



HAL
open science

RNA inverse folding can be solved in linear time for structures without isolated stacks or base pairs

Théo Boury, Laurent Bulteau, Yann Ponty

► **To cite this version:**

Théo Boury, Laurent Bulteau, Yann Ponty. RNA inverse folding can be solved in linear time for structures without isolated stacks or base pairs. 2024. hal-04589901v1

HAL Id: hal-04589901

<https://hal.science/hal-04589901v1>

Preprint submitted on 27 May 2024 (v1), last revised 5 Jun 2024 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

1 RNA inverse folding can be solved in linear time 2 for structures without isolated stacks or base pairs

3 Théo Boury 

4 Laboratoire d'Informatique de l'Ecole Polytechnique (LIX; UMR 7161), Institut Polytechnique de
5 Paris, France

6 Laurent Bulteau  

7 LIGM, CNRS, Université Gustave Eiffel, France

8 Yann Ponty¹

9 Laboratoire d'Informatique de l'Ecole Polytechnique (CNRS/LIX; UMR 7161), Institut
10 Polytechnique de Paris, France

11 — Abstract —

12 Inverse folding is a classic instance of negative RNA design which consists of finding a sequence that
13 uniquely folds into a target secondary structure with respect to energy minimization. A breakthrough
14 result of Bonnet *et al* shows that, even in simple base pairs-based (BP) models, the decision version
15 of a mildly constrained version of inverse folding is NP-hard.

16 In this work, we show that inverse folding can be solved in linear time for every target structure
17 that contains no isolated BP and no isolated stack BP. (i.e. when helices have all sizes more than
18 $h = 3$). For structures featuring shorter helices, our linear algorithm is no longer guaranteed to
19 produce a solution, but still does so for a large proportion of instances.

20 Our approach introduces a notion of modulo m -separability, generalizing a property introduced
21 by Hales *et al*. Separability is a sufficient condition for the existence of a solution to the inverse
22 folding problem. We show that, for any input secondary structure of length n , a modulo m -separated
23 sequence can be produced in time $\mathcal{O}(n \cdot 2^m)$ anytime such a sequence exists. Meanwhile, we show
24 that any structure such that $h = 3$ is either trivially non-designable, or always admits a modulo-2
25 separated sequence ($m = 2$). Solution sequences can thus be produced in linear time, and even be
26 uniformly generated within the set of modulo-2 separable sequences.

27 **2012 ACM Subject Classification** Applied computing → Molecular structural biology

28 **Keywords and phrases** RNA structure, String Design, Parameterized Complexity, Uniform Sampling

29 **Digital Object Identifier** [10.4230/LIPIcs.WABI.2024.XXX](https://doi.org/10.4230/LIPIcs.WABI.2024.XXX)

30 **Supplementary Material** The source code for the project are on [GitLab:Linear-BP-Design](https://gitlab.com/Linear-BP-Design).

¹ Both second and third authors should be considered as corresponding authors



1 Introduction

RNA inverse folding is a fascinating algorithmic problem which, given a target secondary structure T , consists of designing one or several sequences, all of which should uniquely fold into the target T according to a reference folding prediction algorithm. Considering a folding prediction algorithm as a mathematical function $\Phi : \{A, C, G, U\}^* \rightarrow \mathcal{S} \cup \{\perp\}$ mapping an RNA sequence to a unique predicted structure (or \perp if equally likely alternatives exist), inverse folding can be abstracted as the search for a preimage $w \in \Phi^{-1}(T)$ of the target structure T . This naturally generalizes into a variety of design tasks which, given a predictive algorithm implementing a function Φ , aim to create one or multiple instances predicted to behave in a certain way. Such a formulation is, in general, overly broad (*e.g.* it encompasses the concept of one-way functions in cryptography) to inspire reasonable hopes for a general solution. Still, a restriction of the inverse problem to certain types of computable functions/algorithms (*e.g.* amenable to dynamic programming) appears realistic and generally relevant to (synthetic) biology, yet poorly studied to this day.

In the specific case of RNA, despite being the object of substantial attention since its formal introduction in the early 1990s [8], the complexity of RNA inverse folding has remained elusive for almost three decades. A generalization of RNA inverse folding, including the energy model as part of the input, was shown to be NP-hard by Schnall-Levin *et al* [18]. However, their reductions critically relied on (ab)using the energy model to encode a 3SAT instance, leaving the hardness of the problem largely open for a fixed energy model. The classic complexity of inverse folding was only settled, in 2018, when Bonnet *et al* [2] finally showed the NP-hardness of RNA folding in a classic base pairs maximization setting. Such computational intractability (retrospectively) legitimizes a very large quantity of heuristic or exponential-time methods, based on local search [8, 3, 1, 22, 16], bio-inspired metaheuristics [11, 4, 9, 12], global sampling [15, 21], constraint programming [5, 7] and, more recently, neural networks-inspired generative models [17].

In parallel to complexity studies, Hales *et al* [6] revisited the problem from a structural angle, attempting to characterize designable or undesignable families of secondary structures. The authors showed that saturated structures, having all positions paired, are designable if and only if their multiloop degrees do not exceed 4. They also introduced a notion of separability, a sufficient, yet not necessary in general, condition for a sequence to be a design for a given target. This notion allowed them to show that any target structure either features an occurrence of a locally-undesignable motif $\{m_{3\bullet}, m_5\}$, or can always be transformed into a separable structure by adding at most one base pair per helix. More strikingly, they proposed linear-time algorithms for producing a single solution for each characterized class of designable structures, painting a – puzzling – contrasted picture of general hardness (as per Bonnet *et al* [2]) and practical facility for inverse folding.

In this work, we further those studies and show that, while conceptually simpler, the existence of a separated design for a given structure remains NP-hard. Conversely, any structure with helices of length greater than 3 base pairs is either trivially undesignable (*i.e.* contains $\{m_{3\bullet}, m_5\}$), or separable and can be designed in linear-time. This constraint is relevant to the objectives of RNA design, as targeted secondary structures are typically stable and tend to avoid shorter – unstable – helices. This result hinges on the introduction of a modulo m version of separability, coinciding with general separability whenever $m \geq n/2$, for which we give a Fixed-Parameter Tractable (FPT) algorithm running in time $\mathcal{O}(n \cdot 2^m)$. We proved that this algorithm solves all instances with minimal helix lengths of 3 BPs when invoked with $m = 2$ and, even in this restricted setting, solves many instances with shorter

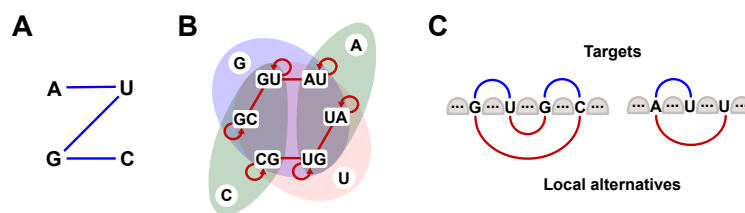


Figure 1 Local design rules. Base pair compatibility graph (A) and incompatibility graph for base pairs and unpaired nucleotides occurring within a loop (B): Connected base pairs, when jointly occurring within a loop of the target structure, can refold to form a local, an alternative structure having same number of base pairs as the target (C, left). Unpaired nucleotides may also interfere with some (A or C) or every (G or U) base pairs, leading to local alternatives (C, right).

78 helices in practice. Based on an unambiguous dynamic programming, our algorithm can be
 79 adapted into a random generator of separated designs. Finally, we show through empirical
 80 studies that separated sequences, despite being only guaranteed to constitute designs with
 81 respect to base pair maximization, are also likely to represent designs in the more realistic
 82 Turner energy model, and are far superior in this setting than compatible sequences.

2 Problem statement, definitions, and prior work

84 Algorithmically, RNA can be abstracted as a nucleotide sequence, *i.e.* a string $w \in$
 85 $\{A, C, G, U\}^n$ where n denotes the length of ω . Given a length n , a (non crossing/pseudoknot-
 86 free) secondary structure is a set $T \subset [1, n]^2$ consisting of base pairs such that:

- 87 ■ Each position in $[1, n]$ is involved in at most one base pair;
- 88 ■ Base pairs in T are pairwise non-crossing: $\forall (i, j) \neq (k, l) \in T, i < k$, either $i < k < l < j$
 89 or $i < j < k < l$.

90 The set \mathcal{S}_w of secondary structures compatible with an RNA sequence w is defined as:
 91 $\mathcal{S}_w := \{\text{Secondary structure } T \mid \forall (i, j) \in T, \{w_i, w_j\} \in \{\{G, C\}, \{A, U\}, \{G, U\}\}\}$.

92 Without loss of generality, a secondary structure can be represented as a tree $T =$
 93 $(V(T), E(T))$, whose nodes $V(T)$ are in bijection with base pairs (internal nodes²) and
 94 unpaired regions (leaves), and whose edges represent the inclusion of base pairs. Given a
 95 node $v \in V(T)$, we denote by $\text{parent}(v)$ the parent of v in T , and by $\text{children}(v)$ the list of
 96 children of v in T . A *loop* is the subtree restricted to node and its (direct) children. The tree
 97 is rooted in a special **Root** node, associated with the whole sequence interval. An *helix* of
 98 length ℓ of the tree is a maximal path v_1, \dots, v_ℓ of base pair nodes such that each v_i with
 99 $i < \ell$ has a single child v_{i+1} (no leaf attached). A helix of length 1 is an *isolated base pair*. A
 100 helix of length 2 is an *isolated stack*. We define h as the minimum length over all helices of T .
 101 As the target tree is always explicit and unmodified through proofs and algorithms we do
 102 not specify it explicitly in the notations.

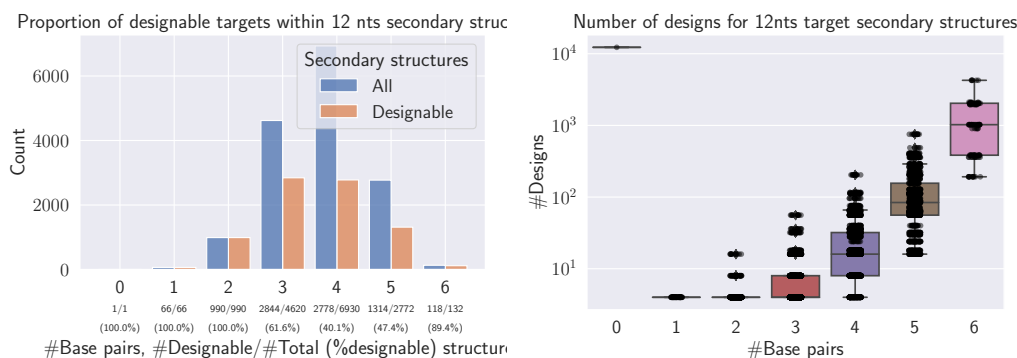
103 RNA inverse folding starts from a target secondary structure T , and attempts to construct
 104 a sequence $\omega \in \{A, C, G, U\}^n$ whose only base-pair maximizing secondary structure is T .

► Problem 1. INVERSE-FOLDING_{BP}

106 **Input:** Target secondary structure T , sequence length n

107 **Output:** Sequence $w \in \{A, C, G, U\}^n$ satisfying both:

² Base pairs may also be leaves of the tree when involving consecutive positions, which happens rarely in practice. We thus qualify as *internal node* any node in bijection with a base pair.



■ **Figure 2 Exhaustive designability analysis of 12nts RNA sequences/structures.** (Left) For a minimum base pair span of $\theta = 0$, there exists 15 511 secondary structures over 12 nucleotides, of which little over half (8 111) admits at least a solution to the inverse folding problem. (Right) The number of valid solutions varies substantially between targets and appears to depend on the number of base pairs. Overall, out of the 16 777 216 RNA sequences of length 12, only 399 348 ($\approx 2.4\%$) represent a valid design for some structure.

- 108 ■ *Compatibility with target structure:* $T \in \mathcal{S}_w$;
- 109 ■ *Uniqueness of the target as the optimal fold for the sequence:* $\forall T' \in \mathcal{S}_w, T' \neq T, |T'| < |T|$.
- 110 *or \perp if no such sequence exists.*

111 Nevertheless, INVERSE-FOLDING_{BP}, mildly extended to allow further restrictions on individual
 112 sequence positions, was shown to be NP-hard by Bonnet *et al* [2].

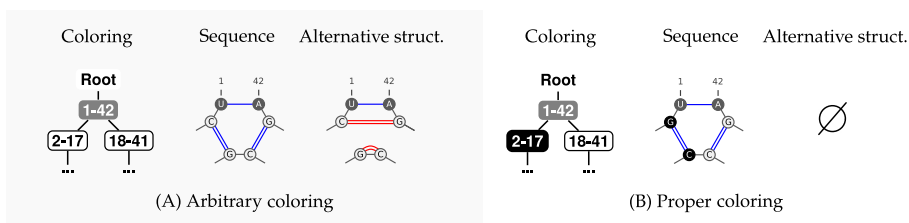
113 A sequence is called a design for a structure T if it represents a solution to the inverse
 114 folding problem for the input T . Note that the uniqueness condition can be tested in
 115 polynomial time using a variant of the Nussinov algorithm [13, 6]. In addition to showing
 116 that INVERSE-FOLDING_{BP} is in P, such an algorithm enables, for moderate sequence lengths,
 117 a systematic folding of all sequences in order to characterize the set of structures admitting
 118 a solution. For instance, Figure 2 shows that, while only 2.4% of RNA sequences of length
 119 12 represent a design for some target, roughly half of the secondary structure admits at least
 120 one solution sequence, and ≈ 49 on average, for the inverse folding problem.

121 We remind that, as noted by Halès *et al* [6], two key motifs are not designable in a *base*
 122 *pair maximization* setting:

- 123 ■ The m_5 motif consists of 5 base pairs occurring on the same loop (not counting the Root).
 124 No sequence can be designed for such a motif, since exposing 5 base pairs on a loop
 125 always allows for local refolding to have the same number of base pairs. This follows from
 126 the inspection of Figure 1, where the largest set of mutually compatible base pairs clearly
 127 has cardinality 4;
- 128 ■ The $m_{3\bullet}$ motif consists of 3 base pairs (excluding the Root) and at least one unpaired
 129 position. Indeed, as shown in Figure 1, the presence of an unpaired nucleotide either
 130 forbids the co-occurrence of any adjacent base pair (G or U), or only allows three (C or
 131 A). Since at most two of those base pairs can co-occur in a successful loop design, $m_{3\bullet}$ is
 132 not designable.

133 Any occurrence of these structures (or of any other undesignable structure, *cf* [20]) as a
 134 subgraph of an instance makes the instance undesignable.

XXX:4 Exact linear-time RNA design for min Helix length 3



■ **Figure 3 A proper coloring is necessary towards design.** In (A), having two \circ children implies that the sequence derived from this coloring features a motif where G and C can reconfigure locally. In that case, they form an alternative structure that contains the same number of base pairs. Conversely, in (B), the proper coloring ensures that locally no alternative of equal (or better) energy exists by forcing some consecutive incompatibilities.

135 2.1 Inverse folding as a tree coloring problem

136 We start by reminding the coloring framework introduced by Halès *et al* [6].

137 ► **Definition 1 (Coloring).** A coloring of a (secondary structure) tree T is a function $\chi : V(T) \rightarrow \{\bullet, \circ, \emptyset\}$ associating a color to each node (except the root and the leaves which always get \emptyset).

140 A coloring of a tree T typically induces multiple RNA sequences that are compatible with, but not guaranteed to fold into, the given secondary structure through letters assignment rules. Namely, in any sequence w derived from a coloring χ , we have for each $(i, j) \in T$:

- 143 ■ If $\chi((i, j)) = \bullet \rightarrow (w_i, w_j) = (G, C)$;
- 144 ■ If $\chi((i, j)) = \circ \rightarrow (w_i, w_j) = (C, G)$;
- 145 ■ If $\chi((i, j)) = \bullet \rightarrow (w_i, w_j) \in \{(A, U), (U, A)\}$.

146 For \bullet nodes, the freedom in choosing (A, U) or (U, A) depends on the context: the choice may be unconstrained (*e.g.* when isolated within a helix), or forced (*e.g.* when two gray nodes are involved in a multiloop or stack). However, this property will only impact the number of sequences associated with the coloring, but bears no consequence on the existence of a solution to INVERSE-FOLDING_{BP}, since the problem asks for the production of a single sequence.

152 Denote by \bar{c} the inverse of a color c , defined as $\bar{\circ} = \bullet$, $\bar{\bullet} = \circ$ and $\bar{\bullet} = \bullet$.

► **Definition 2 (Proper Coloring).** A coloring χ is proper when, for each node $v \in V(T)$, the vector of colors C , assigned to the node and its children, respects the following constraints:

$$|C|_{\bullet} \leq 1, |C|_{\circ} \leq 1 \text{ and } |C|_{\bullet} \leq 2 \text{ with } C := [\bar{\chi(v)}].[\chi(v') \mid v' \in \text{children}(v)].$$

153 These conditions must also hold for the colorless Root, but with C being restricted to the colors of children(Root).

155 In terms of RNA design, the proper condition is necessary for an associated sequence to be a solution to inverse folding. Indeed, any coloring that is not proper will be associated with sequences that can be locally reconfigured, this without losing any base pair (see Figure 3 for an example).

159 ► **Definition 3 (Levels).** Given a coloring χ of a tree T , the level $L : V(T) \rightarrow \mathbb{Z}$ of a node v is $L(v) := |p|_{\bullet} - |p|_{\circ}$ where p denotes the shortest node sequence from parent(v) to Root.

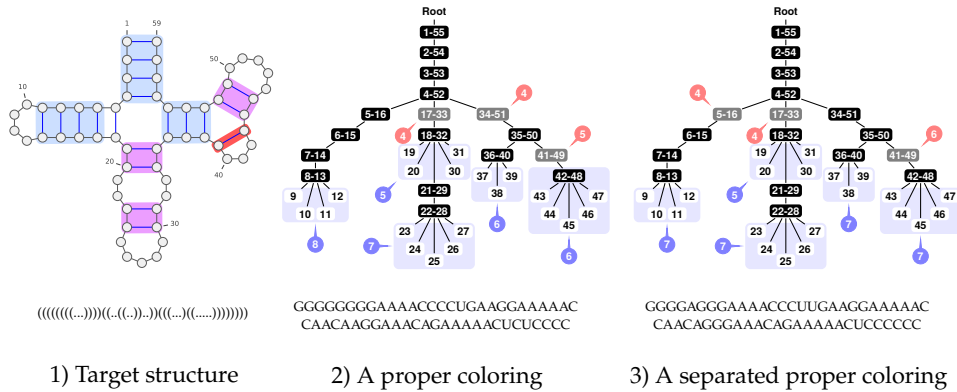


Figure 4 1) 2D and dot-bracket representations of a secondary structure. Helices of sizes respectively 1, 2 and 3 are represented in light red, purple and blue. 2) Same secondary structure as a tree. The tree is colored with a proper non-separated coloring as the level of leaf 19 is the same as the level of the \bullet node 34-51. Note that even if the current coloring is non-separated, it does not mean that the tree is not separable: a direct retrieving of some corresponding nucleotide letters leads to a design in that case. 3) Same secondary structure as a tree but the coloring is separated (necessarily proper). Thus, this coloring yields one or multiple designs (depending on the choice of AU or UA with the \bullet nodes). Even more, this coloring is 2-separated as every level of the leaves is odd while every level of the \bullet nodes is even.

161 On an RNA level, the concept of level helps categorize, and possibly control, the set
 162 of alternative structures to the target. Indeed, consider a sequence w generated from a
 163 coloring χ . First remark that, in order for an alternative structure to be competitive, every
 164 occurrence of C must be paired. Whenever two positions i and j interact to form a base pair,
 165 it can be shown that the inner interval $]i, j[$ interval contains $L(i) - L(j)$ more G than C.
 166 Meanwhile the outermost interval $[1, i[\cup]j, n]$ features the opposite imbalance ($L(i) - L(j)$
 167 more C than G). In other words, any structure that contains a base pair $(i, j) \notin T$ already
 168 has $2 \times |L(i) - L(j)|$ fewer base pairs than the target structure. Thus only structures made
 169 of pairs (i, j) such that $L(i) = L(j)$ need to be considered as viable alternatives to T . This
 170 property can be exploited as a design principle, as formalized by the following property.

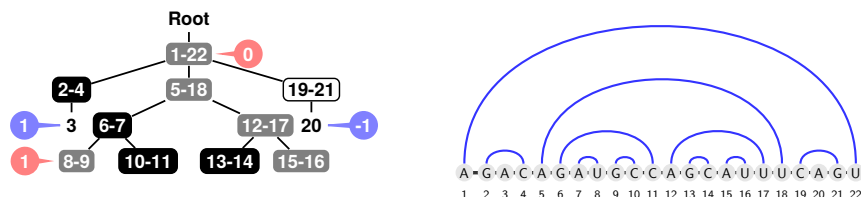
► **Definition 4** (Separated coloring). *A coloring χ is separated for a target T if and only if it is proper and the levels of \bullet -colored nodes and leaves do not overlap:*

$$\{L(v) \mid \chi(v) = \bullet\} \cap \{L(v) \mid v \text{ is a leaf}\} = \emptyset$$

171 This immediately suggests a design strategy that associates A to unpaired positions and
 172 assigns \bullet and \circ colors such that \bullet nodes end up as different levels as the leaves. Indeed,
 173 in this setting, Hales *et al* [6] showed that the proper coloring of a saturated structure
 174 (without unpaired position) yields a sequence that uniquely folds with respect to base pair
 175 maximization. It follows that a competitive/alternative structure may only result from a base
 176 pair $(i, j) \notin T$, a position of which is a \bullet node while the other is a leaf. Ensuring that all \bullet
 177 nodes and leaves are found at different levels is thus sufficient to guarantee the designability
 178 of T , *i.e.* the existence of a solution to this instance of the inverse folding problem.

179 More generally, we say that a target secondary structure T is *separable* if there exists a
 180 coloring χ such that χ is separated for T . We recall the main results of Halès *et al* [6] here.

181 ► **Theorem 1** (Separable \implies Designable (Halès *et al*, 2017)). *If a tree/secondary structure*
 182 *T is separable, then T is designable.*



■ **Figure 5 Designability does not imply separability.** Left: A target structure that does not admit any separated coloring instance. Note that the coloring χ shown here puts the \bullet node 8-9 and the leaf 3 both at level 1. Right: Sequence w compatible with the coloring χ , which provably admits T as its single base pair-maximization structure (i.e. w is a design for T).

183 Moreover, given a separated coloring, an RNA sequence that uniquely folds into T , i.e. a
 184 solution to the inverse folding problem, can be found in linear time.

185 ► **Remark 2.** Note that any design sequence w , generated through a separated coloring,
 186 avoids any alternative structure featuring GU base pair(s). Indeed, every G and C need to be
 187 paired to achieve the number of base pairs featured in the MFE. Meanwhile, the formation
 188 of any GU base pair, leaves one C and one A unpaired, resulting in the overall loss of at least
 189 one base pair. Structures featuring GU base pairs can thus be safely ignored.

190 3 Separability: Intrinsic and computational limits

191 Despite utilizing separability to explore a design of approximative instances, the work of Halès
 192 *et al* [6] left open the complexity of searching for a separated coloring, as well as the existence
 193 of designable, yet non-separable, structures. An exhaustive search for all structures with
 194 up to 12 bases, summarized in Figure 2, shows that for such small instances, all designable
 195 instances are separable.

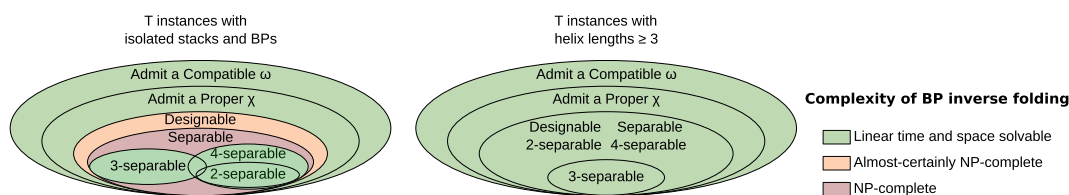
196 However, we show that non-separable designable instances can be constructed.

197 ► **Proposition 1** (Designable $\not\Rightarrow$ Separable). *There exists a target structure which: i) does*
 198 *not admit a separated coloring; and ii) admits a solution to the inverse folding problem.*

199 **Proof.** We use the tree T of Figure 5 as a counterexample to the notion that separability
 200 fully captures designability. First, note that a separated coloring χ of T would be extremely
 201 constrained. Node 5 – 18 should be \bullet and the nodes 2 – 4 and 19 – 21 are \bullet and \circ
 202 respectively, or vice-versa due to their respective leaf. Thus, we have two leaves at levels 1
 203 and -1 . At least, one of the two children of 5 – 18, w.l.o.g 6 – 7 is \bullet or \circ . One child of
 204 6 – 7 is then necessarily \bullet , leading to a \bullet child of level $+1$ or -1 . With two leaves at level
 205 $+1$ and -1 , a direct consequence is that T is non-separable.

206 Now, we show that T is designable. We propose the sequence w of Figure 5. Using a
 207 simple dynamic programming algorithm, it is possible to check that the best folding for w is
 208 unique and corresponds to the secondary structure encoded as the tree T . Intuitively, the
 209 only competitive alternative base pair is the one corresponding to the overlap of the levels. It
 210 consists of joining the U from 8 – 9 with the A at position 3. By doing so, note that the base
 211 pair 5 – 18 will be disconnected with no way to pair A with another U due to the connection
 212 between 5 and 7. ◀

213 Notice that, despite not being separated, the coloring shown in Figure 5 is compatible with
 214 a sequence that is a design for its target. This illustrates the fact that, while not being



■ **Figure 6 Instances of $\text{INVERSE-FOLDING}_{\text{BP}}$ and main results.** For unconstrained instances (Left), solving $\text{INVERSE-FOLDING}_{\text{BP}}$ is probably NP-hard, as suggested by the hardness of a mildly constrained version [2]. Finding a design for a separable target is also NP-hard but, for any fixed modular level m , m -separable targets can be designed in $\Theta(n)$ time. Testing increasing values of m gives an algorithm, FPT on m , for all separable structures. When $h \geq 3$ (Right), Thm 6 applies and the hierarchy collapses: any instance becomes 2-separable, thus separable and designable, and $\text{INVERSE-FOLDING}_{\text{BP}}$ can be solved in $\Theta(n)$ time.

215 guaranteed to uniquely fold as their intended target, sequences produced from non-separated
 216 colorings may still represent solutions for the inverse folding problem.

217 Regarding computational complexity, although looking for a separable coloring is not
 218 directly equivalent to finding a design for a structure, we show that this decision problem
 219 (formalized below) is also NP-complete.

220 ► **Problem 2. SEPARABILITY**

221 **Input:** Target tree T (with no m_3 or m_5 motif)

222 **Output:** Coloring χ of the tree T such that χ is separated

223 ► **Theorem 3. SEPARABILITY is NP-complete.**

224 The proof can be found in the appendix. It is obtained by reduction from BIN PACKING,
 225 with a tree using one branch per item. Leaves and \bullet nodes enforce that items must be
 226 packed in consecutive ranges of levels (with \bullet levels at transitions between successive items
 227 and other levels saturated with leaves). Then, separating \bullet nodes are placed to enforce
 228 that series of consecutive items sum up to the target bin size, thus enforcing that items are
 229 ordered according to a correct bin packing.

230 4 Modulo separability as a parameterized tractable alternative

231 Then, we introduce a stratified version of separability, called modulo m -separability, or
 232 m -separability in short, which prescribes different modular values for the levels of \bullet and
 233 leaves nodes. Figure 6 describes the relative positioning of classes of instances and associated
 234 complexity results.

► **Definition 5 ((Modulo) m -separability).** Let m be an integer. A coloring χ is m -separated (or separated with modulus m) for a target secondary structure T , if and only if χ is proper and

$$\{L(v) \bmod m \mid \chi(v) = \bullet\} \cap \{L(v) \bmod m \mid v \text{ is a leaf}\} = \emptyset$$

235 using for negative levels $l < 0$ the classic $l \bmod m := (l + \lceil -x/m \rceil \times m) \bmod m$.

236 Structure T is m -separable if it admits an m -separated coloring.

237 Clearly, modulo separability implies classic separability: if a coloring χ is m -separated for
 238 a target structure T , then χ is separated for T . Conversely, if a target structure admits a
 239 separated coloring, assigning levels in $[-a, b]$ to \bullet and leaf nodes, then the same coloring
 240 is provably m' -separated for $m' := (b + a + 1)$ (since, for $l, l' \in [-a, b]$, $l \neq l'$ implies that

241 $l \bmod m' \neq l' \bmod m'$). Note that, since there are at most $n/2$ base pairs/internal nodes in
 242 a target tree, then $0 \leq a, b \leq n/2$, and we have $m' \leq n$.

243 The concept of m -separability thus provides an angle to address the generation of separated
 244 colorings, so we introduce below the associated formalized algorithmic problem.

245 ► **Problem 3. MODULO SEPARABILITY**

246 **Input:** A tree T (with no $m_{3\bullet}$ or m_5 motif), a modulus $m \in \mathbb{N}$

247 **Output:** A coloring χ of T that is m -separated, or \perp if no such coloring exists.

248 As noted above, the problem specializes in the SEPARABILITY problem when $m = n$, implying
 249 that MODULO SEPARABILITY remains NP-complete. However, it can be efficiently solved for
 250 moderate values of m , as shown below. In practice, we can focus on small values of m since
 251 in practice 99% of instances without isolated base pairs are separable with modulus $m \leq 6$
 252 (cf Table 1).

253 4.1 Fixed parameter tractable algorithm for modulo-separability

254 We now show that, for any fixed modulus m , MODULO SEPARABILITY can be solved in linear
 255 time. In particular, the problem is Fixed Parameter Tractable (FPT) for the parameter m .

256 Towards that goal, we consider a constrained version of MODULO SEPARABILITY, where
 257 the modular values of levels are prescribed. Formally, we enforce that leaves only occur at
 258 modular levels in $\xi_L \subseteq [0, m[$, and \bullet nodes only occur at levels $[0, m[\setminus \xi_L$. In this constrained
 259 version of MODULO SEPARABILITY, the existence of a valid solution can be solved in linear
 260 time using dynamic programming.

261 Namely, let us denote by $d_{v \rightarrow c, \ell}^{\xi_L}$ the existence of a valid assignment (*i.e.* solution) for
 262 a subtree of T rooted at internal node v , with v occurring at level ℓ , and being assigned a
 263 prior color c . Provably, $d_{v \rightarrow c, \ell}^{\xi_L}$ can be computed recursively by progressing along the tree,
 264 keeping track of the current level and checking that leaves and \bullet end up being assigned at
 265 modular levels ξ_L and $[0, m[\setminus \xi_L$ respectively. This leads to the following formula:

$$\begin{aligned}
 266 \quad d_{v \rightarrow c, \ell}^{\xi_L} = & \begin{cases} \text{False} & \text{if } \ell \in \xi_L \wedge c = \bullet \\ & \text{or } \ell' \notin \xi_L, \text{ and } \exists \text{ leaf in children}(v) \\ \text{True} & \text{if children}(v) = \emptyset \\ \bigvee_{\substack{c' \text{ proper} \\ \text{coloring of} \\ \text{children}(v) \\ \text{given } v \rightarrow c}} \bigwedge_{v' \in \text{children}(v)} d_{v' \rightarrow c'(v'), \ell'}^{\xi_L} & \text{otherwise.} \end{cases} \\
 267 \quad & \text{with } \ell' := \ell + \delta(c) \bmod m
 \end{aligned}$$

268 where δ denotes the level increment induced by a color c , defined as $\delta(\bullet) = +1$, $\delta(\circ) = -1$
 269 and $\delta(\circ) = 0$. Moreover, in the outermost loop, the color assignment explored for children is
 270 meant to be locally proper: the colors $c(v')$ of the children, in conjunction with the color
 271 c of v must obey the conditions of Definition 2. Note that, in the absence of $m_{3\bullet}$ and m_5 ,
 272 the number of (proper) assignments is bounded by a constant, so this conjunctive loop
 273 does not impact the complexity. The existence of a ξ_L coloring for the full tree is then
 274 $\text{Separable}_{\xi_L} := d_{\text{Root} \rightarrow \emptyset, 0}^{\xi_L}$.

275 The decision version of the problem can thus be solved in $\Theta(m.n)$ time. Indeed, the
 276 number of left-hand side terms scales in $\Theta(m.n)$, the number of proper coloring for children

277 is bounded by a constant (since avoiding $m_{3\bullet}$ and $m_5 \implies |\text{child}(v)| < 5$), and the total
 278 number of executions of the conjunctive loops is in overall $\Theta(n)$. A backtracking procedure
 279 could also be defined to reconstruct a solution coloring in $\Theta(n)$ if such a solution exists
 280 ($\text{Separable}_{\xi_L} = \text{True}$) or return \perp otherwise ($\text{Separable}_{\xi_L} = \text{False}$).

281 An algorithm for MODULO SEPARABILITY can then be obtained by explicitly considering
 282 all the possible subsets of admissible modular levels for leaves:

- 283 ■ If T contains $m_{3\bullet}$ or m_5 , return \perp
- 284 ■ For each $\xi_L \subseteq [0, m[$:
 - 285 ■ If $\#\text{Designs}_{\xi_L} > 0$, then backtrack to produce ξ_L -separated design
- 286 ■ Return \perp

287 The algorithm is correct since any ξ_L solution is also m -separated, and any m -separated
 288 coloring implies a partition of the leaves and \bullet nodes into disjoint levels ξ_L and $\chi_{\bullet} \subseteq [0, m[\setminus \xi_L$
 289 respectively. A m -separated coloring is thus always found by invoking the DP algorithm over
 290 the 2^m subsets $\xi_L \in [0, m[$. The overall complexity of the algorithm is in $\Theta(n.m.2^m)$ time
 291 and $\Theta(m.n)$ memory, and we conclude with the parameterized complexity of the problem
 292 with respect to m .

293 ► **Theorem 4.** MODULO SEPARABILITY is Fixed Parameter Tractable for the modulus m

294 4.2 Random generation of m -separated RNA sequences

295 We then turn to the uniform random generation of m -separated sequences, defined as a
 296 design w for T , featuring A on unpaired positions, and such that the coloring χ_w , obtained by
 297 replacing base pairs with suitable color ($(G, C) \rightarrow \bullet$, $(C, G) \rightarrow \circ$ and (A, U) or $(U, A) \rightarrow \bullet$),
 298 is m -separated.

► **Problem 4.** UNIFORM MODULO SEPARATED GENERATION

Input: Target tree T (with no $m_{3\bullet}$ or m_5 motif)

Output: RNA sequence w , associated with m -separated coloring χ_w , such that

$$\mathbb{P}(w \mid \chi_w \text{ is } m\text{-separated}) = \frac{1}{|\{w' \mid \chi_{w'} \text{ is } m\text{-separated}\}|}$$

299 Again, we approach this problem by first solving a more constrained version where the
 300 modular levels of leaves are explicitly given as a set ξ_L . Then, in the spirit of Reinharz *et*
 301 *al* [15], we adapt the above recurrence, through a simple algebra change, to count the number
 302 $p_{v \rightarrow \mu, l}^{\xi_L}$ of RNA sequences, associated with a ξ_L separated coloring (for a subtree of T rooted
 303 at v , with v occurring at level l , and being assigned a nucleotide assignment μ).

$$304 \quad p_{v \rightarrow \mu, l}^{\xi_L} = \begin{cases} 0 & \text{if } l \in \xi_L \text{ and } \mu \in \{(A, U), (U, A)\} \\ 0 & \text{if } l' \notin \xi_L \text{ and } v \text{ has a leaf attached} \\ 1 & \text{if } \text{children}(v) = \emptyset \\ \sum_{\substack{\mu' \text{ proper assignment} \\ \text{children}(v) \rightarrow \Sigma^2 \cup \{\emptyset\}}} \prod_{v' \in \text{children}(v)} p_{v' \rightarrow \mu'(v'), l'}^{\xi_L} & \text{otherwise } (l' := l + \delta(\mu) \bmod m). \end{cases}$$

305 where μ' is a nucleotide assignment to the children of v , consistent with a proper coloring
 306 and additionally respecting natural constraints on the content $((A, U)$ or $(U, A))$ of pairs of
 307 \bullet nodes (same for both if one parent of other, different content if siblings). Once again, the
 308 colorless Root node needs to be distinguished, and the overall number of designs is given by
 309 $\#\text{Designs}_{\xi_L} := p_{\text{Root} \rightarrow \emptyset, 0}^{\xi_L}$.

XXX:10 Exact linear-time RNA design for min Helix length 3

310 The following backtrack procedure then produces a uniform random RNA sequence that
 311 corresponds to a m -separated coloring for a fixed set ξ_L . In that case, by abuse of language,
 312 we say that the sequence is ξ_L *separated*. More precisely, $\text{backtrack}(v, c, \ell)$ produces a random
 313 sequence, associated with a ξ_L separated coloring, for the subtree anchored in v , reached at
 314 height ℓ , where the root is assigned a pair of bases $\mu \in \Sigma^2$. It first picks a random proper
 315 assignment μ' for the children, weighted by the corresponding number of solutions (namely,
 316 $\prod_{v' \in \text{children}(v)} p_{v' \rightarrow \mu'(v'), \ell'}^{\xi_L}$, with $\ell' := \ell + \delta(\mu) \pmod m$). The resulting sequence is then

$$\prod_{v \in \text{children and leaves}(v)} \begin{cases} A & \text{If } v' \text{ is a leaf} \\ b.\text{backtrack}(v', \mu'(v'), \ell').b' & \text{otherwise, with } \mu'(v') = b.b' \end{cases}$$

317 The resulting algorithm, consisting of precomputing all $p_{v \rightarrow \mu, \ell}^{\xi_L}$, followed by a sequence of
 318 k backtracks, provably returns k random, uniformly-distributed and independent designs
 319 that are ξ_L separated in time $\Theta(n.m + k.n)$.

320 To leverage the uniform generation for a fixed ξ_L into a uniform generation of m -separated
 321 designs, we implement a strategy (see [14, pp 77] for details), proven in Appendix C, which
 322 start by generating some ξ_L , and then uses a suitable rejection to correct the emissions
 323 probabilities of sequences compatible with several ξ_L .

324 ► **Theorem 5.** UNIFORM MODULO SEPARATED GENERATION *can be performed in an*
 325 *average-case complexity that is Fixed Parameter Tractable for the modulus parameter m .*

5 Structures without isolated stacks and base pairs are 2-separable

327 Although separability does not give a full characterization of designability in general (cf
 328 Prop. 1), we obtain a much stronger result for structures without small helices, as hinted by
 329 the fact that all counter-examples and hardness gadgets heavily use isolated base pairs in
 330 their construction. Indeed, we show that a 2-separated coloring can be constructed for *all*
 331 structures without forbidden motifs $(m_{3\bullet}, m_5)$ and $h \geq 3$, so indeed all such structures are
 332 designable. Since avoiding $(m_{3\bullet}, m_5)$ is a necessary condition for designability, we obtain the
 333 stronger characterization stated in Corollary 9.

334 ► **Theorem 6.** *Every $(m_{3\bullet}, m_5)$ -avoiding target T , having $h \geq 3$, admits a 2-separated*
 335 *coloring*

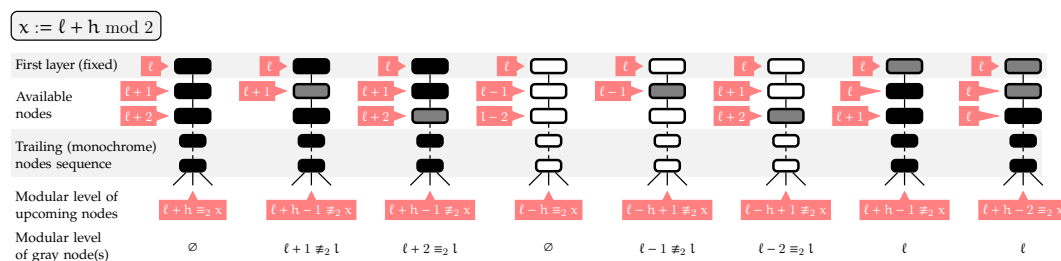
336 **Proof.** First, let us remark that helices can be treated as atomic objects, and compacted
 337 into the edges of a *helix tree*, whose edges are helices (sequence of consecutive BP nodes),
 338 and whose internal nodes are either:

- 339 ■ Multiloops, consisting of 2 or 3 children/BPs/Helices, and no leaf (so $m_{3\bullet}$ does not occur);
- 340 ■ Internal/Bulges/Hairpin (IBH) loops, consisting of at most 1 BP/Helix and featuring at
 341 least one leaf/unpaired node.

342 Remark that, while constructing a separated coloring assigning a modular level ξ_L to leaves,
 343 those two motifs are the only sources of immutable constraints:

- 344 ■ Any proper coloring of a multiloop features at least one \bullet node, so the levels of chil-
 345 dren/nodes need to be set to a level $\xi_L := \xi_L + 1 \pmod 2$;
- 346 ■ Any IBH loop features at least one leaf within its children, which needs to be set to a
 347 modular level ξ_L .

348 Conversely, beyond their first BP, helices may be colored with very limited constraints and
 349 can be used to *offset* multiloops and IBH loops.



■ **Figure 7** Existence of coloring for helices consisting of $h \geq 3$ base pairs, such that the modular level of the following nodes is offset as needed. Additionally, those colorings can be chosen to respect a prescribed level for \bullet nodes and, a predetermined color for the first node/base pair of the helix.

350 ▶ **Lemma 7.** Let $\overline{\xi_L}$ denote the prescribed modular level for \bullet nodes. Consider an helix H
 351 consisting of $h \geq 3$ BPs, whose first BPs is assigned some color $c \in \{\bullet, \circ, \ominus\}$.

352 Then for each modular level $l \in [0, 1]$ for the first BP of H ($c = \bullet$ only if $l = \overline{\xi_L}$), and
 353 targeted exit modular level $l' \in [0, 1]$, there exists a coloring for the rest of H such that:

- 354 ■ The modular level of the upcoming nodes, i.e. those immediately following H , is l' ;
- 355 ■ Base pairs can only be \bullet -colored at modular level $\overline{\xi_L}$.

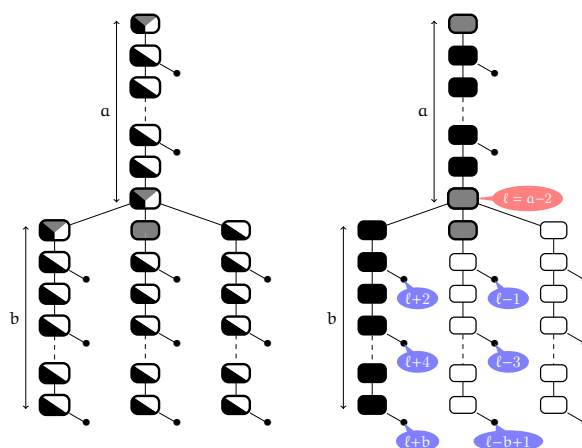
356 **Proof.** The proof is essentially based on case decomposition, and summarized in Figure 7.
 357 We show that, for any l and h , there exists a color assignment to the first 3 nodes of the
 358 helix, such that the modular level of upcoming nodes is either 0 or 1, so l' can be reached.
 359 Moreover, if such a coloring starts with \bullet or \circ , and uses a single \bullet node, then there exists
 360 an alternative coloring placing this \bullet node at the opposite modular level, so one of them
 361 places their \bullet node at the intended level $\overline{\xi_L}$. Finally, if the first node is set to \bullet , then the
 362 consistency condition above implies that $l \bmod 2 = \overline{\xi_L}$, so that \bullet nodes are naturally found
 363 at an admissible modular level. ◀

364 It follows that any helix tree starting with an initial helix H can be colored into a 2-separated
 365 coloring. Starting at initial level $l = 0$ and having initial BP color c ($c \neq \bullet$ if $\xi_L = 0$), color
 366 the rest of H as shown in the proof of Lemma 7, depending on $\overline{\xi_L}$ and the type of upcoming
 367 loop (target $l' = \overline{\xi_L}$ for Multiloops; $l' = \xi_L$ for IBH loops), while ensuring that \bullet nodes end
 368 up at $\overline{\xi_L}$ modular level (which can always be done from Lemma 7). The remaining nodes of
 369 the loop are then colored in a proper/greedy manner, and we iterate the process recursively
 370 on the children helices of the loop (if any) until the full tree is colored.

371 Since its level cannot be offset, the Root node must be treated as a special case. Indeed,
 372 if the Root has at least one leaf/unpaired position, then the modular value 0 is taken by
 373 the leaf, so we must have $\xi_L = 0$. Conversely, if the Root supports at least 3 helices, then
 374 at least one needs to start with a \bullet node, so we must have $\xi_L = 1$. Regardless of this
 375 restriction on ξ_L , in both cases the first base pair of each helix (if any) supported by the Root
 376 can be properly colored, and helices can be independently colored using the above strategy,
 377 ultimately yielding a 2-separated coloring. ◀

378 ▶ **Corollary 8.** INVERSE FOLDING, restricted to instances with $h \geq 3$ (containing no isolated
 379 base pair and no isolated stacks) is solvable in linear time and space.

380 It is a direct consequence of Theorem 6 and of the DP scheme introduced in Section 4.1.
 381 Indeed, for $m = 2$, the DP algorithm only needs to be run twice ($\xi_L = 0$ and $\xi_L = 1$) in linear
 382 time/space, to produce a 2-separated coloring whenever such a coloring exists (guaranteed
 383 by Theorem 6). The coloring can then be transformed into a design, i.e. a solution to the



■ **Figure 8** Main gadget used to build non-separable instances with $h = 2$. Left: Admissible colors for each node (up to branch symmetries). Right: Example coloring and levels of a selection of leaves and \bullet nodes. Note that along with the \bullet node at level ℓ , there always exists a leaf at level $\ell + m$ or $\ell - m$ for $2 \leq m \leq b$, ruling out modulo separability for small m .

384 INVERSE FOLDING problem. Similarly, UNIFORM MODULO SEPARATED GENERATION can
 385 also be performed in linear expected time and space as long as input instances contain only
 386 helices of size 3 or more.

387 ► **Corollary 9.** *Let T be a target structure with $h \geq 3$, then the following are equivalent: i)*
 388 *T is designable; ii) T is 2-separable; and iii) T avoids $(m_{3\bullet}, m_5)$.*

389 With this result, the hierarchy of instances collapses as depicted on the left of Figure 6 A
 390 natural follow-up question is whether the bound 3 on the helix length is tight. Indeed, there
 391 are non-separable and designable instances with $h = 1$ (Proposition 1), but the question
 392 remains for $h = 2$. In Proposition 10 we give a non-separable instance without isolated base
 393 pairs, so $h = 3$ is indeed tight to ensure separability.

394 ► **Proposition 10.** *There exist non-separable structures with $h = 2$.*

395 The full proof relies on a counterexample built from the gadget in Figure 8. Intuitively,
 396 $T(a, b)$ saturates all levels modulo b with leaves, so that none remains available for \bullet nodes.
 397 Meanwhile, the presence of multiloops forces proper colorings to use \bullet nodes, so a collision
 398 occurs and the gadget is not m -separable for any $m \leq b$. By assembling 5 copies of $T(a, b)$
 399 with large b and increasing values of a , we obtain a target that is not separable for any m .

400 6 On the relevance of separated sequences towards realistic designs

401 While the existence of a linear-time algorithm for a reasonable restriction of the inverse folding
 402 problem is already notable, its practical relevance may be perceived as hindered by several
 403 limitations: our algorithms are only guaranteed to produce design solutions for helices beyond
 404 3 base pairs; Proper colorings only allows the design of highly-constrained (multi)loops; and
 405 solutions to the base pair inverse folding are not guaranteed to represent good solutions
 406 in more realistic energy models, such as the Turner nearest-neighbor model. To assess the
 407 promises offered by separated designs in realistic settings, we performed computational
 408 experiments to assess the potential of the concept of separated colorings to serve as a basis
 409 for the design of future RNA design methods.

m^\ominus – Min. separating modulus	2	3	4	5	6	7	8	9	10+
Number of instances	7720	202	1101	537	349	77	9	5	0

■ **Table 1** Minimal modulus \ominus required for 10 000 instances of size $n = 100$ with min hairpin loop $\theta = 3$ containing some isolated stacks (helices restricted to 2 base pairs) to be separable modulo m . All instances we found to be modulo separable.

6.1 Targets with isolated BPs/stacks are frequently separable

While our algorithm is only guaranteed to produce a design when $h \geq 3$, it also produces (guaranteed correct) solutions for input with smaller helices, as long as a separated coloring exists for them.

For very small targets, an exhaustive analysis is feasible, consisting of folding/testing the unicity of the MFE folding for all sequences of length $n = 12$ (see Figure 2). Moreover, once a design w is found for a target T , it is easy to test if the associated coloring χ_w is separated, and to compute minimal modulus value m^\ominus such that χ_w is m^\ominus separated. We found that *all of the 8 111 designable targets are also separable*, despite a very large proportion of them featuring isolated stacks and base pairs. Moreover, all designable targets admit separated solutions associated with very small values of the modulus m (7 690 for $m = 2$, 420 for $m = 3$ and $m = 1$ only for the empty structure).

To further measure the proportion of separable structures within larger targets featuring isolated stacks, we implemented a uniform random generation algorithm [14]. We produced random target secondary structures of length 100 with a min base pair span of $\theta = 3$. We used rejection to produce a synthetic dataset consisting of 10 000 targets having at least one helix of size 2 while avoiding $m_{3\bullet}$ and m_5 . For each target T , we ran an in-house implementation of the algorithm in Section 4.1 with increasing modulus, to find the minimal modulus m^\ominus such that T admits a m^\ominus separated coloring. Table 1 summarizes our results, which we discuss below.

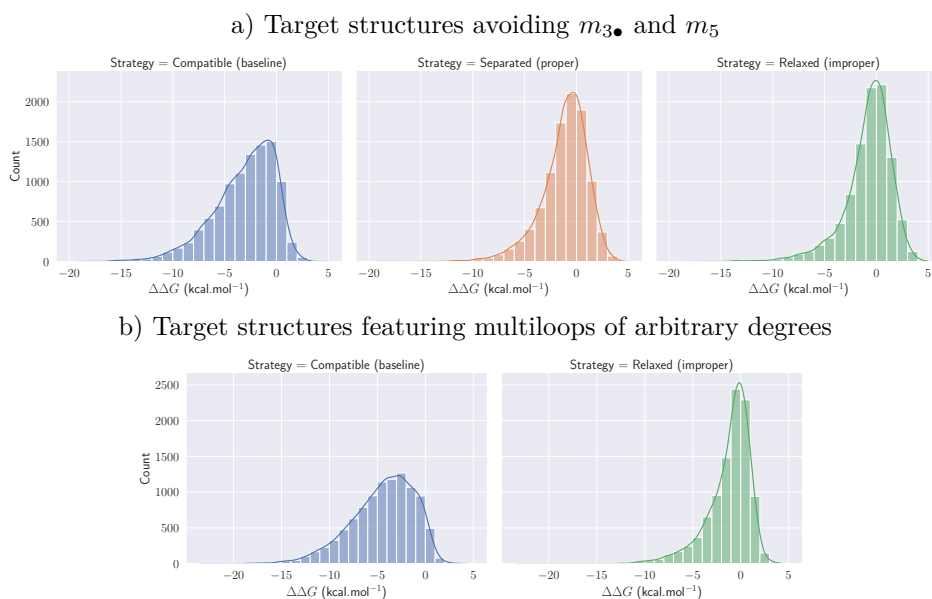
Remarkably, all of the 10k targets in the datasets could be designed using our algorithm, and thus admit a separable coloring. Moreover, roughly three-quarters (77%) of the targets were found to be 2-separable, and less than 1% of the targets required the consideration of values for m^\ominus beyond 6. The max value for m^\ominus in this dataset was 9, an order of magnitude lower than the sequence length. Clearly, since we have shown the existence of non-separable instances with isolated stacks and no isolated base pair, this observation does not generalize to arbitrary sequence lengths. However, the large size of these counterexamples suggests that the proportion of separable structures, despite ultimately decaying exponentially [20], may remain non-negligible for relevant RNA target sizes.

6.2 Separated designs are promising candidates in the Turner model

We now consider a more realistic setting, where the inverse folding problem is now considered with respect to the Turner nearest-neighbor energy model [19]. To assess the value of a sequence in the Turner model, we introduce a metrics which we call the (signed) *energy distance* $\Delta\Delta G(w, T)$ of a target T to its *most stable distant alternative* for the sequence w :

$$\Delta\Delta G(w, T) := \Delta G(w, \text{Alt}_{d_{\min}}(w, T)) - \Delta G(w, T), \text{Alt}_{d_{\min}}(w, T) := \min\{\Delta G(w, T') \mid |T', T| \geq d_{\min}\}$$

where $\Delta G(w, T)$ is the Turner free-energy, $|T, T'| := |T \Delta T'|$ denotes the base-pair distance, and d_{\min} represents the minimum base pair distance to T . Both ΔG and $\text{Alt}_{d_{\min}}(w, T)$ can be obtained by appropriate calls to the ViennaRNA package [8], namely RNAeval and



■ **Figure 9** Comparison of compatible (baseline), separated, and relaxed models for targets having $\theta = 3, h = 3$. For energy distance parameters, we took $d_{min} = 3$ and $E = 5$.

443 RNAsubopts, using max energy distance parameter $E = 5$ (so our estimation of $\Delta\Delta G(w, T)$ is
 444 bounded by 5). A positive energy distance confirms that w is a solution to the Turner version
 445 of inverse folding, and dominates its competitors by $\Delta\Delta G(w, T)$ kcal.mol⁻¹. Meanwhile,
 446 a negative energy distance indicates that the target T is dominated by some alternative
 447 structure, having $\Delta\Delta G(w, T)$ kcal.mol⁻¹ lower free-energy than the target.

448 We consider three strategies for sampling sequences: i) The *compatible* model uniformly
 449 generates random sequences compatible with the target (A for unpaired positions; AU, UA,
 450 GC or CG for base pairs); ii) The *separated* model uses the sampler described in Section 4.2
 451 to generate sequences that are 2-separated and proper; iii) The *relaxed* model generates
 452 sequences that are 2-separated, but not necessarily proper by assigning uniform random
 453 pairs to the base pairs of a multiloop. The *relaxed* model enables a heuristic extension of
 454 our algorithms supporting multiloops of arbitrary degrees, noting that the local refolding
 455 (see Figure 3) occurring in the BP model for non-proper sequences are either unrealistic or
 456 outright impossible, in the Turner energy model.

457 Separated sequences substantially improve over compatible random sequences.

458 We first asked a moderately ambitious question: *Are separated sequences better candidates for*
 459 *design in the Turner model than sequences whose only property is to be compatible with the*
 460 *target?* Indeed, separated sequences are only guaranteed to represent a design for their target
 461 in the BP max. model. We first considered instances of target structures of size $n = 100$,
 462 known to admit a solution to INVERSE-FOLDING_{BP} ($\theta = 3$, no $m_{3\bullet}$ or m_5 , and $h \geq 3$). We
 463 generated 10 000 such random targets and, for each target, sampled a single sequence using
 464 each of the 3 strategies above and computed the energy distance.

465 The results, summarized in Figure 9.top suggest that separated sequences represent a
 466 substantial improvement over merely compatible sequences. Indeed, while 10% of compatible
 467 sequences ended up being good design candidates ($\Delta\Delta G > 0$), the proportion of successful
 468 designs increases to approximately one-third (31%) for separated sequences, and further to
 469 40% for relaxed design. A similar trend can be observed for the average $\Delta\Delta G$ (distance to

470 the first alternative/competitor) among successful designs, being of 0.75/1.05/1.13 kcal.mol⁻¹
471 in the compatible, separated and relaxed models respectively. The surprisingly good behavior
472 of the relaxed model, which was mostly introduced to overcome unrealistic limitations on
473 multiloops, remains to be explained.

474 **Relaxed sequences enable designs for multiloops having higher degrees.** We also
475 tested the capacity of the relaxed model to generate solutions for multiloops of higher degrees,
476 noting that the avoidance of $m_{3\bullet}$ and m_5 restricts the maximum degree of a multiloop to 4.
477 We used the above-mentioned generation algorithm to generate uniform design targets of
478 size $n = 100$, featuring at least one (but frequently many) occurrence of $m_{3\bullet}$ and m_5 . As
479 shown in Figure 9.bottom, compatible sequences are again substantially outperformed by the
480 relaxed separated model in this setting, with 31.5% of the separated/non-proper sequences
481 (4.6% of compatible sequences) representing successful designs ($\Delta\Delta G > 0$), on average 0.84
482 kcal.mol⁻¹ more stable than their best competitor.

483 **7 Conclusion**

484 Adapting a coloring perspective initially introduced by Halès *et al* [6], we have shown that
485 the inverse folding problem can be solved in linear time for all target secondary structures
486 having minimum helix length equal to 3. Towards that main result, we have established the
487 existence of designable, yet non-separable, instances of inverse folding, and the NP-hardness
488 of finding a separable design in the initial sense of Halès *et al*. We have also introduced
489 concrete algorithms for the problem of finding a m modulo-separated coloring, which we
490 have shown to be NP-hard yet FPT-solvable for m . Already for $m = 2$, the scope of our
491 algorithms encompasses all targets without isolated base pairs and stacks, but also extends
492 much beyond, in a way that remains to be fully characterized. Beyond base pair maximization,
493 modulo-separated sequences may also represent a solid foundation towards concrete design
494 methodologies. Namely, we empirically showed that, for the Turner energy model, separated
495 sequences tend to represent better design candidates than merely compatible sequences,
496 and that the limitations on loop degrees (intrinsic to the BP maximization model) can be
497 overcome by relaxing our design model while retaining substantial performances.

498 Future work should focus on how much of designable sequences are covered by sequences
499 obtained with (modulo-)separated colorings. More importantly, does the space of (modulo-
500)separated colorings always/often contain a design with respect to the nearest-neighborhood
501 Turner energy model? Even if it unlikely to hold unconditionally, it is plausible that some
502 extensions of separability and m -separability will achieve theoretical and practical solutions
503 of the problem in more general energy models. As a first step, separability in a stacking
504 energy model seems a relevant goal, even if less ambitious than the Turner model. It would
505 probably require to go beyond the current coloring formalism, implying the introduction a
506 more general notion of defect to encapsulate imbalance at the dinucleotide level.

507 — References —

- 508 1 Mirela Andronescu, Anthony P. Fejes, Frank Hutter, Holger H. Hoos, and Anne Condon. A new
509 algorithm for rna secondary structure design. *Journal of Molecular Biology*, 336(3):607–624,
510 2004. URL: <https://www.sciencedirect.com/science/article/pii/S0022283603015596>,
511 [doi:10.1016/j.jmb.2003.12.041](https://doi.org/10.1016/j.jmb.2003.12.041).
- 512 2 Édouard Bonnet, Paweł Rzażewski, and Florian Sikora. Designing rna secondary structures is
513 hard. *Journal of Computational Biology*, 27(3):302–316, 2020. PMID:32160034. [arXiv:https://](https://arxiv.org/abs/2019.0420)
514 [doi:10.1089/cmb.2019.0420](https://doi.org/10.1089/cmb.2019.0420), [doi:10.1089/cmb.2019.0420](https://doi.org/10.1089/cmb.2019.0420).
- 515 3 Anke Busch and Rolf Backofen. INFO-RNA—a fast approach to inverse RNA folding. *Bioin-*
516 *formatics*, 22(15):1823–31, 2006.
- 517 4 Ali Esmaili-Taheri and Mohammad Ganjtabesh. ERD: a fast and reliable tool for RNA design
518 including constraints. *BMC Bioinform.*, 16:20:1–20:11, 2015.
- 519 5 Juan Antonio Garcia-Martin, Ivan Dotu, and Peter Clote. RNAiFold 2.0: a web server
520 and software to design custom and Rfam-based RNA molecules. *Nucleic Acids Research*,
521 43(W1):W513–W521, 05 2015. [arXiv:https://academic.oup.com/nar/article-pdf/43/W1/](https://arxiv.org/abs/2015.05.05)
522 [W513/7476300/gkv460.pdf](https://doi.org/10.1093/nar/gkv460), [doi:10.1093/nar/gkv460](https://doi.org/10.1093/nar/gkv460).
- 523 6 Jozef Hales, Alice Héliou, Ján Manuch, Yann Ponty, and Ladislav Stacho. Combinatorial RNA
524 design: Designability and structure-approximating algorithm in watson-crick and nussinov-
525 jacobson energy models. *Algorithmica*, 79(3):835–856, 2017.
- 526 7 Stefan Hammer, Wei Wang, Sebastian Will, and Yann Ponty. Fixed-parameter tractable
527 sampling for RNA design with multiple target structures. *BMC bioinformatics*, 20:209, April
528 2019. [doi:10.1186/s12859-019-2784-7](https://doi.org/10.1186/s12859-019-2784-7).
- 529 8 Ivo L Hofacker, Walter Fontana, Peter F Stadler, L Sebastian Bonhoeffer, Manfred Tacker,
530 and Peter Schuster. Fast folding and comparison of RNA secondary structures. *Monatshefte*
531 *für Chemie/Chemical Monthly*, 125(2):167–188, 1994.
- 532 9 Robert Kleinkauf, Martin Mann, and Rolf Backofen. antaRNA: ant colony-based RNA sequence
533 design. *Bioinformatics*, 31(19):3114–3121, 05 2015. [doi:10.1093/bioinformatics/btv319](https://doi.org/10.1093/bioinformatics/btv319).
- 534 10 William Andrew Lorenz and Yann Ponty. Non-redundant random generation algorithms for
535 weighted context-free grammars. *Theoretical Computer Science*, 502:177–194, 2013. Generation
536 of Combinatorial Structures. URL: [https://www.sciencedirect.com/science/article/pii/](https://www.sciencedirect.com/science/article/pii/S0304397513000443)
537 [S0304397513000443](https://doi.org/10.1016/j.tcs.2013.01.006), [doi:10.1016/j.tcs.2013.01.006](https://doi.org/10.1016/j.tcs.2013.01.006).
- 538 11 Rune B. Lyngsø, James W. J. Anderson, Elena Sizikova, Amarendra Badugu, Tomas Hyland,
539 and Jotun Hein. Frnakenstein: multiple target inverse RNA folding. *BMC Bioinform.*, 13:260,
540 2012.
- 541 12 Nono S. C. Merleau and Matteo Smerlak. arnaque: an evolutionary algorithm for inverse
542 pseudoknotted RNA folding inspired by lévy flights. *BMC Bioinform.*, 23(1):335, 2022.
- 543 13 R Nussinov and A B Jacobson. Fast algorithm for predicting the secondary structure of
544 single-stranded rna. *Proceedings of the National Academy of Sciences*, 77(11):6309–6313,
545 1980. URL: <https://www.pnas.org/doi/abs/10.1073/pnas.77.11.6309>, [arXiv:https://](https://arxiv.org/abs/1980.11.6309)
546 www.pnas.org/doi/pdf/10.1073/pnas.77.11.6309, [doi:10.1073/pnas.77.11.6309](https://doi.org/10.1073/pnas.77.11.6309).
- 547 14 Yann Ponty. *Ensemble Algorithms and Analytic Combinatorics in RNA Bioinformatics and*
548 *Beyond*. Habilitation à diriger des recherches, Université Paris-Saclay, May 2020. URL:
549 <https://theses.hal.science/tel-03219977>.
- 550 15 Vladimir Reinharz, Yann Ponty, and Jérôme Waldispühl. A weighted sampling algorithm
551 for the design of RNA sequences with targeted secondary structure and nucleotide dis-
552 tribution. *Bioinformatics*, 29(13):i308–i315, 06 2013. [arXiv:https://academic.oup.com/](https://arxiv.org/abs/2013.06.06)
553 [bioinformatics/article-pdf/29/13/i308/50704314/bioinformatics_29_13_i308.pdf](https://doi.org/10.1093/bioinformatics/btt217),
554 [doi:10.1093/bioinformatics/btt217](https://doi.org/10.1093/bioinformatics/btt217).
- 555 16 Matan Drory Retwitzer, Vladimir Reinharz, Alexander Churkin, Yann Ponty, Jérôme
556 Waldispühl, and Danny Barash. incaRNAfbinv 2.0: a webserver and software with
557 motif control for fragment-based design of RNAs. *Bioinformatics*, 36(9):2920–2922,

- 558 01 2020. [arXiv:https://academic.oup.com/bioinformatics/article-pdf/36/9/2920/](https://academic.oup.com/bioinformatics/article-pdf/36/9/2920/48986446/bioinformatics_36_9_2920.pdf)
559 [48986446/bioinformatics_36_9_2920.pdf](https://academic.oup.com/bioinformatics/article-pdf/36/9/2920/48986446/bioinformatics_36_9_2920.pdf), doi:10.1093/bioinformatics/btaa039.
- 560 17 Frederic Runge, Danny Stoll, Stefan Falkner, and Frank Hutter. Learning to design RNA. In
561 *Proceedings of ICLR 2019*, 2019.
- 562 18 Michael Schnall-Levin, Leonid Chindelevitch, and Bonnie Berger. Inverting the viterbi
563 algorithm: an abstract framework for structure design. In *ICML*, volume 307 of *ACM*
564 *International Conference Proceeding Series*, pages 904–911. ACM, 2008.
- 565 19 Douglas H. Turner and David H. Mathews. NNDB: the nearest neighbor parameter data-
566 base for predicting stability of nucleic acid secondary structure. *Nucleic Acids Research*,
567 38(suppl_1):D280–D282, 10 2009. [arXiv:https://academic.oup.com/nar/article-pdf/38/](https://academic.oup.com/nar/article-pdf/38/suppl_1/D280/11217894/gkp892.pdf)
568 [suppl_1/D280/11217894/gkp892.pdf](https://academic.oup.com/nar/article-pdf/38/suppl_1/D280/11217894/gkp892.pdf), doi:10.1093/nar/gkp892.
- 569 20 Hua-Ting Yao, Cedric Chauve, Mireille Regnier, and Yann Ponty. Exponentially few RNA
570 structures are designable. In *ACM-BCB 2019 - 10th ACM Conference on Bioinformatics,*
571 *Computational Biology, and Health Informatics*, pages 289–298, Niagara-Falls, United States,
572 September 2019. ACM Press. URL: <https://inria.hal.science/hal-02141853>, doi:10.
573 [1145/3307339.3342163](https://inria.hal.science/hal-02141853).
- 574 21 Hua-Ting Yao, Jérôme Waldispühl, Yann Ponty, and Sebastian Will. Taming Disruptive Base
575 Pairs to Reconcile Positive and Negative Structural Design of RNA. In *Proc. of the 25th*
576 *Annual International Conferences on Computational Molecular Biology (RECOMB'21)*, 2021.
577 URL: <https://inria.hal.science/hal-02987566>.
- 578 22 Joseph N. Zadeh, Brian R. Wolfe, and Niles A. Pierce. Nucleic acid sequence design via
579 efficient ensemble defect optimization. *Journal of Computational Chemistry*, 32(3):439–452,
580 2011. doi:10.1002/jcc.21633.

581 **A NP-completeness of general separability (Proof of Theorem 3)**

582 SEPARABILITY is clearly in NP, since any coloring (certificate) can be checked in linear
 583 time. We prove hardness by reduction from BIN PACKING which we formulate as an interval
 584 packing problem.

585 ► **Problem 5.** INTERVAL PACKING

586 **Input:** set of pairwise distinct integers $A = \{a_1, \dots, a_n\}$, integers k and B

587 **Output:** function x from A to intervals of $[0, kB - 1[$ such that:

- 588 ■ $x(a_i)$ is an interval of size a_i
- 589 ■ $x(a_i)$ and $x(a_j)$ are disjoint for $i \neq j$
- 590 ■ $x(a_i)$ does not contain both $jB - 1$ and jB for any i, j .

591 This is a reformulation of BIN PACKING: fitting items for a total size of B is equivalent to
 592 finding a partition of a size- B interval into smaller intervals. The problem remains NP-hard
 593 even when input integers are encoded in unary (which corresponds to the fact that BIN
 594 PACKING is strongly NP-hard). We further require that all items have size $a_i \geq 5$

595 **Object and border gadgets.** We first give the main gadgets for our reduction, see figure 10
 596 for more details.

597 ► **Definition 6.** An object gadget of size $q \geq 3$ is a chain of $q + 3$ nodes c_0, \dots, c_{q+2} with a
 598 child attached to c_1 and c_{q+1} and leaves attached to all other nodes c_i .

599 A period- p border gadget of size q is a chain of q nodes c_0, \dots, c_{q-1} with a child attached
 600 to c_i for all $i \equiv 0 \pmod p$ and leaves attached to all other nodes c_i .

601 ► **Proposition 11.** If an object gadget of size q appears in a tree with a separated coloring χ ,
 602 with $\ell = \min\{L(c_i) \mid 1 \leq i \leq q\}$ such that

- 603 ■ there are ● nodes at levels $\ell + 2$ and $\ell + (q + 2)$
- 604 ■ there are leaves at levels $\ell + i$ for all $1 \leq i \leq q + 3$, $i \neq 2, q + 2$.

605 If a period- p gadget of size q appears in a tree with a separated coloring χ , with the root
 606 at level ℓ , then there exists some direction $d \in \{-1, 1\}$ such that

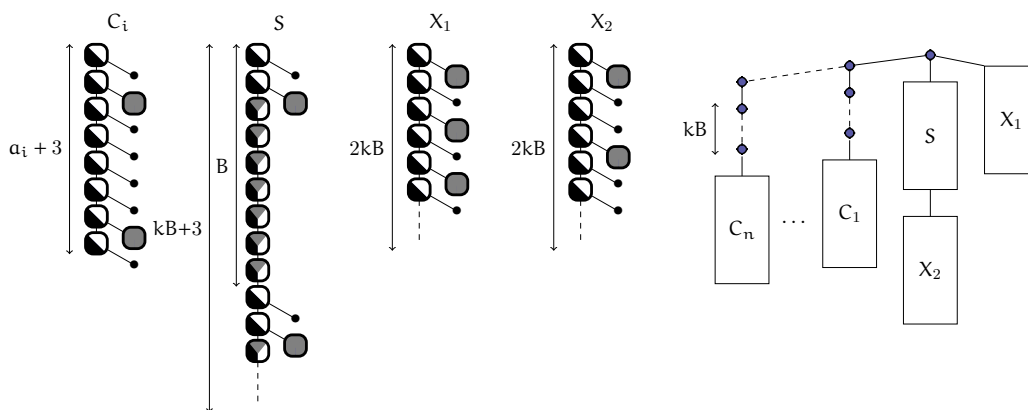
- 607 ■ there are ● nodes at levels $\ell + d \cdot i + 1$ for all $1 \leq i \leq q$, $i \equiv 0 \pmod p$;
- 608 ■ there are leaves at levels $\ell + d \cdot i$ for all $1 \leq i \leq q + 3$, $i \not\equiv 0 \pmod p$.

609 **Proof.** First note that in either gadget, all nodes c_i have the same non-● color. Indeed,
 610 nodes with a leaf attached or a leaf sibling cannot be ●, so all c_i are ● or ○. Furthermore,
 611 by the proper coloring constraints, consecutive nodes must be of the same color, so all c_i
 612 have the same color. Thus, writing ℓ_r for the root level, we have that the level below each
 613 node c_i is $\ell_r + di$, with $d = 1$ if the whole chain is ● and $d = -1$ otherwise.

614 Furthermore, all nodes attached to the chain must be ● by the proper coloring constraints.
 615 This directly gives the desired property for border gadgets. For object gadgets, the minimum
 616 level ℓ along the chain is either ℓ_r (if $d = 1$) or $\ell_r - q - 3$ (if $d = -1$), and in both cases, for
 617 each level $\ell + i$ with $1 \leq i \leq q + 3$, there is either a ● node ($i = 2$ or $i = q + 2$) or a leaf
 618 (otherwise). ◀

619 **Reduction.** Given an instance A, k, B of INTERVAL PACKING, we build a tree T as follows:

- 620 ■ We start with a chain P of $n + 1$ nodes denoted p_0, \dots, p_n .
- 621 ■ For each $i \geq 1$ we attach a chain (denoted P_i) of Bk nodes to p_i , and an object gadget
 622 C_i of size a_i to the end of the chain.
- 623 ■ We attach a period-2 border gadget of size $2kB$ to p_0 , denoted X_1 .



■ **Figure 10** Left: details of the four main parts of the reduction, i.e. an object gadget C_i of size a_i (in this example with $a_i = 5$), border gadgets X_1 and X_2 with respective periods 2 and 3, and the separator chain S). Right: general layout of the tree built in the reduction.

- 624 ■ We attach a chain S of $kB + 3$ nodes to p_0 with:
- 625 ■ a leaf to the $(iB + 1)$ st node of S for each $0 \leq i \leq k$,
 - 626 ■ a second child, called *separator*, to the $(iB + 2)$ nd node of S for each $0 \leq i \leq k$,
 - 627 ■ a period-3 border gadget of size $2kB$ at the end of S , denoted X_2 .

628 We will now show that there exists a solution for unary bin packing if and only one can find
629 a separated coloring for T .

630 From interval packing to separated coloring.

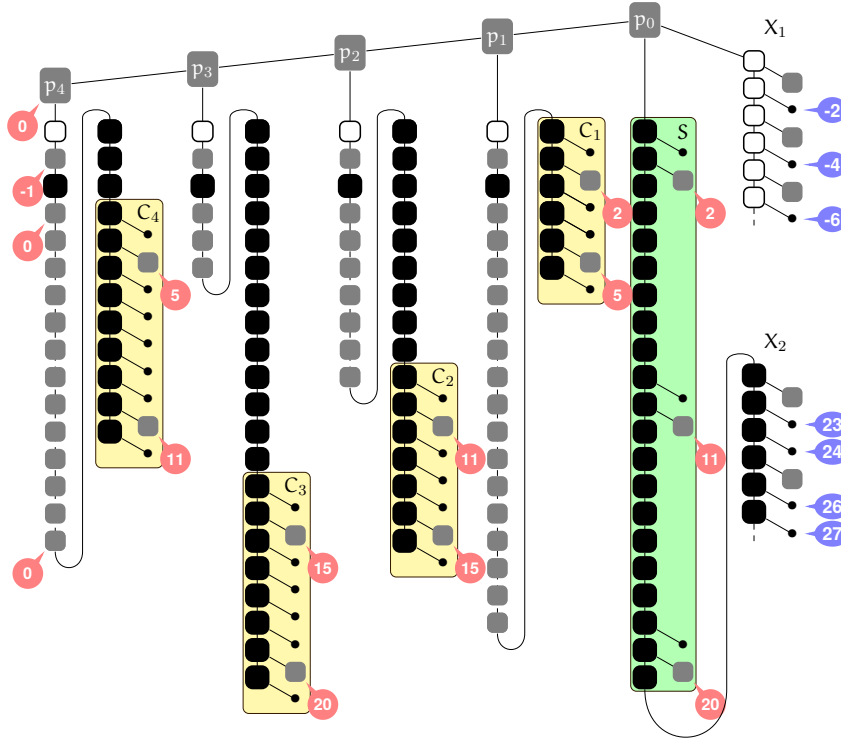
631 In this section, we consider an interval packing x assigning an interval of $[0, kB - 1[$ to
632 each item a_i . We write x_i such that $x(a_i) = [x_i, x_i + a_i - 1[$, and we color the tree T as
633 follows (see Figure 11):

- 634 ■ All nodes c_i in object gadgets, all non-separator nodes in S and all nodes c_i in X_1 are
635 colored ●,
- 636 ■ All nodes c_i in X_2 are colored ○.
- 637 ■ The first three nodes of P_i are colored ○ ● ●, and the last x_i nodes of P_i are colored ● ●
638 (note that P_i has length $kB \geq x_i + 3$ since $x_i + a_i < kB$ and $a_i \geq 5$).
- 639 ■ All remaining nodes are colored ●.

640 We show that this coloring is separated, in particular, we show that the level of each ●
641 node is of one of the following types, and that leaves are *not* of these types:

- 642 a. 0, 2 and $kB + 2$
- 643 b. $x_i + 2$ for each $1 \leq i \leq n$
- 644 c. $j - 1$ for $j \leq 0$, $j \equiv 0 \pmod{2}$
- 645 d. $kB + j + 4$ for $j \geq 0$, $j \equiv 0 \pmod{3}$

646 For the chain P , all nodes are ● and have level 0 (type a). For each P_i , there are ● nodes
647 at levels -1 and 0 (types a and c), and the chain ends at level x_i . For each object gadget
648 C_i , there are ● nodes at levels $x_i + 2$ (type b), and $x_i + a_i + 2$ (type b or a, since this
649 corresponds to the start of the next interval or to $kB + 2$). There are also leaves in C_i at
650 each level $x_i + j$ for $j = 1, 3, 4, \dots, a_i, a_i + 1, a_i + 3$ which are all values between 1 and $kB + 3$
651 and indeed do not correspond to any of the four types above. For gadget X_1 , there are ●
652 nodes at odd levels from -1 down to $-2kB + 1$ (type c), and leaves at even negative levels.
653 For the chain S , there are ● nodes attached at levels $iB + 2$ for each $0 \leq i \leq k$, which are
654 necessarily of the form $x_i + 2$ (type b) for some i (since each iB must be the start of some



■ **Figure 11** Example of the reduction with $n = 4$ items with sizes $\{3, 4, 5, 6\}$ to be sorted into $k = 2$ size-9 bins. A separated coloring is shown, corresponding to the solution $\{3, 6\}, \{4, 5\}$ (a selection of leaf and \bullet levels are depicted). Each item is mapped into a branch P_i followed by an object gadget C_i , containing 2 \bullet nodes separated by the size of the item. Leaves in object gadget enforce that any two gadgets may overlap only if the \bullet nodes are aligned. The bins are implemented using the separator sequence S , with \bullet nodes at every B th position, enforcing that series of consecutive items are packed into size- B bins. Finally, border gadgets X_1 and X_2 may not overlap with any other gadget, and enforce that all object gadgets and separators are packed together in a size- kB range of levels.

655 interval of x). Leaves in S are at level 1 and $kB + 1$, which are not of any type (in particular
 656 for type **b**, this is true since $a_i \geq 5$). Finally, for gadget X_2 , the \bullet nodes are of type **d**, and
 657 the leaves occupy remaining levels beyond $kB + 4$.

658 **From separated coloring to interval packing**

659 Suppose now that T admits a separated coloring χ , and consider the gadget X_1 . Its root is
 660 at level $\ell_{X_1} \in \{-1, 0, 1\}$, and by Proposition 11, there exists some $d_{X_1} \in \{-1, 1\}$ such that,
 661 for each level $\ell_{X_1} + d_{X_1}j$, there is a leaf (for even j) or a \bullet node (odd j). Without loss
 662 of generality, we assume that $d_{X_1} = -1$ (i.e., the chain in X_1 is \circ): if this is not the case
 663 we swap \circ and \bullet colors overall. Thus, there are leaves and \bullet nodes at alternating levels
 664 between -2 and $-2kB + 1$ (at least).

665 Consider the chain S . For any $0 \leq i \leq k$, the $(iB + 2)$ nd node of the chain cannot be \bullet
 666 (since it has a leaf sibling) so one of its two children must be \bullet . We write $s_0 \leq s_1 \leq \dots \leq s_k$
 667 for the levels of such \bullet nodes in ascending order: from the position of the nodes we have
 668 $s_{j+1} \leq s_j + B$. Furthermore, $s_0 \leq 3$ and $s_k \leq kB + 3$ (using the distances to the root).

669 Consider now X_2 . Its root is at most one level away from a separator, so at level ℓ_{X_2}

670 with $s_0 - 1 \leq \ell_{X_2} \leq s_k + 1$. By Proposition 11, there exists some $d_{X_2} \in \{-1, 1\}$ such that,
 671 for each level $\ell_{X_2} + d_{X_2}j$ with $1 \leq j \leq 2kB$, there is a \bullet node ($i \equiv 0 \pmod{3}$) or a leaf
 672 (otherwise). In particular, we necessarily have $d_{x_2} = 1$, since otherwise there would be two
 673 consecutive \bullet levels among levels $\{-2, -3, -4\}$, which would raise a conflict with X_1 .

674 For any $i \in [1, n]$, consider object gadget C_i . Its minimum level is ℓ_i with $-kB - n - a_i - 3 \leq$
 675 $\ell_i \leq kB + a_i + n + 3$, and by Proposition 11, for each level $\ell_i + j$ with $1 \leq j \leq a_i + 3$, there
 676 is a \bullet node ($j = 2, a_i + 2$) or a leaf (otherwise). In particular, $\ell_i \geq s_0 - 5$ (as otherwise
 677 there would be consecutive leaves at consecutive levels under $s_0 - 2$, in conflict with X_1)
 678 and $\ell_i + a_j \leq s_k + 5$ (otherwise there would be leaves at consecutive levels higher than
 679 $s_k + 3$, in conflict with X_2). Finally, since levels s_0 and s_k have \bullet nodes and $a_i \geq 5$, then
 680 for i such that $\ell_i \leq s_0 - 2$, we have $\ell_i = s_0 - 2$. Similarly, for i such that $\ell_i + a_i + 2 \geq s_k$,
 681 we have $\ell_i + a_i + 2 = s_k$. And for any i and j , if $\ell_i + 2 \leq s_j \leq \ell_i + a_i + 2$, we have
 682 $s_j \in \{\ell_i + 2, \ell_i + a_i + 2\}$.

683 Pick any two object gadgets $C_i, C_{i'}$ with $\ell_i \leq \ell_{i'}$. Then $\ell_i \neq \ell_{i'}$ (otherwise, since $a_i \neq a_{i'}$,
 684 there would be a conflict at level $\ell_i + \min\{a_i, a_{i'}\} + 2$), and $\ell_{i'} \geq \ell_i + a_i$ (otherwise, there
 685 would be a conflict at level $\ell_{i'} + 2$).

686 We now have all the tools to build an interval packing. We write $x_i = \ell_i - s_0 + 2$ and
 687 $\sigma_j = s_j - s_0$. By the remarks above, we have that intervals $[x_i, x_i + a_i - 1[$ are pairwise
 688 disjoint. Furthermore, they are all included in interval $[0, \sigma_k - 1[$. Since they have total
 689 size $\sum_{i=1}^n a_i = kB$ and $\sigma_k = s_k - s_0 \leq kB$, we have $\sigma_k = kB$, which is only possible with a
 690 fully \bullet chain S : so we get $\sigma_j = jB$ for all $0 \leq j \leq k$. And finally, if $\sigma_j \in [x_i, x_i + a_i - 1[$,
 691 then $\ell_i + 2 \leq s_j \leq \ell_i + a_i + 2$ which yields $s_j \in \{\ell_i + 2, \ell_i + a_i + 2\}$. This translates into
 692 $\sigma_j \in \{x_i, x_i + a_i\}$, so necessarily $\sigma_j = x_i$ and $\sigma_j - 1 \notin [x_i, x_i + a_i[$. Overall gadget levels
 693 relative to the first separator s_0 give a valid partition of $[0, kB - 1[$ into pairwise disjoint
 694 size- a_i intervals non-overlapping block border positions jB , so they give a valid INTERVAL
 695 PACKING solution.

696 **B** Non-separable target w/o isolated BPs (Proof of Proposition 10)

697 We start with the following remark:

698 **► Proposition 2.** *If u_0, \dots, u_k is a path in T and each u_i for even i has a leaf attached to it*
 699 *then, for any coloring χ of the path, we have $\chi(u_0) \in \{\bullet, \circ\}$ and $\chi(u_i) = \chi(u_0)$ for all i .*

700 **Proof.** Indeed, by the proper coloring constraint, every node with an attached leaf or with a
 701 leaf sibling may not be \bullet , so all $\chi(u_i) \in \{\bullet, \circ\}$ for all i . Moreover, there can be no direct
 702 edge between \circ and \bullet nodes, so $\chi(u_i) = \chi(u_{i-1})$ for all i which gives the desired property
 703 by induction. ◀

704 We now build a non-separable instance I without size-1 helix nor $(m_{3\bullet}, m_5)$ motif. Let
 705 $a \geq 2$ and $b \geq 2$ be even numbers. Let $T(a, b)$ be the gadget from Fig 8, containing a length- a
 706 path from the root to an internal node denoted t , and three length- b branches attached to t .
 707 Further attach a leaf to every node at an even distance from the root (except t itself). Note
 708 that all helices in $T(a, b)$ have length 2. The *level* of a copy of some $T(a, b)$ gadget is the
 709 level reached under node t of this gadget.

710 We build the instance I as a tree containing 5 copies of the gadget $T(a, b)$, precisely
 711 $I = (((T[10, 100], T[20, 100])), ((T[30, 100], T[40, 100])), T[50, 100])$.

712 First note that for a copy of gadget $T(a, b)$ at level ℓ in any separable coloring, there is
 713 a \bullet node at level ℓ , since the node t has three children and at least one must be \bullet . Also,
 714 there exist two integers u, v such that, for every $x \in [1, b]$, there is a leaf at level $\ell + ux$ if x

XXX:22 Exact linear-time RNA design for min Helix length 3

715 is odd, and level $\ell + vx$ if x is even. Indeed, pick one gray child U of t , and one non-gray
 716 child V . All vertices under U form an all-white or all-black branch by Proposition 2 (we let
 717 respectively $u = -1$ and $u = 1$), and vertices at levels $l + u, l + 3u, \dots, l + bu$ (or $l + (b - 1)u$)
 718 have a pending leaf. We similarly define $v = 1$ if V is black and $v = -1$ if V is white, and
 719 vertices at levels $l + 2v, l + 4v, \dots, l + bv$ (or $l + (b - 1)v$) have a pending leaf. From the
 720 above, if there are \bullet nodes at levels ℓ_1 and ℓ_2 with $\ell - b \leq \ell_1 < \ell_2 \leq \ell + b$, then $\ell_1 \not\equiv \ell_2$
 721 mod 2 (since otherwise, one of ℓ_1, ℓ_2 could be written as $\ell + ux$ with even x , so that level
 722 would be a leaf level).

723 Aiming at a contradiction, assume that I admits a separable coloring. Let $\ell_1 \leq \ell_2 \leq$
 724 $\ell_3 \leq \ell_4 \leq \ell_5$ be the levels of all five copies of the $T[a, b]$ gadgets of I , in ascending order.
 725 Then from the length of the branches from the root, we have $\ell_i \in [-50, 50]$ and $\ell_i \neq \ell_j$.
 726 Then by the remark above applied to the gadget with level ℓ_2 , we have $\ell_1 \not\equiv \ell_3 \pmod{2}$, and
 727 similarly using gadgets with level ℓ_4 we have $\ell_3 \not\equiv \ell_5 \pmod{2}$ and $\ell_1 \not\equiv \ell_5 \pmod{2}$, leading
 728 to a contradiction (any three integers such as ℓ_1, ℓ_3 and ℓ_5 may not have pairwise distinct
 729 parities).

730 **C** Leveraging random generators at fixed modular levels into a 731 uniform random generation of separated sequences

732 **► Theorem 12.** UNIFORM MODULO SEPARATED GENERATION *can be performed in an*
 733 *average-case complexity that is Fixed Parameter Tractable for the modulus parameter m .*

734 We consider a rejection-based approach, which starts by precomputing all $\#\text{Designs}_{\xi_L}$ in
 735 time $\Theta(n.m.2^m)$ (see Section 4.2), and accumulates them into $\mathcal{Z}_m := \sum_{\xi'_L \subseteq [0, m[} \#\text{Designs}_{\xi'_L}$.
 736 It then iterates the following steps until a suitable sequence is returned:

- 737 1. Choose some $\xi_L \subseteq [0, m[$ with probability $\mathbb{P}(\xi_L) = \#\text{Designs}_{\xi_L} / \mathcal{Z}_m$
- 738 2. Generate a ξ_L separated sequence w
- 739 3. Compute the number Ξ_w of $\xi'_L \subseteq [0, m[$ such that w is ξ'_L separated
- 740 4. Accept/return w with probability $1/\Xi_w$; Reject/restart from 1. otherwise.

Due to the full reset on each rejection, the emission probability p_w of any suitable w does
 not depend on the prior sequence of rejections (folklore, proven in [14, pp 77]), and we have:

$$p_w \propto \sum_{\substack{\xi_L \text{ such that } w \\ \text{is } \xi_L \text{ separated}}} \mathbb{P}(\xi_L) \times \mathbb{P}(w \mid \xi_L) \times \frac{1}{\Xi_w} = \sum_{\substack{\xi_L \text{ such that } w \\ \text{is } \xi_L \text{ separated}}} \frac{\#\text{Designs}_{\xi_L}}{\mathcal{Z}_m} \times \frac{1}{\#\text{Designs}_{\xi_L}} \times \frac{1}{\Xi_w}$$

741 Some terms directly cancel out and, by definition, we have $\sum_{\substack{\xi_L \text{ such that } w \\ \text{is } \xi_w \text{ separated}}} 1 = \Xi_w$. It follows
 742 that $p_w \propto 1/\mathcal{Z}_m$, a term that no longer depends on w , from which we conclude that the
 743 generation is uniform.

744 Complexity-wise, a prior accumulation of the 2^m terms $\#\text{Designs}_{\xi_L}$, each smaller than
 745 4^m , into a suitable data structure (see Lorenz and Ponty [10] for details) enables a random
 746 choice of ξ_L (Step 1.) in $\Theta(n.m)$. Once ξ_L is chosen, the above DP algorithm uniformly
 747 generates w in time $\Theta(m.n)$ (Step 2). The computation of Ξ_w (Step 3) is trivial and consists
 748 in identifying, in time $\Theta(n + m)$, the subset $\Phi_w \subseteq [0, m[$ of modular levels that are populated
 749 by neither leaves nor \bullet nodes in χ_w . Indeed, those levels represent the only degrees of
 750 freedom available while choosing a compatible ξ_L , the others modular values being forced
 751 to either \bullet or leaves. Since such modular values can be independently chosen to be in
 752 or out of ξ_L , then we have $\Xi_w = 2^{|\Phi_w|}$. Clearly, we have $\Xi_w \leq 2^m$, so the expectation
 753 of the number of (independent) rejections admits an upper bound in 2^m , and the overall
 754 average-case complexity is in $\Theta(n.m.2^m)$.