotides and corresponding arcs that must be pruned from the labelled perfect floral tree τ' to obtain the labelled floral tree τ . It follows that the secondary structure R

corresponding to T is a maximum-size arc set for S : if S admits a larger arc set than R then that arc set, together with the red arc set of S' , gives a larger arc set for S' (contradicting that the secondary structure corresponding to T' is a maximum-size arc set for S'). It also follows that S is a design for R : if S admits an alternative maximum-size arc set to R then that arc set, together with the red arc set of S' , gives an alternative maximum-size arc set for S' (contradicting that S' is a design). \square




Figure 10: Illustration of the proof of Corollary 3.4

Recall from Section 1 that we have assumed the Watson-Crick model (under which minimizing the total energy is equivalent to maximizing the number of $\{G,C\}$ and $\{A,U\}$ base pairs). We show in Corollary 3.5 that the result of Corollary 3.4 carries over to the Nussinov-Jacobson model with $\alpha = \beta = \gamma = -1$, under which minimizing the total energy is equivalent to maximizing the number of $\{G,C\}$ and $\{A,U\}$ and $\{G,U\}$ base pairs.

Corollary 3.5. *The secondary structure R corresponding to a P -unsaturated floral tree T is designable under the model in which minimizing the total energy is equivalent to maximizing the number of $\{G,C\}$ and $\{A,U\}$ and $\{G,U\}$ base pairs.*

Proof. Let R be the secondary structure corresponding to a P -unsaturated floral tree T . By Corollary 3.4, there is a design S for R under the Watson-Crick model (containing no $\{G,U\}$ base pairs), and S corresponds to a labelled subtree of the natural labelling of a perfect floral tree. Since all G

and U nucleotides are paired in the natural labelling, the number of arcs in R equals the number g of G nucleotides in S plus the number u of U nucleotides in S .

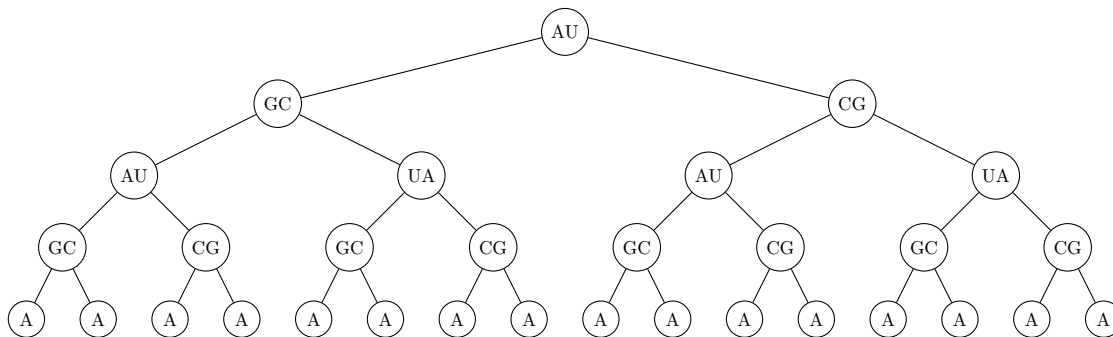
We shall show that S is also a design for R under the model which involves maximizing the number of $\{G,C\}$ and $\{A,U\}$ and $\{G,U\}$ base pairs, by showing that an alternative arc set to R that includes at least one $\{G,U\}$ base pair must contain fewer arcs than the $g + u$ contained in R . Let the number of $\{G,U\}$ base pairs in an arc set under this more general model be $n > 0$. Then there are at most $g - n$ nucleotides G remaining to form $\{G,C\}$ base pairs, and at most $u - n$ nucleotides U remaining to form $\{A,U\}$ base pairs, giving a total base pair count of at most $(g - n) + (u - n) + n < g + u$. \square

4 Illustration of proof techniques

Although the natural labelling of a P-unsaturated perfect floral tree follows a simple pattern, and it is relatively straightforward to show that Theorem 3.3 holds when T has small height, finding a general argument that applies to all tree heights appears to be very delicate. In Section 5 we will prove Theorem 3.3 according to the following outline (where R is the secondary structure corresponding to a P-unsaturated perfect floral tree T , and S is the nucleotide sequences corresponding to the natural labelling of T). Let R , S , and T be as given in Theorem 3.3. We must establish that S has a unique maximum-size arc set, and that this arc set equals R . We first show that all nucleotides in S that are unpaired in T (those at depth $n + 1$) must also remain unpaired in every maximum-size arc set. In view of this, we show how to reduce to the case of a saturated perfect binary tree. Our main tool is a running difference of strategically chosen subsets of the nucleotides in S , and a key insight is to focus on the parity of these running differences.

Having reduced a P-unsaturated perfect floral tree to a saturated perfect binary tree, we could then apply Theorem 2.4 to show that the secondary structure is designable. However, to keep our proof self-contained and to demonstrate the versatility of the method of running differences, we instead give our own proof in which we show that for every depth j of the reduced tree, nucleotides at depth j must all pair with one another, and that this forces a maximum-size arc set for S to equal R .

We now give a detailed illustration of the proof of Theorem 3.3, using the example of a P-unsaturated perfect binary tree T of height 4. (We use a binary tree in this illustration simply for convenience: the arguments readily extend to a P-unsaturated perfect floral tree.) The natural labelling τ of T is:



The corresponding sequence S (represented on two lines to assist reading) is:

A G A G A A C C A A G U U G A A C C A A G A C
 C A G A A C C A A G U U G A A C C A A G A G U

We shall show that S admits a unique maximum-size arc set, and that this arc set equals R . We introduce the following definition.

Definition 4.1 (Balanced set). *A subset of nucleotides from a nucleotide sequence forms a balanced set if it comprises either an equal number of Gs and Cs or an equal number of As and Us.*

For example, in Figure 11 the subset of G and C nucleotides shown boxed and coloured forms a balanced set \mathcal{B} within the nucleotide sequence S . The tag shown in Figure 11 above each nucleotide N (not in \mathcal{B}) is the difference between the number of Gs and the number of Cs occurring in \mathcal{B} to the left of N , namely the *running difference* with respect to \mathcal{B} . We use the corresponding definition for the running difference with respect to a balanced set of A and U nucleotides.




Figure 11: A balanced set of nucleotides and its corresponding running differences

In order to prove Theorem 3.3 for T , we iteratively define a succession of balanced sets of nucleotides, occurring as disjoint subsets of the sequence S , and use the running difference with respect to these balanced sets to constrain the possible arcs of a maximum-size arc set.

Suppose a maximum-size arc set M is applied to the sequence S . The labelled tree τ from which S is derived shows that S admits an arc set for which only A nucleotides are unpaired; therefore in M all G, C, and U nucleotides of the sequence must be paired.


Iteration 1:

S has an equal number of Gs and Cs, forming a balanced set \mathcal{B} which is boxed and coloured as shown in Figure 11. To achieve the maximum number of arcs, all elements of \mathcal{B} must pair with one another; this is indicated by the colouring, which will be retained in all subsequent iterations (whereas the boxes are a temporary notation used only in the current iteration). To avoid inducing an arc crossing, every arc in M joining uncoloured nucleotides (As and Us) must therefore enclose equally many of the two types of boxed nucleotides (Gs and Cs). Tag the uncoloured nucleotides with the running difference with respect to \mathcal{B} , as in Figure 11.

Then the uncoloured nucleotides with a given tag cannot pair with uncoloured nucleotides with a different tag, and cannot pair with coloured nucleotides (because those must all pair with one another). Therefore, the uncoloured nucleotides with a given tag can pair only with one another. This implies the weaker result that uncoloured nucleotides whose tags have the same parity can pair only with one another. (We have not yet shown that the uncoloured nucleotides with a given tag must all pair with one another, nor what the pairings must be.)

As a visual aid, we copy the boxing, colouring, and tagging applied to the nucleotides of S onto the corresponding labelled tree τ : the nucleotides belonging to \mathcal{B} (the Gs and Cs) are boxed and

coloured, and the uncoloured nucleotides (the As and Us) are tagged as above. (This does not assume the result we wish to prove, namely that applying a maximum-size arc set to S results in the labelled tree τ .)



As described, all uncoloured nucleotides whose tags are odd (this being the opposite parity to that of the nucleotides at depth 4) can pair only with one another. But these nucleotides form a new balanced set \mathcal{B}' (comprising A and U nucleotides), and so must all pair with one another to achieve the maximum number of arcs; this will be indicated by colouring in the next iteration. Remove the tags and boxes but retain the colouring.

Iteration 2:

Box and colour the nucleotides of the new balanced set \mathcal{B}' , as shown below.

A G A G A A C C A A G U U G A A C C A A G A C

C A G A A C C A A G U U G A A C C A A G A G U

To avoid inducing an arc crossing, every arc in M joining two nucleotides not in \mathcal{B}' must enclose an equal number of the two nucleotide types in \mathcal{B}' (boxed As and boxed Us). Tag the uncoloured nucleotides with the running difference with respect to \mathcal{B}' .


0 1 1 1 1 -1 -1 -1 -1

A G A G A A C C A A G U U G A A C C A A G A C

 1 1 1 1 -1 -1 -1 -1 0

C A G A A C C A A G U U G A A C C A A G A G U

Uncoloured nucleotides whose tags have the same parity can pair only with one another. The corresponding version of τ (again shown only as a visual aid) is:



All uncoloured nucleotides whose tags are even (this being the opposite parity to that of the nucleotides at depth 4) can pair only with one another. There are only two such nucleotides, namely the initial A and final U of the sequence S . They form a trivial new balanced set \mathcal{B}'' and so must pair with one another. Remove the tags and boxes but retain the colouring.

Iteration 3:

Box and colour the nucleotides of the new balanced set \mathcal{B}'' .

A G A G A A C C A A G U U G A A C C A A G A C

C A G A A C C A A G U U G A A C C A A G A G U

At this point, only the nucleotides at depth 4 are uncoloured, which is the stopping criterion. Since all coloured nucleotides must pair with one another, and all nucleotides at depth 4 (the uncoloured nucleotides) are As, the nucleotides at depth 4 must remain unpaired in M .

We will use this fact to determine M completely. Before doing so, we observe several properties that allow us to simplify the argument presented so far. Consider the balanced sets \mathcal{B} , \mathcal{B}' , \mathcal{B}'' in iterations 1 to 3. Observe that the nucleotides in S occurring at a given depth in τ are either all in, or all not in, the current balanced set. Furthermore the tags (with respect to this balanced set) of all nucleotides at a given depth share the same parity, and the parity switches between successive tagged depths. As a result, the labelled tree τ for each iteration 1 to 3 can be condensed as follows:





Figure 12: Simplified representation of iterations 1 to 3 using condensed tree

In the condensed tree, all vertices at a given depth in the original tree τ are represented by a single vertex labelled with either AU or GC according to the nucleotide types appearing at that depth in τ . The numbers placed at uncoloured depths give the parity of the running difference with respect to the current balanced set.


We now show that M must equal the arc set R corresponding to T . Since all nucleotides in S occurring at depth 4 in τ remain unpaired in M , we may remove them from τ to obtain the labelled height 3 saturated perfect binary tree $\hat{\tau}$, namely:



The corresponding nucleotide sequence \hat{S} is:

A G A G C C G U U G C C G A C C A G C C G U U G C C G A G U

We claim that all nucleotides in \hat{S} occurring at a given depth in $\hat{\tau}$ must pair with one another in M . By then considering arcs joining nucleotides in \hat{S} that occur at depth 0 to 3 in $\hat{\tau}$ in that order, and applying the condition of no arc crossings, we obtain the following nested sequence of arcs:



This arc set corresponds exactly to the arc set of M (because the nucleotides occurring at depth 4 in τ remain unpaired in M). Therefore M equals R , as required.

It remains to prove the claim. To show that all nucleotides in \hat{S} occurring at depth j ($0 \leq j \leq 3$) in $\hat{\tau}$ must pair with one another in M , we modify the previous process involving boxes, colouring, and tags so that the last remaining uncoloured depth is j (rather than 4, as in Figure 12). We illustrate this modified process for the case $j = 1$, using a condensed tree representation for $\hat{\tau}$. The argument for other values of j is similar.




Figure 13: Sequence of condensed trees for the case $j = 1$

In iteration 1 of Figure 13, box and colour the even depths (having opposite parity to $j = 1$). The nucleotides occurring at these depths belong to a balanced set \mathcal{B} , and so they must all pair with one another. Place at uncoloured depths the parity of the running difference with respect to \mathcal{B} . Nucleotides whose running difference has even parity (opposite to that of the nucleotides at depth $j = 1$) form a new balanced set \mathcal{B}' .

In iteration 2 of Figure 13, box and colour the depth at which the nucleotides of the new balanced set \mathcal{B}' occur. All nucleotides at coloured depths must pair with one another. The only depth remaining uncoloured is $j = 1$, all of whose nucleotides can therefore pair only with one another. The existence of $\hat{\tau}$ (in which every nucleotide is paired) then shows that the nucleotides at depth $j = 1$ must all pair with one another. This establishes the claim.

We shall prove Theorem 3.3 in Section 5 by generalizing this example.

5 Proof of main result (Theorem 3.3)

In this section we prove Theorem 3.3 according to the model described in Section 4.

We shall require two preliminary results. Observation 5.1 states that we cannot “jump over” a depth of a labelled tree as it is traversed to produce its corresponding nucleotide sequence. Lemma 5.2 counts the number of nucleotides at a fixed depth that occur in a subsequence of the nucleotide sequence corresponding to the natural labelling of a P-unsaturated perfect floral tree. The four cases of Lemma 5.2 are illustrated in Figure 14 using a tree of height 3, with nucleotides N_1 and N_2 marked in red in the tree representation (left) and corresponding arc representation (right), and the nucleotides occurring at depth h between N_1 and N_2 marked in green.

Observation 5.1. *Let S be the nucleotide sequence corresponding to a labelled tree. Then adjacent nucleotides in S occur either at the same depth or at depths that differ by one.*

Lemma 5.2. *Let S be the nucleotide sequence corresponding to the natural labelling of a P-unsaturated perfect floral tree of height $n + 1$. Let N_1 be a nucleotide at depth k . Suppose at least one nucleotide following N_1 in S is at depth ℓ , and let N_2 be the first such nucleotide.*

Further suppose the subsequence of S starting at N_1 and ending at N_2 passes through depth $h \notin \{k, \ell\}$ at least once. Then the total number of nucleotides at depth h occurring between N_1 and N_2 in S is

$$\begin{cases} 2 & \text{if } \ell = k \text{ and } h < k \\ 2^{h-k+1} & \text{if } \ell = k \text{ and } k < h < n + 1 \\ 1 & \text{if } k < h < \ell \\ \text{odd} & \text{if } \ell < h < k. \end{cases}$$

Proof. Consider the path through the tree corresponding to the subsequence of S that starts at N_1 and ends at N_2 .


Case 1: $\ell = k$ and $h < k$. The path passes through depth h exactly once in the upward direction and exactly once in the downward direction.

Case 2: $\ell = k$ and $k < h < n + 1$. The nucleotides N_1 and N_2 form a base pair (by definition of N_2), and the path traverses the entire subtree rooted at the vertex labelled by the base pair N_1N_2 .


Case 3: $k < h < \ell$. By definition of N_2 , the path passes through depth h exactly once.

Case 4: $\ell < h < k$. Consider the subtrees rooted at the vertices at depth h . Nucleotide N_1 occurs in one of these subtrees, say σ , and in S it lies between the two nucleotides labelling the root of σ . To reach N_2 from N_1 , the path passes through the second of these two nucleotides and then through both the nucleotides labelling the root of all the subtrees (if any) lying to the right of σ in the tree before reaching depth ℓ .


□




Case 1: $\ell = k$ and $h < k$ (illustrated by $\ell = k = 2$ and $h = 1$).



Case 2: $\ell = k$ and $k < h < n + 1$ (illustrated by $\ell = k = 1$ and $h = 2$).



Case 3: $k < h < \ell$ (illustrated by $\ell = 2$ and $k = 0$ and $h = 1$).



Case 4: $\ell < h < k$ (illustrated by $\ell = 0$ and $k = 2$ and $h = 1$).

Figure 14: Illustration of the four cases of Lemma 5.2

The crucial property which we shall use from Lemma 5.2 is that the nucleotide count is even if $\ell = k$ and is odd if $\ell \neq k$. Note that we have excluded from Lemma 5.2 the case where $\ell = k$ and $h = n + 1$, because the count of the nucleotides at depth $n + 1$ of the tree is then not determined by the conditions of the lemma.

We now prove Theorem 3.3, subject to two Claims whose proof will be given directly after the main argument. Let T be a P-unsaturated perfect floral tree of height $n + 1 \geq 1$, and let S be the nucleotide sequence corresponding to the natural labelling τ of T . Suppose a maximum-size arc set M is applied to the sequence S .

Claim 1: All nucleotides in S occurring at depth $n + 1$ in τ remain unpaired in M .

By Claim 1, we may remove the nucleotides at depth $n + 1$ in τ to obtain a labelled saturated perfect binary tree $\hat{\tau}$ of height $n \geq 0$. Let \hat{S} be the sequence corresponding to $\hat{\tau}$.

Claim 2: For each j satisfying $0 \leq j \leq n$, all nucleotides in \hat{S} occurring at depth j in $\hat{\tau}$ must pair with one another in M .

Apply Claim 2 to successive values of j starting from $j = 0$. The case $j = 0$ forces the leftmost and rightmost nucleotide of \hat{S} to be paired, giving the outermost arc shown below. Each successive case $j \geq 1$, together with the condition that there are no arc crossings, forces a further 2^j arcs of M to take the nested form shown below.




Figure 15: Forced secondary structure

Note that this repeated application of Claim 2 to \hat{S} establishes that \hat{S} admits the unique maximum-size arc set shown above, corresponding to a saturated perfect binary tree, and so \hat{S} is a design. Furthermore, by Claim 1 this arc set corresponds exactly to the arc set M . Therefore S is a design whose unique maximum-size arc set M equals the secondary structure R corresponding to T . It remains to prove Claims 1 and 2.

Proof of Claim 1. We use the following algorithm to iteratively tag the nucleotides of the labelled tree τ . This algorithm is modelled on the examples of Section 4 involving boxes, colouring, and tags. However, instead of indicating successive balanced sets by $\mathcal{B}, \mathcal{B}', \mathcal{B}'', \dots$ and using boxes to indicate the current balanced set, we update the current balanced set \mathcal{B} by replacing it with a set \mathcal{B}' that will be shown to be balanced.

Algorithm 1.

Input: Labelled tree τ and corresponding sequence S .

Step 1. Initialize the set \mathcal{B} to be all nucleotides in S occurring in τ at depths whose parity is opposite to that of $n + 1$.

Step 2. Colour all nucleotides in \mathcal{B} .

Step 3. If only the nucleotides at depth $n + 1$ are uncoloured, stop.

Step 4. Tag all uncoloured nucleotides by the running difference with respect to \mathcal{B} .

Step 5. Let \mathcal{B}' be the set of all tagged nucleotides whose running difference has opposite parity to that of the leftmost nucleotide at depth $n + 1$. Replace \mathcal{B} by \mathcal{B}' .

Step 6. Remove tags (but not the colouring) and go to Step 2.

Output: Sequence S with some nucleotides coloured.

We shall establish that the following six properties hold for Algorithm 1.

Property 1. *For each set \mathcal{B} , nucleotides at a given depth are either all in \mathcal{B} , or all not in \mathcal{B} .*

Property 2. *For each set \mathcal{B} , the nucleotides at depth $n + 1$ are not coloured at Step 2 (and so, if Step 4 is reached, these nucleotides receive a tag and therefore Step 5 is well-defined).*

Property 3. *For each set \mathcal{B} , the tags (with respect to \mathcal{B}) of all nucleotides at a given depth share the same parity.*

Property 4. *Parity alternates between successive tagged depths.*

Property 5. *Each set \mathcal{B} is balanced.*

Property 6. *Elements of each balanced set \mathcal{B} must all pair with one another in M .*

Assume Properties 1 to 6 hold, and consider how Algorithm 1 terminates at Step 3. Step 2 colours all nucleotides in set \mathcal{B} , which by Property 1 comprises all nucleotides at one or more depths. Therefore each application of Step 2 strictly decreases the number of uncoloured depths. This number is bounded below by 1 because, by Property 2, the nucleotides at depth $n + 1$ are never coloured. Therefore eventually this number is reduced to 1 at Step 2, at which point only the nucleotides at depth $n + 1$ are uncoloured and Algorithm 1 immediately terminates at Step 3. At this point, all nucleotides at other depths have been coloured (at some iteration of Step 2), and so belong to some balanced set by Property 5; by Property 6 these nucleotides must all pair with one another in M . The nucleotides at depth $n + 1$ all have the same type (A or C) and so cannot pair with one another; therefore they remain unpaired in M , establishing Claim 1.

To complete the proof of Claim 1, we must prove Properties 1 to 6.

We first show that Properties 1 to 4 hold for the initial set \mathcal{B} (defined in Step 1).

Property 1: The initial set \mathcal{B} is all nucleotides at depths whose parity is opposite to that of $n + 1$.

Property 2: All nucleotides at depth $n + 1$ lie outside \mathcal{B} and so are uncoloured.

Properties 3 and 4: From Step 4, the tagged nucleotides are exactly the uncoloured nucleotides and so occur at depths whose parity is the same as that of $n+1$. Let N_1 be an uncoloured nucleotide, and suppose another uncoloured nucleotide follows N_1 in S . Let N_2 be the first such uncoloured nucleotide. We may assume that at least one nucleotide from \mathcal{B} occurs between N_1 and N_2 in S (otherwise N_1 and N_2 occur at the same depth and receive identical tags with respect to \mathcal{B}).

Let N_1 occur at depth k and N_2 at depth ℓ . By Property 1 for \mathcal{B} , all nucleotides at depth ℓ remain uncoloured at Step 2, and so N_2 is the first nucleotide at depth ℓ following N_1 in S . Observation 5.1 implies that $\ell \in \{k-2, k, k+2\}$, and all nucleotides in \mathcal{B} occurring in the subsequence of S starting at N_1 and ending at N_2 occur at a single depth $h \in \{k-1, k+1\}$.

The possibilities for (ℓ, h) are $(k, k-1)$, $(k, k+1)$, $(k+2, k+1)$, and $(k-2, k-1)$. By definition of \mathcal{B} we have $h \neq n+1$ (which in fact implies that the case $(\ell, h) = (k, k+1)$ cannot occur). Then by Lemma 5.2, the number of nucleotides in \mathcal{B} occurring in the subsequence of S starting at N_1 and ending at N_2 is even if $\ell = k$ and is odd if $\ell \neq k$. Therefore the tags with respect to \mathcal{B} of all uncoloured nucleotides at a given depth share the same parity, and the parity alternates between successive tagged depths.

We now assume Properties 1 to 4 hold for the current set \mathcal{B} , and show that they hold for the set \mathcal{B}' defined in Step 5.

Property 1: By Properties 2 and 3 for \mathcal{B} , all nucleotides at depth $n+1$ are tagged with respect to \mathcal{B} and have the same parity. So \mathcal{B}' comprises all tagged nucleotides whose running difference with respect to \mathcal{B} has opposite parity to that of all the nucleotides at depth $n+1$. Then by Property 3 for \mathcal{B} , we see that Property 1 holds for \mathcal{B}' .

Property 2: nucleotides at depth $n+1$ do not belong to \mathcal{B}' (from Step 5) and so are not coloured (in Step 2).

Properties 3 and 4: We represent the application of Steps 2 and 4 to \mathcal{B}' by the following sequence of “condensed trees”.




Figure 16: Application of Steps 2 and 4 to new balanced set \mathcal{B}'

Figure 16(a) represents the tree prior to applying Step 5 to produce \mathcal{B}' ; only the depths containing nucleotides that are tagged with respect to \mathcal{B} (and which are necessarily uncoloured) are shown. (Previously coloured nucleotides may occur at other depths not shown.) By Property 1 for \mathcal{B} , all nucleotides at a given such depth may be represented by a single vertex. By Property 3 for \mathcal{B} , each such vertex may be assigned a single parity (p or $1 - p$), representing the parity of the tag shared by all nucleotides at that depth. By Properties 2 and 4 for \mathcal{B} , the nucleotides at depth $n + 1$ appear as a tagged vertex in the condensed tree, say with parity p , and the parity of vertices at successive tagged depths alternates. The set \mathcal{B}' therefore comprises all nucleotides at depths tagged with parity $1 - p$. These depths are coloured in Figure 16(b) to indicate that the nucleotides at these depths are coloured when Step 2 is applied to \mathcal{B}' .

Figure 16(c) shows the parity of the tags assigned to nucleotides when Step 4 is applied to \mathcal{B}' , as we now describe; this establishes Properties 3 and 4 for \mathcal{B}' . The uncoloured nucleotides are tagged in Step 4 by the running difference with respect to \mathcal{B}' . Let N_1 be an uncoloured nucleotide, and suppose another uncoloured nucleotide follows N_1 in S . Let N_2 be the first such uncoloured nucleotide. We may assume that at least one nucleotide from \mathcal{B}' occurs between N_1 and N_2 in S (otherwise N_1 and N_2 occur at the same depth and receive identical tags with respect to \mathcal{B}').

Let N_1 occur at depth k and N_2 at depth ℓ . By Property 1 for \mathcal{B}' , all nucleotides at depth ℓ remain uncoloured at Step 2, and so N_2 is the first nucleotide at depth ℓ following N_1 in S . Observation 5.1 shows that k and ℓ are either the same depth or are successive uncoloured depths. In both cases, all nucleotides in \mathcal{B}' occurring in the subsequence of S starting at N_1 and ending at N_2 occur at a single depth h (see Figure 16(b)).

In order to apply Lemma 5.2, we verify the appropriate conditions on k, h , and ℓ . By Property 1 for \mathcal{B}' , all nucleotides at depth h are in \mathcal{B}' (coloured at Step 2 for \mathcal{B}') and all nucleotides at depths k and ℓ are not in \mathcal{B}' (uncoloured). Therefore $h \notin \{k, \ell\}$, and by Property 2 for \mathcal{B}' we have $h \neq n + 1$. By Observation 5.1, if $k \neq \ell$ then either $k < h < \ell$ or $\ell < h < k$.

Then by Lemma 5.2, the number of nucleotides in \mathcal{B}' occurring in the subsequence of S starting at N_1 and ending at N_2 is even if $\ell = k$ and is odd if $\ell \neq k$. Therefore the tags with respect to \mathcal{B}' of all uncoloured nucleotides at a given depth share the same parity, and the parity alternates between successive tagged depths.

This completes the proof that Properties 1 to 4 hold for Algorithm 1.

We next show that Properties 5 and 6 hold for the initial set \mathcal{B} (defined in Step 1).

Property 5: The initial set \mathcal{B} comprises all G and C nucleotides if $n + 1$ is even, or all A and U nucleotides if $n + 1$ is odd, and so is balanced (by definition of the natural labelling).

Property 6: The labelled tree τ from which the sequence S is derived shows that S admits an arc set in which only nucleotides of a single type (A if $n + 1$ is even, or C if $n + 1$ is odd) are unpaired; therefore in the maximum-size arc set M all nucleotides from the initial set \mathcal{B} must pair with one another.

We now assume Properties 5 and 6 hold for the current set \mathcal{B} , and show that they hold for \mathcal{B}' .

Property 5: The initial application of Step 2 colours either all G and C nucleotides, or all A and U nucleotides. Since colouring is never removed, the set \mathcal{B}' comprises only (a subset of the) A and U nucleotides, or G and C nucleotides, respectively. Furthermore, the set \mathcal{B}' comprises all nucleotides occurring at some subset of depths not including depth $n + 1$ (see Figure 16(b)). Therefore, from the definition of the natural labelling, the nucleotides at this subset of depths form a new balanced set.

Property 6: The elements of the balanced set \mathcal{B} must all pair with one another, by Property 6.

So each arc joining two nucleotides not in \mathcal{B} must enclose an equal number of the two nucleotide types in \mathcal{B} (otherwise the arc will induce a crossing). Therefore nucleotides tagged in Step 4 with the same running difference with respect to \mathcal{B} can pair only with one another. In particular, nucleotides tagged in Step 4 whose running differences with respect to \mathcal{B} have the same parity can pair only with one another. It follows that the elements of the set \mathcal{B}' defined in Step 5 (whose tags with respect to \mathcal{B} all share the same parity) can pair only with one another.

Now the labelled tree τ from which the sequence S is derived shows that S admits an arc set in which only nucleotides of a single type (A or C) are unpaired; therefore in the maximum-size arc set M all nucleotides of the other three types must be paired. By Property 5 we know that \mathcal{B}' forms a balanced set, and we have shown that its elements can pair only with one another. Therefore the elements of \mathcal{B}' must all pair with one another.

This completes the proof that Properties 5 and 6 hold for Algorithm 1. \square

Proof of Claim 2. Fix j satisfying $0 \leq j \leq n$. We shall apply Algorithm 1 to $\hat{\tau}$ with $n+1$ replaced throughout by j . We shall show that Properties 1 to 6 hold for this modified algorithm, again with $n+1$ replaced by j . By the same argument as previously, these properties imply that all nucleotides at depths other than j must pair with one another in M . The nucleotides at depth j can therefore pair only with one another in M , and the existence of $\hat{\tau}$ (in which every nucleotide is paired) then shows that the nucleotides at depth j must all pair with one another in M , proving Claim 2.

In order to establish that Properties 1 to 6 hold for the modified algorithm, we highlight only the places in which the argument differs from that given previously.

Changes throughout:

Replace $n+1$ by j , τ by $\hat{\tau}$, and S by \hat{S} .

Changes to the proof that Properties 3 and 4 hold for \mathcal{B}' :

Replace Figure 16 by Figure 17.




Figure 17

Changes to the proof that Property 6 holds for the initial set \mathcal{B} :

The labelled tree $\hat{\tau}$ from which the sequence \hat{S} is derived shows that \hat{S} admits an arc set in which all nucleotides are paired; therefore in the maximum-size arc set M all nucleotides from the initial set \mathcal{B} (comprising all G and C nucleotides or all A and U nucleotides) must pair with one another.

Changes to the proof that Property 6 holds for \mathcal{B}' :

As previously, the elements of \mathcal{B}' can pair only with one another. The labelled tree $\hat{\tau}$ from which the sequence \hat{S} is derived shows that \hat{S} admits an arc set in which all nucleotides are paired; therefore in the maximum-size arc set M all nucleotides must be paired. Therefore the elements of \mathcal{B}' must all pair with one another.

□

Claims 1 and 2 have now been established and so the proof of Theorem 3.3 is complete.

6 Open Questions

- Q1. Theorem 3.3 shows the existence of at least one design for the secondary structure corresponding to a P-unsaturated perfect floral tree T . How many such designs are there?
- Q2. Theorems 2.4 and 2.5 identify obstructions whose presence in a rooted tree prevents the corresponding secondary structure from being designable. Can other obstructions be identified? Can it be shown that the number of minimal obstructions is finite?
- Q3. Theorem 2.5 shows that the secondary structures corresponding to rooted unsaturated trees are not designable in general, although Corollary 3.4 shows that those corresponding to a subclass of rooted unsaturated trees are. Is the secondary structure corresponding to a rooted unsaturated binary tree always designable?
- Q4. In Section 3 we showed that Theorems 2.6 and 3.3 can each be used to establish the designability of a secondary structure that the other cannot. Can these criteria be subsumed into a more general result?
- Q5. Haleš et al. [6] proposed the combinatorial RNA design problem as an idealized version of the RNA design problem, in the expectation that it would produce algorithmic insights applicable to more sophisticated models (see Section 1). As a next step, how would the results of this paper change if, instead of seeking sequences which admit a unique maximum-size arc set, we allowed sequences to admit at most two maximum-size arc sets?
- Q6. As discussed in Section 1, the practical applicability of our results is limited by the restricted class of sequences and the energy model used. To what extent can the methods of this paper be extended to investigate the designability of wider classes of sequences using more realistic energy scores?

Acknowledgements

Our research was carried out on the unceded Aboriginal territories of the Coast Salish people, including the Musqueam, Tsleil-Waututh, and Squamish First Nations. Thanks to Ladislav Stacho

for introducing us to this problem, to Marni Mishna for detailed comments on the manuscript, and to Stefan Hannie for interesting discussions. We appreciate Yann Ponty’s valuable assistance in giving helpful feedback on our results and methods, providing careful responses to our questions, and suggesting Corollary 3.5 and the comparison in Section 3 between Theorems 2.6 and 3.3. We are grateful to the reviewers for their very careful reading of the paper and for several constructive suggestions for improvement.

References

- [1] R. Aguirre-Hernández, H.H. Hoos, and A. Condon. Computational RNA secondary structure design: empirical complexity and improved methods. *BMC Bioinformatics*, 8(1):34, 2007.
- [2] É. Bonnet, P. Rzażewski, and F. Sikora. Designing RNA secondary structures is hard. 2017. arXiv:1710.11513.
- [3] A. Busch and R. Backofen. Info-RNA—a fast approach to inverse RNA folding. *Bioinformatics*, 22(15):1823–1831, 2006.
- [4] A. Esmaili-Taheri, M. Ganjtabesh, and M. Mohammad-Noori. Evolutionary solution for the RNA design problem. *Bioinformatics*, 30(9):1250–1258, 2014.
- [5] J.A. Garcia-Martin, P. Clote, and I. Dotu. RNAiFold: A constraint programming algorithm for RNA inverse folding and molecular design. *Journal of Bioinformatics and Computational Biology*, 11:1350001, 2013.
- [6] J. Haleš, A. Héliou, J. Maňuch, Y. Ponty, and L. Stacho. Combinatorial RNA Design: Designability and structure-approximating algorithm in Watson-Crick and Nussinov-Jacobson energy models. *Algorithmica*, 79(3):835–856, 2017.
- [7] R.B. Lyngsø, J.W.J. Anderson, E. Sizikova, Am. Badugu, T. Hyland, and J. Hein. FRNAkenstein: multiple target inverse RNA folding. *BMC Bioinformatics*, 13(1):260, 2012.
- [8] D.H. Mathews, J. Sabina, M. Zuker, and D.H. Turner. Expanded sequence dependence of thermodynamic parameters improves prediction of RNA secondary structure. *Journal of Molecular Biology*, 288(5):911–940, 1999.
- [9] M.E. Nebel, A. Scheid, and F. Weinberg. Random generation of RNA secondary structures according to native distributions. *Algorithms for Molecular Biology*, 6:24, 2011. Available online: <https://almob.biomedcentral.com/articles/10.1186/1748-7188-6-24>.
- [10] R. Nussinov and A.B. Jacobson. Fast algorithm for predicting the secondary structure of single-stranded RNA. *Proc. National Academy of Sciences*, 77(11):6309–6313, 1980.
- [11] T.J. Petrie. The combinatorial RNA design problem for binary trees. Master’s thesis, Simon Fraser University, 2017. Available online: <http://summit.sfu.ca/item/17222>.
- [12] Y. Ponty. Personal communication, August 2016.
- [13] M. Schnall-Levin, L. Chindelevitch, and B. Berger. Inverting the Viterbi algorithm: an abstract framework for structure design. In *Proceedings of the 25th International Conference on Machine Learning*, pages 904–911. ACM, 2008.

- [14] A. Taneda. MODENA: a multi-objective RNA inverse folding. *Advances and Applications in Bioinformatics and Chemistry*, 4:1–12, 2011.
- [15] J.N. Zadeh, B.R. Wolfe, and N.A. Pierce. Nucleic acid sequence design via efficient ensemble defect optimization. *Journal of Computational Chemistry*, 32(3):439–452, 2011.
- [16] M. Zuker and P. Stiegler. Optimal computer folding of large RNA sequences using thermodynamics and auxiliary information. *Nucleic Acids Research*, 9(1):133–148, 1981.