

# Lynx: A Programmatic SAT Solver for the RNA-Folding Problem

Vijay Ganesh<sup>1</sup>, Charles W. O’Donnell<sup>1</sup>, Mate Soos<sup>2</sup>, Srinivas Devadas<sup>1</sup>,  
Martin C. Rinard<sup>1</sup>, and Armando Solar-Lezama<sup>1</sup>

<sup>1</sup> Massachusetts Institute of Technology  
{vganesh, cwo, devadas, rinard, asolar}@csail.mit.edu  
<sup>2</sup> Security Research Labs  
mate@srlabs.de

**Abstract.** This paper introduces Lynx, an incremental programmatic SAT solver that allows non-expert users to introduce domain-specific code into modern conflict-driven clause-learning (CDCL) SAT solvers, thus enabling users to guide the behavior of the solver.

The key idea of Lynx is a *callback interface* that enables non-expert users to *specialize* the SAT solver to a class of Boolean instances. The user writes specialized code for a class of Boolean formulas, which is periodically called by Lynx’s search routine in its inner loop through the callback interface. The user-provided code is allowed to examine partial solutions generated by the solver during its search, and to respond by adding CNF clauses back to the solver dynamically and incrementally. Thus, the user-provided code can specialize and influence the solver’s search in a highly targeted fashion. While the power of incremental SAT solvers has been amply demonstrated in the SAT literature and in the context of DPLL(T), it has not been previously made available as a programmatic API that is easy to use for non-expert users. Lynx’s callback interface is a simple yet very effective strategy that addresses this need.

We demonstrate the benefits of Lynx through a case-study from computational biology, namely, the RNA secondary structure prediction problem. The constraints that make up this problem fall into two categories: structural constraints, which describe properties of the biological structure of the solution, and energetic constraints, which encode quantitative requirements that the solution must satisfy. We show that by introducing structural constraints on-demand through user provided code we can achieve, in comparison with standard SAT approaches, upto 30x reduction in memory usage and upto 100x reduction in time.

## 1 Introduction

Conflict-driven clause-learning (CDCL) Boolean SAT solvers have had a huge impact on a variety of domains ranging from program analysis to AI [3]. This success can partly be attributed to their simple interface and powerful heuristics. In many cases, a straightforward translation from a program analysis or AI problem into Boolean formulas in CNF (conjunctive normal form) format is sufficient to leverage the power of the solver. Unfortunately, there are many other important domains (e.g., biology) where straightforward translation of problems to CNF clauses leads to formulas that are too

large or complex for solvers to handle. For many of these domains, however, small domain-specific modifications to the solver can make SAT-based solution feasible. The challenge addressed by this paper is to enable users to make these small adaptations with minimal effort and without breaking subtle invariants in the solver implementation. The solution we provide allows for the resultant specialized solver to be adaptive, efficient for the problem-at-hand, and easy to build and maintain. Equally important, users are not burdened with knowing too much about the internals of SAT solvers and related technologies.

## 1.1 Our Contributions

- To address the problem described above, we created the solver Lynx that extends CryptoMiniSat [23] with an API allowing user-provided code to examine partial solutions generated by the SAT solver and add CNF clauses back to the solver in response. The added code is called inside the inner loop of the SAT solver, allowing the user to tightly integrate problem-specific clause-generation heuristics into the solver.

We call solvers extended in this way *programmatic*, i.e., the user can programmatically influence solver behavior and adapt it to their specific problem domain in ways that are difficult to achieve otherwise. Programmatic solvers address the “solvers are unpredictable black boxes” problem by giving users more control over their search heuristics.

- Using Lynx we developed the first SAT based tool for solving the RNA-folding prediction problem. We present a detailed experimental evaluation of our technique in comparison with standard approaches. We use the above-mentioned callback interface in efficiently translating the RNA prediction problem into Boolean formulas. The interface allows Lynx to incrementally translate the RNA-folding structure inside the inner loop of the SAT solver, allowing a tighter, highly targeted and more efficient integration of the SAT solver and the translator.

## 1.2 Existing Approaches to Incremental and Adaptive Solving

Incremental solvers, that use some form of *abstraction-refinement* [3], have been proposed as a solution to the above-mentioned issue of *simple but inefficient translations* from problems to Boolean formulas. Instead of translating the entire input problem-instance into a potentially very large Boolean formula in one step, *abstraction-refinement approaches* translate the input instance into Boolean formulas incrementally and call the solver on these incrementally generated formulas. Such formulas are abstractions of the input instance and are often easier to solve than the entire input instance. The solver terminates if it gets the correct result to the input instance by solving an abstraction. Otherwise the solver iteratively refines the abstractions as necessary until it gets the correct result. Typically these abstractions and their refinements are performed by a layer outside the inner loop of the SAT solver. For an excellent reference on abstraction-refinement strategies refer to the Handbook of Satisfiability [3].

Such incremental SAT solvers with an outside abstraction-refinement loop are relatively easy to build. However, the problem with such an approach is that it may not be

the most efficient for the problem-at-hand. Indeed, Ohrimenko et al. [18] have proposed incremental translation of problems to SAT where the integration of the solver and the incremental translation is much tighter and more efficient than an outer layer translator. However, their implementation is non-adaptive, and is specific to a class of difference logic formulas — they do not provide an API for users to easily adapt or extend the solver for a previously unknown class of Boolean formulas.

An example of an API that allows users to adapt or extend solvers is the powerful idea of DPLL(T) [11] aimed at solving Boolean combination of formulas in rich theories such as integer linear arithmetic, uninterpreted functions and datatypes (aka SMT solvers [3]). In this approach, there is a tight integration of a CDCL SAT solver with a *theory solver* (aka a T-solver) that can handle conjunction of constraints represented in a rich logic. The CDCL SAT solver does the search on the Boolean structure of the formula without knowing the semantics of the literals, while the T-solver reasons about the literals themselves adding any new derived literals back to the Boolean CDCL solver appropriately. The tight integration enables the T-solver to influence the CDCL solver's behavior in ways not possible otherwise, and the resultant combination is typically a solver that can handle arbitrary Boolean combination of theory formulas efficiently.

A lay non-expert user could implement a “T-solver” using the DPLL(T) framework that reasons about a specific domain (say, theory of RNA folding) and adds constraints incrementally to the SAT solver. The resultant combination can be a powerful incremental domain-specific solver. However, the DPLL(T) API imposes strict requirements on the user-specified code (T-solver) to ensure that the resultant combination is sound and complete. Such requirements make perfect sense for constructing powerful SMT solvers with complex T-solvers, the problem for which the DPLL(T) approach was originally proposed. However, for the lay non-expert users such requirements may be onerous, and may not be essential. Lynx, by contrast provides a simple interface which is relatively easy to prove correct and is tailored for problem-specific extensions.

### 1.3 RNA-Folding with Lynx

To explore the benefits of using the Lynx's callback interface, we applied the technique to the problem of RNA folding. This is an application of significant practical relevance: understanding RNA folding is crucial to understanding a number of biological processes, including the replication of single-strand RNA viruses such as the poliovirus which causes polio in humans. Moreover, RNA prediction actually shares important similarities with other structure prediction problems of biological interest. This problem is particularly suitable to benchmark our approach. First, a SAT based solution to this problem is desirable because it gives researchers the ability to easily experiment with different formulations for the basic problem. Moreover, previous work in the literature has succeeded in formalizing the problem in a form that lends itself very naturally to solution with a Boolean SAT solver. SAT based solutions, however, have been elusive because the standard encoding leads to Boolean SAT instances that are too big for solvers to handle. Using Lynx's callback interface allowed us to encode instances of the RNA folding problem in a memory efficient manner, producing the first successful SAT based solution to this problem. The resultant incremental (or online abstraction-refinement) solver led to a 30-fold reduction in the amount of memory required to solve

some of these problems compared to standard SAT approach, and demonstrated dramatic time improvements over standard abstraction refinement techniques.

**Paper Layout.** In Section 2 we provide a detailed overview of our incremental approach. In Section 3 we provide a self-contained description of the RNA-folding structure prediction problem. In Section 4 we provide detailed description of our experimental setup and results. We review the related work in Section 5, and conclude in Section 6.

## 2 Incrementality in Lynx

This section details how the callback interface in Lynx makes the solver incremental, what we sometimes also refer to as online abstraction-refinement or OAR. In order to facilitate the description, let us introduce a simple running example which shares some features with the more complex biology application.

The running example is a formula of the form  $P(\mathbf{x}) \wedge C(\mathbf{x})$  over a vector  $\mathbf{x} = \langle x_0, x_1, \dots, x_N \rangle$  of Boolean variables, where  $P(\mathbf{x})$  consists of some arbitrary set of constraints and  $C(\mathbf{x})$  is a cardinality constraint that says that no more than 2 bits in  $\mathbf{x}$  can be set to 1.

$$C(\mathbf{x}) \equiv \forall_{i \neq j \neq k} (\neg x_i \vee \neg x_j \vee \neg x_k)$$

The above definition of  $C(\mathbf{x})$  can be trivially encoded as a set of  $N^3$  CNF clauses — too many for large values of  $N$ . For this specific case, more efficient encodings exist using only  $O(N)$  clauses, but they are more complicated and require the introduction of additional SAT variables. By contrast, online abstraction refinement allows us to use the simple encoding without having to pay the price of introducing  $N^3$  clauses.

The first step in using OAR is to divide the problem into a core set of clauses added to the solver from the very beginning, and a different set of *dynamic* clauses added to the solver incrementally by a *callback function*. The callback function is a user-provided function  $\mathcal{M}$  producing a set of clauses given a partial assignment to the variables of the solver’s input instance. A partial assignment sets each variable in the problem to either 1, 0, or  $\perp$  (undefined), and is represented as a vector  $\mathbf{t} \in \{0, 1, \perp\}^N$ .

In the case of the example, we define  $P(\mathbf{x})$  to be the core clauses, and  $C(\mathbf{x})$  to be the clauses added dynamically by a callback function defined as:

$$\mathcal{M}(\mathbf{t}) \equiv \{(\neg x_i \vee \neg x_j \vee \neg x_k) \mid i \neq j \neq k \wedge t_i = t_j = t_k = 1\}$$

This callback function receives a partial assignment  $\mathbf{t}$ , and returns a set of clauses of the form  $(\neg x_i \vee \neg x_j \vee \neg x_k)$  where  $x_i$ ,  $x_j$  and  $x_k$  are variables set to 1 in the partial assignment (i.e.,  $t_i = t_j = t_k = 1$ ). The clauses produced by the callback function eliminate those incorrect solutions that would have been eliminated by  $C(\mathbf{x})$ , so running the solver with constraints  $P(\mathbf{x})$  and callback function  $\mathcal{M}$  is the same as solving  $P(\mathbf{x}) \wedge C(\mathbf{x})$ .

Lynx incorporates the callback function into the solution process by invoking it periodically with the current partial assignment. If the callback function returns any clause, these are incorporated into the problem. This process continues until an assignment  $q$

is found such that: a)  $q$  satisfies all the core constraints, b)  $q$  satisfies all the constraints ever produced by the callback function, and c) the callback function produces an empty set of clauses when applied to  $q$  indicating that the process can be terminated. If the input problem is unsatisfiable, the solver with the callback function is guaranteed to report unsatisfiable and terminate. It is possible for the user-code, without any restrictions, to render the combination of base solver plus user-code incomplete. However, we can impose some minimal conditions on the user-code such that the combination is guaranteed to be a complete decision procedure. In particular, one such condition is as follows: assume the desired input instance to be solved is  $P(x) \wedge C(x)$ , and  $P(x)$  is input to the base solver. Then, the user-code must "encode"  $C(x)$  exactly. Imposing this particular condition on the user-code is guaranteed to render the combination complete.

### 3 Biological Problem Overview

RNA is a versatile polymer essential to all of life. A chain of covalently bound nucleotides, RNA classically acts as a cellular messenger which duplicates DNA sequence information in the nucleus/nucleoid and transports that code to ribosomes for the construction of proteins. However, this chain can also fold in on itself into a 3-dimensional globular molecule which catalyzes biological reactions by itself. In fact, modern studies have suggested that such non-coding RNA (ncRNA) may play even a bigger cellular role than messenger RNA, with significant effects on metabolism, signal transduction, gene regulation, and chromosome inactivation. Such RNA function is determined by its nucleotide composition and 3-dimensional structure, however, relatively little ncRNA structural data is known [25], severely limiting our understanding of these mechanisms. Therefore, algorithmic prediction of RNA structure from its nucleotide sequence has been a longstanding computational goal.

#### 3.1 Structure Prediction via SAT

The computational problem we address is "how to correctly attribute a unique structural state to each nucleic acid (or groups of nucleic acids) within an RNA polymer sequence". This problem has a long history of solutions based on many different algorithmic models — the most successful of which using a recursive, grammatical approach introduced by Zuker [26]. In this biophysical model, each nucleotide is allowed to form a pairwise bond with another, and each pair is assigned an energetic cost based on spatially adjacent nucleotide types [16]. The most likely structure is predicted by optimizing pairing configuration according to a fixed thermodynamic scoring system (energy minimization). Efficient computation is made possible through the imposition of specific, often biologically-inspired model restrictions — for example, limiting base-pairs to be sequentially nested (i.e. no "pseudoknots") and scoring only a subset of all potential energetic interactions (i.e. only Watson-Crick or wobble base-pairs). Unfortunately, this entangles the optimization techniques used with a particular set of biological assumptions. While these methods have shown good predictive accuracy, changes to the algorithm can be difficult to implement as new scientific data comes to light. For example, it has been shown that a more complex description of the RNA interaction energetics can lead to greatly improved results [19].

We propose a declarative approach for the structure prediction problem, providing a decoupled platform for reasoning about biological concepts in clear, succinct rules, backed by the powerful generic optimization of CDCL SAT solvers. This allows biological models to be tested and flexibly refined using a constraint-based philosophy, independent of performance improvements to the underlying solver.

To study this approach, we have implemented an RNA structure prediction algorithm using Lynx. Rather than comparing the benefits and disadvantages of different biological models, we base our implementation on an RNA scoring model recently proposed by Kato, et al. for integer programming optimization [20]. Although other models outperform this scoring system's accuracy, we believe our results are easily generalizable to greater classes of RNA structures [4] and more complex (non-RNA) structure prediction problems in general.

To implement energy minimization as a SAT-based decision procedure, we pose the question of whether an assignment exists that is lower than a certain energy threshold and perform iterative binary search. Despite this search routine, this approach can often be more efficient than the dynamic programming methods used by grammatical models as the problem can be finely partitioned into smaller jobs that are run in parallel. Further, when a sub-optimal solution is sufficient, this method quickly short-circuits, along with a guarantee of how near the solution is to optimality.

### 3.2 RNA Secondary Structure Prediction with Pseudoknots

The RNA prediction algorithm described here differentiates itself from classical prediction methods in its goal of predicting pseudoknots. Earlier grammar-based predictors allowed only base-pairs to occur in a recursively nested fashion (i.e. for every base-pair  $i$ - $j$  there exists no base-pair  $k$ - $l$  such that  $i < k < j < l$ ) to enable highly efficient energy minimization via dynamic programming. However, pseudo-knotted structures which break this restriction are known to be essential to a number of functions, such as the Diels-Alder ribozyme and mouse mammary tumor virus [24]. However, predicting pseudoknotted structures is computationally much harder with fewer solutions [17,20,21]. In fact, the prediction of truly arbitrary pseudoknots has been shown NP-complete [14], and classes of pseudoknotted structures are often more easily defined by the algorithms which recognize them rather than their biological significance [7]. This motivates the use of a declarative approach, allowing easy exploration of different trade-offs between representation and optimization, especially if the underlying scoring system is changed from the standard Watson-Crick/wobble base-pair models to more complex interactions [19]. However, in the remainder of this work we restrict ourselves to the model proposed by Kato, et al. [20].

### 3.3 Encoding RNA Structure Prediction in SAT

Our SAT encoding is formulated by two sets of constraints, structural and energetic, that control the assignment of a vector of free variables which represent the final structural solution. The assignment of each free variable indicates whether two nucleotides are base-paired in the final RNA structure, fixed by structural constraints and an associated energetic score. Figure 1 depicts this formulation.

**Solution Variables.** The set of all properly-nested base-pairs within the final output RNA structure is represented by the variables  $X_{i,j}$ : where  $i$  and  $j$  indicate the sequence position of two nucleotides, a value  $X_{i,j} = \mathbf{T}$  indicates a hydrogen bond base-pair exists between nucleotides at  $i$  and  $j$ , and  $X_{i,j} = \mathbf{F}$  indicates that no base-pairing occurs between positions  $i$  and  $j$ . The set of pseudoknotted base pairs that cannot be properly nested are similarly represented by the independent variables  $Y_{i,j}$ . In this way pseudoknots are represented solely by the alignment of properly-nested  $X_{i,j}$  pairs and properly-nested  $Y_{i,j}$  pairs. Since RNA structure permits any nucleotide position  $i$  to pair with any other position  $j$ , a valid biological structure requires a complete assignment of all  $X_{i,j}$ s and  $Y_{i,j}$ s for every  $i, j$  ( $0 \leq i, j < \text{length}(\text{sequence})$ ). Therefore, the number of solution variables, the number of resultant constraints, and thus the difficulty of the SAT problem depends directly on the sequence length of the input RNA.

**Structural Constraints.** The structural representation places requirements on the assignment of the solution bits  $X_{i,j}$  and  $Y_{i,j}$  to ensure a biologically consistent structure. Therefore, we declare the following constraints, which must be satisfied in any valid solution:

- Every position  $i$  can at most pair with one other position  $j$ , independent of whether that pairing is properly-nested or a pseudoknot (Figure 1(a-d)). Four straightforward constraints ensure this:

$$\begin{aligned} \forall i, j, k, \quad i < j < k \\ (X_{i,j} \wedge X_{j,k}) &= \mathbf{F} \wedge (Y_{i,j} \wedge Y_{j,k}) = \mathbf{F} \wedge \\ (X_{i,j} \wedge Y_{j,k}) &= \mathbf{F} \wedge (Y_{i,j} \wedge X_{j,k}) = \mathbf{F} \end{aligned}$$

- All base-pairs  $i, j$  are properly nested or a pseudoknot, but not both (Figure 1(e)):

$$\forall i, j \quad (X_{i,j} \wedge Y_{i,j}) = \mathbf{F}$$

- We define all  $X_{i,j}$  and  $Y_{i,j}$  base-pairs to be independently knot-free (Figure 1(f-g)):

$$\begin{aligned} \forall i, j, k, l, \quad i < k < j < l \\ (X_{i,j} \wedge X_{k,l}) &= \mathbf{F} \wedge (Y_{i,j} \wedge Y_{k,l}) = \mathbf{F} \end{aligned}$$

- We only permit bifurcations within the “normal” base-pairs in  $X_{i,j}$  since pseudoknots are rare and deserve distinct energetic treatment. Therefore (Figure 1(h)):

$$\forall i, j, k, l, \quad i < k < j < l \quad (Y_{i,j} \wedge Y_{k,l}) = \mathbf{F}$$

- Finally, the class of structures with “double-crossing” pseudoknots are rare and present unusual energetics which are not handled by the energy model we use, thus we constrain pseudoknots to only cross at most once (Figure 1(i-j)):

$$\begin{aligned} \forall i, j, k, l, m, n, \quad i < m < j < k < n < l \\ (X_{i,j} \wedge Y_{m,n}) &\implies (X_{k,l} = \mathbf{F}) \quad \wedge \\ (X_{k,l} \wedge Y_{m,n}) &\implies (X_{i,j} = \mathbf{F}) \end{aligned}$$

**Energetic Constraints.** The total energy of an RNA structure is defined as the sum of experimentally-derived energy parameters [26,20] for every constituent base-pair *stack*, where a stack indicates two adjacent base pairs, e.g.  $X_{i,j}$  and  $X_{i+1,j-1}$ . Energy parameters are given in terms of base-pair stacks because nucleotide  $\pi$ -orbital overlap serves as a dominant stabilizing factor in RNA structure. Thus, an energy value is assigned to every base-pair stack  $X_{i,j}X_{i+1,j-1}$  according to the four nucleotide types at sequence positions  $i, j, i+1$ , and  $j-1$  (Parameters found in [20]). By including a logical adder of all possible energetic assignments, we can then define a valid solution as an assignment of  $X_{i,j}$  and  $Y_{i,j}$  (subject to structural constraints), where the output of the adder overcomes some minimum threshold energy  $E_{threshold}$  (the energy bound). As a logical declaration, we write:

$$\forall i, j, i < j \ (X_{i,j} \wedge X_{i+1,j-1}) = \mathbf{T} \Rightarrow (E_{X_{i,j}} = EnergyConstant_{(i,j,i+1,j-1)}) \ \wedge$$

$$(Y_{i,j} \wedge Y_{i+1,j-1}) = \mathbf{T} \Rightarrow (E_{Y_{i,j}} = EnergyConstant_{(i,j,i+1,j-1)}) \ \wedge$$

$$(X_{i,j} \wedge X_{i+1,j-1}) = \mathbf{F} \Rightarrow (E_{X_{i,j}} = 0) \ \wedge$$

$$(Y_{i,j} \wedge Y_{i+1,j-1}) = \mathbf{F} \Rightarrow (E_{Y_{i,j}} = 0),$$

where  $EnergyConstant(i, j, i + 1, j - 1)$  indicates the energy score of the four nucleotides found at positions  $i, j, i + 1$ , and  $j + 1$  base-pairing and stacking, and


$$\sum_{\forall i,j} (E_{X_{i,j}} + E_{Y_{i,j}}) \geq E_{threshold}.$$

Finally, to enforce that all assigned base-pairs are accounted within the adder by stacking energy parameters, we require:

$$\forall i, j \quad s.t. \ i < j$$

$$\overline{(X_{i-1,j+1} \wedge X_{i,j} \wedge X_{i+1,j-1})} = \mathbf{F} \ \wedge$$

$$\overline{(Y_{i-1,j+1} \wedge Y_{i,j} \wedge Y_{i+1,j-1})} = \mathbf{F}$$



**Fig. 1.** RNA Constraints



## 4 Experimental Results

In this section we describe the results of our experimental evaluation of Lynx and competing approaches over input tests obtained from a set of RNA sequences. As described in detail in 3, we solve the two dimensional RNA optimum structure prediction problem (where the structures may have pseudoknots). We ran all experiments on a 3GHz Intel Xeon X5460 with 64GB of RAM and a 6MB L2 cache with 1 hour timeout per SAT instance.

### 4.1 Description of Input Tests

We acquired a set of benchmark RNA sequences and structures from the PseudoBase website [1]. These RNA sequences are widely used by computational biologists for a variety of structure prediction tasks. The biological accuracy of our lowest-energy structure predictions were verified to agree with Kato, et al. [20], whose scoring model we duplicate. Recall that the optimization problem is treated as a series of decision problems performing a binary search of the energy space. For each RNA sequence, a corresponding SAT instance is therefore constructed containing the energy and structural constraints along with an energetic bound that captures the minimum and maximum allowed energy for that step in the binary search. Given the precision of our energy model a search depth of 10 sufficed to identify the minimum energy structure of any structure tested.

### 4.2 Experimental Methodology

We solve the structure prediction problem using the following three methods:

- **Baseline Approach Using CryptoMiniSat (BA):** A standard encoding of our problem in SAT. We generate the complete SAT encoding (with XOR clauses as appropriate) of the RNA secondary structure prediction problem, then use CryptoMiniSat to solve this problem. We also used MiniSat2 [9], and found that for this problem its performance is similar to CryptoMiniSat [23].
- **Offline Abstraction Refinement (OFFA):** An encoding of our problem using established refinement techniques. Starting with only the energy constraints from the SAT encoding of the RNA structure prediction problem to form the abstracted constraint, we use offline abstraction refinement to obtain a solution to the complete structure prediction problem. Each refinement step uses CryptoMiniSat to solve the current SAT problem, computes the set of constraints from the complete structure prediction problem that are inconsistent with this solution, and generates a new problem by incrementally adding these constraints to the current problem in SAT. The refinement process continues until it produces a solution to the complete input problem.
- **Online Abstraction Refinement (ONA):** The methodology enabled by our tool Lynx. Starting with only the energy constraints from the SAT encoding of the RNA structure prediction problem to form the abstracted constraint, we use online abstraction refinement to obtain a solution to the complete structure prediction problem. After each CryptoMiniSat propagation step, the constraint manager examines

the current partial solution to find the set of constraints from the full structure prediction problem that conflict with the current solution. It then incrementally adds these constraints to the current problem before CryptoMiniSat takes the next partial solution step. The difference between the Offline (OFFA) and Online (ONA) approaches is the granularity of the refinement steps. Each refinement step in the OFFA version takes place only after CryptoMiniSat produces a complete solution to the current problem. Each refinement step in the ONA version, in contrast, takes place at the much finer granularity, every time CryptoMiniSat extends the current partial solution.

**Table 1.** Comparison of running times between Baseline (BA), Offline (OFFA), and Online (ONA) methods. Total cumulative time (across all solver instances during search) is reported, broken down by the amount of time spent in the SAT solver versus the amount of time spent in refinement. The number of refinement steps involved is also given. T.O. indicates that a timeout occurred after 1hr of an individual SAT solver instance.

RNA	sequence length	Baseline (sec)	Offline Tot(sec)=SAT+Ref (# steps)	Online Tot(sec)=SAT+Ref (# steps)
PKB115	24	<b>1.4</b>	<b>1.7</b> = 1.3+0.4 (205)	<b>0.8</b> = 0.6+0.2 (2,538)
PKB102	24	<b>1.3</b>	<b>1.0</b> = 0.7+0.3 (129)	<b>0.6</b> = 0.5+0.1 (1,766)
PKB119	24	<b>2.1</b>	<b>3.6</b> = 3.0+0.6 (266)	<b>1.6</b> = 1.3+0.3 (4,108)
PKB103	25	<b>3.1</b>	<b>6.6</b> = 5.4+1.2 (417)	<b>3.5</b> = 3.1+0.4 (6,191)
PKB123	26	<b>5.6</b>	<b>24.7</b> = 22.7+2.0 (597)	<b>7.4</b> = 6.8+0.6 (8,980)
PKB154	26	<b>2.5</b>	<b>3.8</b> = 3.2+0.6 (236)	<b>1.9</b> = 1.7+0.2 (4,070)
PKB152	26	<b>3.2</b>	<b>6.2</b> = 5.2+1.0 (255)	<b>2.3</b> = 2.0+0.3 (5,528)
PKB126	27	<b>4.0</b>	<b>6.6</b> = 5.5+1.1 (384)	<b>2.8</b> = 2.5+0.3 (5,874)
PKB124	29	<b>4.7</b>	<b>5.1</b> = 4.4+0.7 (262)	<b>2.3</b> = 2.1+0.2 (4,635)
PKB100	31	<b>11.0</b>	<b>52.3</b> = 49.4+2.9 (315)	<b>6.8</b> = 6.0+0.8 (11,890)
PKB105	32	<b>17.0</b>	<b>58.3</b> = 54.0+4.3 (1004)	<b>18.1</b> = 17.0+1.1 (16,817)
PKB118	33	<b>13.7</b>	<b>32.8</b> = 29.6+3.2 (591)	<b>8.2</b> = 7.4+0.8 (12,878)
PKB120	36	<b>36.1</b>	<b>571.1</b> = 560.6+10.5 (652)	<b>24.1</b> = 21.9+2.2 (26,370)
PKB065	46	<b>185.1</b>	<b>11,341.9</b> = 11,298.7+43.2 (1,344)	<b>112.7</b> = 108.1+4.6 (50,508)
PKB205	48	<b>388.6</b>	<b>T.O.</b>	<b>391.6</b> = 381.9+9.7 (72,922)
PKB147	51	<b>1,917.3</b>	<b>T.O.</b>	<b>1,087.9</b> = 1,067.2+20.7 (131,321)
PKB248	66	<b>T.O.</b>	<b>T.O.</b>	<b>T.O.</b>
PKB072	67	<b>5,352.6</b>	<b>T.O.</b>	<b>2,414.1</b> = 2,367.6+46.5 (286,881)

### 4.3 Results

Table 1 presents the total execution times required for the different methods to solve the RNA structure prediction problems. We ran each method with a timeout of 3600 seconds for each SAT solution problem (i.e., each binary search step). Each row in the table corresponds to a single RNA. The first column is the number of base pairs in the RNA sequence. The next column presents the time (in seconds) required for the BA method to solve the problem. Recall that each problem requires the solution of 10 SAT instances; the reported total time is the sum of the 10 individual SAT solution times. The next column presents data from the OFFA method and is of the form  $t = s + c(r)$ .

**Table 2.** Comparison of memory usage between Baseline (BA), Offline (OFFA), and Online (ONA) methods. Given is the maximum memory (in MB) required throughout all SAT solver instances, along with the sum of the total number of clauses (in thousands) both input and generated during refinement. T.O. indicates that a timeout occurred after 1hr of an individual SAT solver instance.

RNA	sequence length	Baseline Mem(MB) / Clauses	Offline Mem(MB) / Clauses	Online Mem(MB) / Clauses
PKB115	24	<b>5.0</b> / 3,223k	<b>5.0</b> / 94k	<b>72.1</b> / 82k
PKB102	24	<b>5.0</b> / 3,219k	<b>5.0</b> / 86k	<b>5.0</b> / 75k
PKB119	24	<b>5.0</b> / 3,240k	<b>5.0</b> / 130k	<b>5.0</b> / 104k
PKB103	25	<b>5.0</b> / 4,142k	<b>16.5</b> / 174k	<b>5.0</b> / 136k
PKB123	26	<b>43.4</b> / 5,244k	<b>19.7</b> / 226k	<b>74.7</b> / 168k
PKB154	26	<b>5.0</b> / 5,204k	<b>5.0</b> / 128k	<b>5.0</b> / 106k
PKB152	26	<b>5.0</b> / 5,220k	<b>16.6</b> / 174k	<b>5.0</b> / 128k
PKB126	27	<b>72.1</b> / 6,544k	<b>74.5</b> / 171k	<b>5.0</b> / 129k
PKB124	29	<b>5.0</b> / 10,076k	<b>5.0</b> / 142k	<b>5.0</b> / 108k
PKB100	31	<b>90.5</b> / 16,937k	<b>23.9</b> / 376k	<b>90.0</b> / 231k
PKB105	32	<b>157.4</b> / 20,584k	<b>75.9</b> / 448k	<b>95.7</b> / 260k
PKB118	33	<b>131.9</b> / 24,870k	<b>23.2</b> / 355k	<b>22.8</b> / 227k
PKB120	36	<b>276.0</b> / 42,698k	<b>76.7</b> / 729k	<b>75.3</b> / 369k
PKB065	46	<b>1,011.8</b> / 196,236k	<b>150.6</b> / 341k	<b>122.9</b> / 595k
PKB205	48	<b>1,221.3</b> / 255,861k	<b>T.O.</b>	<b>145.0</b> / 808k
PKB147	51	<b>1,988.9</b> / 373,294k	<b>T.O.</b>	<b>188.7</b> / 1,322k
PKB248	66	<b>T.O.</b>	<b>T.O.</b>	<b>T.O.</b>
PKB072	67	<b>9,046.5</b> / 2,031,362k	<b>T.O.</b>	<b>313.1</b> / 2,652k

Here  $t$  is the total time required to solve the structure prediction problem (the sum of the solution times for the 10 SAT problems),  $s$  is the amount of time spent in the SAT solver,  $c$  is the amount of time spent in the constraint manager, and  $r$  is the total number of refinement steps (summed over all 10 SAT problems). The last column presents data from the ONA method and is also of the form  $t = s + c(r)$ .

Up to problem PKB124, the solution times for all of the methods are roughly comparable: each is less than ten seconds and within a factor of two for the same RNA sequence. For larger problems the OFFA approach starts to exhibit substantially larger solution times than either BA or ONA approaches; for the largest problems in our benchmark set OFFA times out. For two of the largest three problem sizes BA is roughly a factor of two slower than ONA; BA times out for PKB248.

We note that there is a substantial difference between the number of refinement steps that the ONA and OFFA methods perform — OFFA typically performs hundreds of (relatively coarse grain) refinement steps, while ONA performs thousands of (fine grain) refinement steps. These data indicate that, as expected, the SAT solver can respond much more quickly to fine grain than to coarse grain refinement steps, but that the ONA method requires more fine grain steps to reach a solution.

Table 2 presents the maximum amount of memory required to solve the structure prediction problem (this is the maximum over all runs of the SAT solver of the amount of memory that the SAT solver consumes) and the total number of clauses for each

RNA. For the OFFA and ONA methods, the total number of clauses is the sum over all binary search steps of the number of clauses in the problem at the final refinement step. Each entry of the table is in the form  $m/c$ , where  $m$  is the maximum memory and  $c$  is the number of clauses. Both the OFFA and ONA methods generate problems with substantially smaller numbers of clauses than the BA method (BA typically generates hundred to thousand times as many clauses OFFA and ONA typically generate). For larger RNA sequences, these larger clause sizes translate into substantially larger memory requirements for the BA method — OFFA and ONA never go above several hundred Mbytes, while BA starts requiring more than 1Gbyte of memory for the larger sequences.

#### 4.4 Discussion

These data highlight how the ONA method is able to combine the benefit of small memory requirements, which it shares with OFFA, and feasible execution times, which it shares with BA (further note that ONA often exhibits roughly a factor of two performance advantage over BA). We attribute these characteristics to, first, the ability of the ONA method to effectively find relatively small problems whose solution also happens to be a solution of the complete structure prediction problem, and second, the ability of the ONA method to efficiently guide the SAT solver to the solution through fine-grain corrections to partial solution missteps. A comparison with the OFFA method illustrates how quickly correcting any missteps on the part of the SAT solver (by operating the refinement steps after every intermediate SAT solver decision rather than after every complete solution) can deliver very efficient solution times even in situations where the more coarse OFFA approach fails to solve the problem in an acceptable amount of time.

## 5 Related Work

There has been a lot of recent work on incremental SAT solvers [18], DPLL(T) [11], abstraction-refinement based techniques in the context of model-checking and decision procedures for SMT theories [2]. We summarize the related work, and contrast Lynx with other tools.

**Incrementality, Extensibility and SAT Solvers.** The work that is closest to ours is by Stuckey et al. [18] and the related idea of DPLL(T) [11]. Our work is different from Stuckey et al. in the mechanism employed to implement incrementality, namely, a callback interface. Our approach is more flexible in the sense that it can be used to expose other internals of SAT solvers (e.g., branching heuristics or restart triggers) to lay non-expert users. While DPLL(T) is a very powerful idea, it places more requirements on user-code (to ensure completeness and soundness) and is probably best used by experts.

**Abstraction-Refinement in Decision Procedures.** The idea of counter-example guided abstraction refinement was originally developed in the context of model-checking [6]. Since then the basic idea has been adapted in different ways to solve the satisfiability problems of SMT theories [2]. Kroening, Ouaknine, Seshia, and Strichman [13] were

the first to adapt CEGAR to deciding quantifier-free Presburger arithmetic. More recently, Brummayer and Biere give a new technique that allows early termination of an under-approximation refinement loop even when the original formula is unsatisfiable [5]. Ganesh and Dill proposed the use of abstraction-refinement for deciding the theory of arrays [10].

**RNA Secondary Structure Prediction.** Zuker introduced the first optimal algorithms for RNA secondary structure prediction based on a dynamic programming solution to energy minimization [26], although many improved predictors have followed [15]. Non-thermodynamic approaches have also met success through the use of phylogenetic relationships [12], or via machine learning [8]. The first efficient thermodynamic-based algorithm for predicting RNA pseudoknotted secondary structure was introduced by Rivas and Eddy (PKNOTS [22]). Subsequent algorithms have recognized alternate classes of pseudoknots or improved upon the efficiency of solutions [17,4], including the IP formulation focused on in this paper [20], and heuristics such as HotKnots [21].

## 6 Conclusions

We present Lynx, a programmatic incremental SAT solver that allows non-expert users to easily introduce domain-specific or instance-specific code into modern CDCL SAT solvers, thus enabling users to control the behavior of the solver in ways not possible otherwise. While there has been work on incremental SAT [18] before and related ideas such as DPLL(T), Lynx's interface is simple to use and the requirements placed on user code are minimal. The approach is a template on how to expose other internals of the SAT solver to non-expert users in a easy-to-use and intuitive way. We demonstrate the benefits of Lynx through a first-of-its-kind solver case-study from computational biology, namely, RNA secondary structure prediction.

## References

1. Pseudobase RNA sequence, Most widely used database for research on RNA sequences with Psuedoknots. website, <http://pseudobaseplusplus.utep.edu/>
2. SMTLIB website, <http://combination.cs.uiowa.edu/smtlib/>
3. Biere, A., Heule, M.J.H., van Maaren, H., Walsh, T. (eds.): *Handbook of Satisfiability*. Frontiers in Artificial Intelligence and Applications, vol. 185. IOS Press (February 2009)
4. Bon, M., Vernizzi, G., Orland, H., Zee, A.: Topological classification of RNA structures. *J. Mol. Biol.* 379(4), 900–911 (2008)
5. Brummayer, R., Biere, A.: Effective Bit-Width and Under-Approximation. In: Moreno-Díaz, R., Pichler, F., Quesada-Arencibia, A. (eds.) EUROCAST 2009. LNCS, vol. 5717, pp. 304–311. Springer, Heidelberg (2009)
6. Clarke, E., Grumberg, O., Jha, S., Lu, Y., Veith, H.: Counterexample-guided abstraction refinement for symbolic model checking. *J. ACM* 50(5), 752–794 (2003)
7. Condon, A., Davy, B., Rastegari, B., Chao, S., Tarrant, F.: Classifying RNA pseudoknotted structures. *Theoretical Computer Science* 320, 35–50 (2004)
8. Do, C., Woods, D., Batzoglou, S.: CONTRAfold: RNA secondary structure prediction without energy-based models. *Bioinformatics* 22(14), e90–e98 (2006)

9. Een, N., Sorensson, N.: An extensible SAT-solver. In: Proc. Sixth International Conference on Theory and Applications of Satisfiability Testing, pp. 78–92 (May 2003)
10. Ganesh, V., Dill, D.L.: A Decision Procedure for Bit-Vectors and Arrays. In: Damm, W., Hermanns, H. (eds.) CAV 2007. LNCS, vol. 4590, pp. 519–531. Springer, Heidelberg (2007)
11. Ganzinger, H., Hagen, G., Nieuwenhuis, R., Oliveras, A., Tinelli, C.: DPLL(*T*): Fast Decision Procedures. In: Alur, R., Peled, D.A. (eds.) CAV 2004. LNCS, vol. 3114, pp. 175–188. Springer, Heidelberg (2004)
12. Knudsen, B., Hein, J.: RNA secondary structure prediction using stochastic context-free grammars and evolutionary history. *Bioinformatics* 15, 446–454 (1999)
13. Kroning, D., Ouaknine, J., Seshia, S.A., Strichman, O.: Abstraction-Based Satisfiability Solving of Presburger Arithmetic. In: Alur, R., Peled, D.A. (eds.) CAV 2004. LNCS, vol. 3114, pp. 308–320. Springer, Heidelberg (2004)
14. Lyngsø, R.B., Pedersen, C.N.S.: Pseudoknots in RNA secondary structures. In: Proc. Computational Molecular Biology, RECOMB 2000, pp. 201–209. ACM (2000)
15. Mathews, D., Disney, M., Childs, J., Schroeder, S., Zuker, M., Turner, D.: Incorporating chemical modification constraints into a dynamic programming algorithm for prediction of RNA secondary structure. *Proc. Natl. Acad. Sci.* 101, 7287–7292 (2004)
16. Mathews, D.H., Sabina, J., Zuker, M., Turner, D.H.: Expanded sequence dependence of thermodynamic parameters improves prediction of RNA secondary structure. *Journal of Molecular Biology* 288(5), 911–940 (1999)
17. Mathews, D.H., Turner, D.H.: Prediction of RNA secondary structure by free energy. *Curr. Opin. Struct. Biol.* 16, 270–278 (2006)
18. Ohrimenko, O., Stuckey, P.J., Codish, M.: Propagation = Lazy Clause Generation. In: Bessière, C. (ed.) CP 2007. LNCS, vol. 4741, pp. 544–558. Springer, Heidelberg (2007)
19. Parisien, M., Major, F.: The MC-Fold and MC-Sym pipeline infers RNA structure from sequence data. *Nature* 452, 51–55 (2008)
20. Poolsap, U., Kato, Y., Akutsu, T.: Prediction of RNA secondary structure with pseudoknots using integer programming. *BMC Bioinformatics* 10, S38 (2009)
21. Ren, J., Rastegari, B., Condon, A., Hoos, H.H.: HotKnots: Heuristic prediction of RNA secondary structures including pseudoknots. *RNA* 11, 1494–1504 (2005)
22. Rivas, E., Eddy, S.: A dynamic programming algorithm for RNA structure prediction including pseudoknots. *J. Mol. Biol.* 285, 2053–2068 (1999)
23. Soos, M., Nohl, K., Castelluccia, C.: Extending SAT Solvers to Cryptographic Problems. In: Kullmann, O. (ed.) SAT 2009. LNCS, vol. 5584, pp. 244–257. Springer, Heidelberg (2009)
24. Staple, D.W., Butcher, S.E.: Pseudoknots: RNA structures with diverse functions. *PLoS Biol.* 3(6), e213 (2005)
25. Washietl, S., Hofacker, I., Lukasser, M., Hüttenhofer, A., Stadler, P.: Mapping of conserved RNA secondary structures predicts thousands of functional noncoding RNAs in the human genome. *Nat. Biotechnol.* 23(11), 1383–1390 (2005)
26. Zuker, M., Stiegler, P.: Optimal computer folding of large RNA sequences using thermodynamics and auxiliary information. *Nucleic Acids Research* 9(1), 133–148 (1981)