


DAG-Based Compositional Approaches for LTLf to DFA Conversions

Suguman Bansal 

Georgia Institute of Technology, USA
suguman@gatech.edu

Yash Kankariya

Georgia Institute of Technology, USA
ykankariya3@gatech.edu

Yong Li 

University of Liverpool, UK
yong.li3@liverpool.ac.uk

Abstract—Scalable and efficient conversions of LTL over finite horizon (LTLf) to their deterministic finite automata (DFA) remain a critical bottleneck in several applications of LTLf. Recently, compositional approaches have seen remarkable success in scaling the conversion to large formulas. Here the input formula is first decomposed into smaller subformulas, each of which can be easily converted to their DFAs, then these DFAs are composed to generate the desired DFA. This work proposes a series of simple-yet-effective optimizations to improve the performance of compositional approaches based on reducing the number of composition steps required to generate the desired DFA.

We incorporate these optimizations in a tool called **Lisa2** that builds on one of the state-of-the-art tools for LTLf-to-DFA conversion. A comprehensive empirical evaluation of **Lisa2** demonstrates overall improvements on both parameters of the number of benchmarks solved and runtime. Most remarkably, it demonstrates significant improvement over structured benchmarks where runtime speedups range between 1.5x to 8000x under fair comparisons to prior state-of-the-art tools.

I. INTRODUCTION

Linear Temporal Logic over finite traces [1] (LTLf) is the finite-horizon counterpart of Linear Temporal Logic (LTL) over infinite traces [2]. LTLf is rapidly gaining popularity among real-world applications where behaviors are better expressed over a finite but unbounded horizon. These include applications in planning and synthesis [3], [4], [5], [6], reinforcement learning [7], [8], [9], business processes [10], verification [11].

A critical challenge facing its applications is the conversion of LTLf specifications into their equivalent *deterministic finite automata (DFAs)*. This is not unexpected since the LTLf-to-DFA conversion exhibits a double-exponential blow-up in the size of the input specification in the worst-case [1], [12]. Yet, state-of-the-art LTLf-to-DFA conversion tools like *Lisa* [13] and *Lydia* [14] often succeed at converting large formulas. Their success can be attributed to *compositional algorithms* which are split into two phases. First, in the *decomposition phase* a large LTLf formula is decomposed into smaller subformulas using the formula’s *Abstract Syntax Tree (AST)*. Next, in the *composition phase*, subformulas at the leaves of the AST are converted to their DFAs using direct LTLf-to-DFA tools suitable for smaller formulas, such as *Spot* [15] or *Mona* [16]. Then, these DFAs are composed using language-theoretic and/or automata-based operations to obtain the DFA of the original formula. For DFA composition, the AST of the formula is traversed bottom-up.

Through this work, we propose a series of simple yet effective optimizations to improve the performance of compositional algorithms. Our first optimization aims at striking a balance between the time spent in converting subformulas at the leaves of the AST into their DFAs and the time spent in composing the intermediate DFAs. The deeper the AST is unrolled, the smaller are the subformulas at the leaves of the AST. While these smaller subformulas may be easier to convert into their DFAs, it increases the number of composition steps required to traverse the AST. To this end, we propose to unroll the AST on their outermost boolean operators only. In contrast, both *Lisa* and *Lydia* unroll at the extremes: *Lisa* unrolls on the outermost conjunction only whereas *Lydia* unrolls completely till the propositional literals.

Our other optimizations focus on reducing the number of composition steps required during the bottom-up traversal by modifying the AST. Our first optimization is based on eliminating subformula duplication within the AST. This eliminates multiple computations of the DFA of the same subformula that may be present at multiple locations in the AST of a formula. Such duplication is not uncommon in structured, real-world formulas. Our second optimization is based on using semantics-preserving syntactic transformations to the formula. In particular, subformulas of the form $\phi = \bigwedge_{i=1}^k (\psi \vee \phi_i)$ are rewritten as $\phi = \psi \vee \bigwedge_{i=1}^k \phi_i$, as the latter requires fewer composition steps: The former representation of ϕ will require $2k - 1$ composition steps (k steps to create the DFA for all $(\psi \vee \phi_i)$ and $k - 1$ steps from the outer conjunction of these clauses). Whereas, the latter will require only k steps of which $k - 1$ steps are required to create the DFA for the large conjunction and one more is required to compose with ψ . Neither *Lisa* nor *Lydia* incorporate either of these optimizations.

We have implemented these optimizations in a tool **Lisa2** that builds on the existing tool *Lisa*. The compositional algorithm in **Lisa2** differs from *Lisa* as follows: (a). In the decomposition phase, the formula is unrolled on all outermost boolean operations using the formula’s AST, (b). Next, there is an additional *optimization phase* in which duplicate removal and syntactic transformations modify the AST to a DAG as opposed to the AST, (c). Finally, in the composition phase, formulas at the *leaves* of the DAG are converted to their DFA and then these intermediate DFAs are composed in a

90 bottom-up traversal of the DAG. *Lisa2* also differs from
 91 *Lisa* and *Lydia* in supporting multiple representations for
 92 DFAs. It permits *Spot*'s [15] labeled-graph and Reduced-
 93 Ordered BDD (ROBDD) [17] (which is used by *Lisa*) as well
 94 as *Mona*'s Shared Multi-Terminal BDD (ShMTBDD) (which
 95 is used by *Lydia*) [16]. Permitting both datastructures makes
 96 *Lisa2* more versatile than prior tools. An additional benefit
 97 is that this enables fair comparison with *Lisa* and *Lydia*.
 98 While these prior tools implement differing compositional
 99 approaches, a fair comparison of these algorithms has not
 100 been possible since the performance of these tools is also
 101 affected by the complementary strengths of the underlying
 102 datastructure for DFAs. In particular, ROBDD may be slower
 103 but require less memory whereas ShMTBDD can be blazingly
 104 fast but are memory exhaustive. With the flexibility in choice
 105 of DFA datastructure, *Lisa2* can compare different algorithmic
 106 approaches by ensuring that their underlying datastructures are
 107 identical, hence rendering fairer comparisons.

108 A comprehensive empirical evaluation demonstrates signif-
 109 icant improvements over prior state-of-the-art tools in both
 110 the number of benchmarks solved and their runtime. We
 111 evaluated the performance of *Lisa2* against *Lisa* and *Ly-*
 112 *dia* on LTLf benchmarks (a collection of randomly generated
 113 formulas and structured formulas) from the LTLf track in
 114 SYNTCOMP2023¹. While *Lisa2* outperforms both baselines
 115 comprehensively, its performance on the structured bench-
 116 marks is most remarkable. Not only *Lisa2* solves $\sim 50\%$
 117 more structured benchmarks than prior approaches, it also
 118 demonstrates runtime improvements in the range of 1.5x-
 119 8000x (with more benchmarks recording high runtime im-
 120 provement), highlighting the strength of our tool on realistic
 121 benchmarks.

122 II. PRELIMINARIES AND NOTATIONS

123 A. Linear Temporal Logic over Finite Traces (LTLf)

LTLf [1] extends propositional logic with finite-horizon
 temporal operators. The syntax of LTLf over a finite set of
 propositions *Prop* is identical to LTL, and defined as

$$\varphi := \text{true} \mid \text{false} \mid a \in \text{Prop} \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid X\varphi \mid \varphi_1 U \varphi_2$$

124 where *X* (Next) and *U* (Until), are temporal operators. We
 125 include their dual operators, *N* (Weak Next) and *R* (Release),
 126 defined as $N\varphi \equiv \neg X\neg\varphi$ and $\varphi_1 R \varphi_2 \equiv \neg(\neg\varphi_1 U \neg\varphi_2)$. We
 127 also use typical abbreviations such as $F\varphi \equiv \text{true} U \varphi$, $G\varphi \equiv$
 128 $\text{false} R \varphi$, $\varphi_1 \vee \varphi_2 = \neg(\neg\varphi_1 \wedge \neg\varphi_2)$, $\varphi_1 \rightarrow \varphi_2 \equiv \neg\varphi_1 \vee \varphi_2$.
 129 The semantics of LTLf can be found in [1].

130 Wlog, we assume formulas are given in *negation normal*
 131 *form* (NNF), i.e., the negation operator (\neg) appears in front
 132 of propositions only. In the given syntax, all formulas can be
 133 converted to their NNF with no blow-up in length.

134 Every LTLf formula φ over *Prop* can be converted into
 135 a deterministic finite-state automata (DFA) *D* with alphabet
 136 $\Sigma = 2^{\text{Prop}}$ and at most double-exponential number of states in
 137 $|\varphi|$ such that $\mathcal{L}(D) = \mathcal{L}(\varphi)$ [1].

¹<https://www.syntcomp.org>

B. Abstract Syntax Tree

The *Abstract Syntax Tree* (AST) is an *n*-ary tree represen-
 tation of an LTLf formula. Formally, for an LTLf formula ϕ ,
 (1). Every node corresponds to a subformula of the formula φ .
 In particular, the root of the AST corresponds to the formula
 itself, (2). For every node, the *operator* of the node is given
 by the outermost (primary) operator of its subformula, (3).
 For every node, its children correspond to the immediate
 subformulas of the formula in that node. Formulas with unary
 operators (such as \neg , *X*, *F*, and so on), binary operators (*U*
 and *R*), and *n*-ary operators (\vee and \wedge) consist of one, two,
 and *n*-many children, respectively. The order of children is
 crucial for temporal operators *U* and *R*. For the remaining
 operators, the order of children does not alter the semantics
 of the formula since the operators are commutative.

III. RELATED WORK

Optimizations in compositional LTLf-to-DFA
approaches.: The current state-of-the-art AST-based
 compositional tools, *Lisa* [13] and *Lydia* [14], employ
 various optimizations to counter its inherent non-elementary
 complexity. In addition to aggressive DFA minimization [16]
 and order in which DFAs at each node are composed [13],
 the tools optimize on the depth to which the ASTs are
 unrolled during formula decomposition. For instance,
Lisa decomposes the original formula at its outermost
 conjunction only, thus creating a *k*-ary AST of depth one.
 Thus, the leaves represent subformulas as opposed to atomic
 propositions. On the other hand, *Lydia* [14] generates the
 full *k*-ary AST up to literals in the leaves. This way tools
 attempt to trade-off between resources spent in the direct
 conversion of subformulas at the leaves and composition
 steps. The tools also optimize on the data structure used for
 explicit state-space representations of the DFAs. *Lisa* stores
 DFAs as labeled-graphs and in symbolical state-space using
Reduced-Ordered Binary Decision Diagrams (ROBDD) [17]
 if necessary. *Lydia* stores DFAs using *Shared Multi-Terminal*
Binary Decision Diagrams (ShMTBDD) of *Mona* [16].

Direct LTLf-to-DFA conversions.: *Spot* [15] and
Mona [16] are two popularly used tools for direct LTLf-to-
 DFA conversions of smaller formulas. *Spot* translates the
 LTLf formula into an LTL formula with equivalent semantics,
 converts this formula into a Büchi automaton [18], and then
 transforms this Büchi automaton into a DFA. The *Mona*-
 based approach translates LTLf into *first-order logic over finite*
traces (FOL) and then *Mona* converts the FOL formula into
 a DFA. Both generate minimal DFAs in explicit state-space
 representation.

DAG-based compositional approaches.: *Mona* also uses
 directed acyclic graphs (DAGs) to represent formulas as we
 do in this work [19]. The differences lie in the ways we
 identify equivalent subformulas and the depth of the DAGs.
Mona identifies equivalent formulas ϕ and ϕ' by checking
 whether there is an order-preserving renaming of propositions
 in ϕ such that ϕ and ϕ' are identical, while we check the
 syntax equivalence of two formulas, simpler but much more

138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192

193 efficient. Moreover, *Mona*'s DAG unrolls till its literals while
194 our DAG unrolls on boolean operators only. To the best of
195 our knowledge, *Lisa* and *Lydia* do not use DAG since no
196 intermediate DFAs are stored for later use in their source code.

197 *Optimizations in LTL to automata conversion:* The con-
198 version of LTL (Linear Temporal Logic) [2] to automata
199 forms has received much attention. While the conversion
200 incurs a single-exponential and double-exponential blow-up
201 for the non-deterministic and deterministic versions automata
202 versions, significant work has gone into developing optimiza-
203 tions for LTL to automata. However, these optimizations are
204 too sophisticated for LTLf to automata translations, where
205 relatively simpler translations have been shown to be more
206 effective.

207 IV. OPTIMIZATIONS

208 We propose a series of optimizations to be applied to
209 compositional approaches for LTLf-to-DFA conversion. The
210 first optimization (Section IV-A) determines the depth to
211 which an input formula's AST is unrolled during formula
212 decomposition into smaller subformulas. The remaining two
213 optimizations are geared to reduce the computation required
214 during the composition phase by reducing the number of com-
215 position operations. Here, the first optimization compresses
216 this AST by removing duplicate subformulas (Section IV-
217 B). The second optimization performs semantics-preserving
218 syntactic transformations that are guaranteed to reduce the
219 number of composition steps (Section IV-C).

220 A. Depth of AST Unrolling

221 Our first optimization is based on the depth to which
222 an input LTLf formula is unrolled to obtain smaller sub-
223 formulas. We propose to unroll subformulas only if their
224 outermost operator is a boolean operator. Recall, since we
225 assume formulas are given in NNF, the outermost boolean
226 operators are effectively only the conjunction operator or the
227 disjunction operator as negations appear before propositions
228 only. Consequently, for formulas at the leaves of this tree, the
229 outermost operator could be any of the temporal operators.
230 Figure 1 illustrates such an unrolling of an AST.

231 We make this choice to strike a balance between the
232 conversion of subformulas to DFAs at the leaves vs. the
233 composition of DFAs at intermediate nodes to obtain the final
234 DFA. Prior state-of-the-art approaches *Lisa* and *Lydia* take
235 diametrically opposite routes in this regard. *Lisa* unrolls the
236 AST at its outermost conjunction only. I.e., given an LTLf
237 formula $\varphi = \bigwedge_{i=1}^n \varphi_i$, *Lisa* decomposes the original formula
238 into the n -subformulas given by φ_i s. The disadvantage of
239 this decomposition is that in the worst case, the φ_i s could
240 be too large to be handled by an off-the-shelf LTLf-to-DFA
241 conversion tools like *Spot* or *Mona*. The advantage, however,
242 is that once the DFAs for the φ_i s are created, the approach
243 requires only $n-1$ composition steps, where each composition
244 consists of polynomial-time operations of DFA product and
245 DFA minimization only. In contrast, *Lydia* unrolls the AST
246 completely. I.e., its leaves comprise of literals (propositions

247 or their negation). This ensures that the DFA at the leaf node
248 is obtained trivially. However, not only do the number of
249 composition operations increase dramatically, the composition
250 operations may become more complex. In particular, the
251 composition at nodes with a boolean operator comprise of
252 polynomial-time operations identical to *Lisa*, but the composi-
253 tion at nodes with temporal operators may involve exponential-
254 time operations such as projection and/or determinization.

255 Our choice to unroll only on boolean operators ensures that
256 all composition operations require polytime operations only
257 while also ensuring that the size of subformulas at the leaves
258 are small, hence combining the benefits of *Lisa* and *Lydia*.

259 B. Duplicate Removal

260 For our next optimization, we observe that in several
261 formulas, an intermediate subformula may appear more than
262 once in the formula's AST. This results in redundant com-
263 putation during the composition phase as it generates the
264 DFA for equivalent subformulas multiple times. To eliminate
265 such redundant recomputation, we propose to merge nodes of
266 equivalent subformulas. This is illustrated in Figure 2 where
267 formula θ_1 that appeared twice in Figure 1 has been merged
268 into one node. Such duplication removal will result in the AST
269 being converted to a DAG as the merged nodes are required
270 to serve multiple parent nodes.

271 In order to merge nodes in an AST, we must check if the
272 formulas at two or more nodes are equivalent. LTLf formula
273 equivalence is PSPACE-complete, hence we must resort to
274 computationally inexpensive approaches to identify formula
275 equivalence. Tools such as *Spot* offer inexpensive syntactic
276 checks to examine formula equivalence. We combine these
277 checks with leveraging the parent-child relationship between
278 nodes in the AST to identify equivalent formulas.

279 To elaborate further, first we identify formula equivalence
280 between the leaf nodes of the AST using syntactic checks,
281 and merge each class of equivalent formulas into one leaf
282 node. This converts the AST into a DAG as the merged
283 leaf nodes will now serve multiple parent nodes. Next, it
284 is easy to see that two non-leaf nodes are equivalent if
285 all their children nodes are identical and their operators are
286 identical. All such formula equivalence in non-leaf nodes can
287 be identified efficiently in a single bottom-up traversal of the
288 DAG that simultaneously merges equivalent non-leaf nodes
289 into one node.

290 Note that this procedure may fail to recognize equivalent
291 subformulas that do not adhere to our inexpensive checks.
292 Despite this incompleteness, we observe that sometimes it can
293 reduce the number of nodes in the AST/DAG by 30-40% in
294 negligible time, hence demonstrating its effectiveness.

295 C. Semantics-Preserving Transformation

296 The final optimization aims to reduce the number of com-
297 position steps using *semantics-preserving syntactic transfor-*
298 *mation* of the formula. Lemma 1 motivates our optimization:




Fig. 1: AST unrolled on boolean operators only.




Fig. 2: After duplicate removal. θ_1 has been merged.




Fig. 3: After Transformation. θ_1 has been pulled out.

299 *Lemma 1:* Consider the following formula,

$$\phi = \circ_{i=1}^k \left((\theta_1 \circ \dots \circ \theta_i) \circ (\theta_{i,1} \circ \dots \circ \theta_{i,m_i}) \right) \quad (1)$$

300 where $\circ, \circ' \in \{\vee, \wedge\}$ s.t. $\circ' \neq \circ$, and for all $i \in [k]$, $m_i \geq 0$.

301 Then ϕ is equivalently written as:

$$\phi' = (\theta_1 \circ \dots \circ \theta_l) \circ \left(\circ_{i=1}^k (\theta_{i,1} \circ \dots \circ \theta_{i,m_i}) \right) \quad (2)$$

302 using the laws of associativity. Assuming the DFAs for all θ_i
 303 and $\theta_{i,j}$ are given, the required composition steps to create the
 304 DFA for ϕ is $\mathcal{O}(l \cdot (k-1))$ more than those required to create
 305 the DFA for ϕ' .

306 *Proof:* We begin with a sketch of the argument. In
 307 practice, a product over k DFAs requires $k-1$ products
 308 between two DFAs. Now, notice that the intermediate DFA
 309 for the common segment $\theta_c = (\theta_1 \circ \dots \circ \theta_l)$ is constructed
 310 $k-1$ times more in ϕ than in ϕ' . By pulling out the common
 311 segment θ_c using the laws of associativity in ϕ' , the DFA for
 312 θ_c is constructed only once, amounting to the difference.

313 The formal argument follows: Let us first analyze the
 314 number of products required in ϕ . For every $i \in [k]$, the clause
 315 $(\theta_1 \circ \dots \circ \theta_l) \circ (\theta_{i,1} \circ \dots \circ \theta_{i,m_i})$ requires $l + m_i - 1$ products.
 316 Next, these clauses are combined using products to obtain ϕ .
 317 Since there are k clauses, we require $k-1$ additional products.
 318 Therefore, evaluating ϕ requires $(k-1) + \sum_{i=1}^k (l-1 + m_i) =$
 319 $k \cdot l - 1 + \sum_{i=1}^k m_i$ product operations.

320 Next, we analyse the number of products required in ϕ' .
 321 For every $i \in [k]$, the clause $(\theta_{i,1} \circ \dots \circ \theta_{i,m_i})$ requires $m_i - 1$
 322 products. In addition, these k clauses are combined using $k-1$
 323 products to form the DFA for $\left(\circ_{i=1}^k (\theta_{i,1} \circ \dots \circ \theta_{i,m_i}) \right)$. Hence,
 324 $\left(\circ_{i=1}^k (\theta_{i,1} \circ \dots \circ \theta_{i,m_i}) \right)$ requires $k-1 + \sum_{i=1}^k (m_i - 1) =$
 325 $(\sum_{i=1}^k m_i) - 1$ operations. Combined with l many θ_j s to obtain
 326 ϕ' , we require $l-1 + \sum_{i=1}^k m_i$ products to form ϕ' . Therefore,
 327 constructing the DFA via ϕ' requires $\mathcal{O}(l \cdot (k-1))$ fewer
 328 operations than creating the same DFA via ϕ , where l is the
 329 number of subformulas common to k -many clauses in ϕ . ■

330 Our optimization, therefore, applies the associative law to
 331 transform nodes of the form ϕ to nodes of the form ϕ' in
 332 the DAG obtained after duplicate removal. As earlier, the

333 transformation is carried out by an analysis of the parent-
 334 child relations between nodes. A node v_0 is eligible for the
 335 transformation if the formula it represents is of the form ϕ ,
 336 i.e. : (a) the outermost boolean operator should differ from
 337 the outermost boolean operator of all of its children, and (b)
 338 all its children share a common child of their own, referring
 339 to $\theta_c = (\theta_1 \circ \dots \circ \theta_l)$ in ϕ . In the DAG, the common child
 340 of all children is simply a common grandchild node. The
 341 common grandchildren are obtained from the intersection of
 342 all of children of v_i 's for $i \geq 1$. In Figure 2, node v_0 is eligible
 343 for the transformation with a single common grandchild in
 344 θ . When eligible, the transformation *pulls out* all common
 345 segment θ_c . In the DAG, this translates to promoting all
 346 common grandchildren of v_0 to direct children of v_0 and
 347 the earlier children of v_0 are modified accordingly. Figure 2
 348 to Figure 3 illustrate the transformation. Observe that the
 349 transformation may result in the creation of new nodes such
 350 as v'_0, \dots, v'_k in Figure 3.

351 As earlier, these transformations are carried out in a single
 352 bottom-up traversal (reverse topological order) of the DAG
 353 starting with the leaf nodes. In instances when the transforma-
 354 tion results in the creation of a new node (such as v'_0, v'_1, \dots, v'_k
 355 in Figure 3), the new nodes are examined for duplicates using
 356 the earlier approach. Then the transformation is recursively
 357 applied to v'_0 first and then to v'_1, \dots, v'_k before returning to
 358 the next node as per the reverse topological order.

359 V. COMPOSITIONAL ALGORITHM

360 For the sake of completion, we present an outline of the
 361 compositional algorithm. W.l.o.g., our algorithm receives an
 362 LTLf formula in NNF and outputs a minimal DFA for the
 363 formula. The algorithm proceeds in three phases: First is
 364 the *decomposition phase* in which the input LTLf formula is
 365 decomposed into smaller subformulas based on its AST. The
 366 AST is unrolled on boolean operators only. This is followed
 367 by the *optimization phase* in which the proposed duplication
 368 removal and semantic-transformations are applied. As a result,
 369 the AST is converted to a DAG. Finally, in the *composition*
 370 *phase*, the subformulas at the *leaves* of the DAG, i.e. nodes
 371 with no outgoing edges in the DAG, are converted to their

372 minimal DFA form. Next, the DAG is traversed bottom-up
 373 starting with the leaves, i.e. the DAG is traversed in reverse
 374 topological order. During this traversal, the minimal DFA at a
 375 node is created if the minimal DFA at all its children have
 376 already been constructed. The primary difference from the
 377 AST-based composition is that the DFA at a node in the AST
 378 can be removed from memory as soon as the DFA in its parent
 379 node has been constructed. In the case of a DAG, the DFA at
 380 a node is discarded only after the DFA at *all* its parent nodes
 381 have been generated.

382 VI. IMPLEMENTATION DETAILS

383 We have implemented compositional algorithm in a tool
 384 called *Lisa2*². *Lisa2* takes an LTLf formula in NNF as its
 385 input and outputs its minimal DFA in explicit representation.

386 In brief, *Lisa2* extends a current state-of-the-art tool *Lisa* to
 387 incorporate the optimizations described in Section IV. In de-
 388 tail, *Lisa2* has been written in C++. It uses *Spot* LTLf parser
 389 to parse the input formula. The input formula is decomposed
 390 into a DAG following the optimizations described in Sec-
 391 tion IV. To convert the subformulas at the leaves of the DAG,
 392 *Lisa2* converts the LTLf formulas at the leaf nodes to their
 393 equivalent first-order logic (FOL) and uses *Mona* to convert
 394 the FOL formulas to their minimal DFAs. The DFAs are then
 395 composed as described in Section V. Similar to *Lisa* [13],
 396 *Lisa2* deploys two performance-enhancing heuristics (a)
 397 *aggressive DFA minimization*, i.e. each DFA (intermediate or
 398 final) is minimized as soon as it is created, and (b) *smallest-*
 399 *first* heuristic that always picks the smallest two (minimal)
 400 DFAs to compose during a k -way product construction (for
 401 both conjunction and disjunction).

402 *Lisa2* generates DFA in explicit-state representation, i.e.,
 403 the states are given explicitly and the transitions are given
 404 as labeled formulas over the propositions of the input LTLf
 405 formula. *Lisa2* supports two datastructures to represent the
 406 final and all intermediate DFAs: (a) *Spot*'s labeled graphs and
 407 Reduced Ordered BDD (ROBDD) and (b). *Mona*'s Shared
 408 Multi-Terminal BDD (ShMTBDD). We refer to these two vari-
 409 ants of our tool as *Lisa2-Spot* and *Lisa2-Mona*, respectively.
 410 These tool variants use the DFA manipulation APIs provided
 411 by *Spot* and *Mona*, respectively, for all DFA operations
 412 including the product construction and minimization.

413 *Tool Features:* By supporting both *Spot* and *Mona*,
 414 *Lisa2* is the only LTLf-to-DFA conversion tool that can
 415 support both datastructures, adding to its versatility in applica-
 416 tions. Another motivation to support both DFA datastructures
 417 is to enable fair comparison for future algorithmic advances
 418 in LTLf-to-DFA conversion tools. Prior tools *Lisa* and *Ly-*
 419 *dia* support only one of the two *Spot*'s labeled graphs
 420 + ROBDD combination and ShMTBDD, respectively. These
 421 datastructures have complementary benefits (ROBDD may be
 422 slower but require less memory whereas ShMTBDD are faster
 423 but are memory extensive.) and a bear significant impact
 424 the performance of their tools. As a result, performance

425 comparisons between prior tools are unable to differentiate
 426 between the improvement caused by the algorithm vs. the
 427 improvement caused by the datastructure. Thus the ability to
 428 switch between DFA datastructures within *Lisa2* creates the
 429 opportunity for more fair comparisons of algorithmic advances
 430 in LTLf-to-DFA tools.

431 A. Implementation-Level Optimizations

432 *Lisa2* incorporates several implementation-level optimiza-
 433 tions. Few key ones are described below.

434 First, formulas of the form $G(\bigwedge_{i=0}^m \phi_i)$ and $F(\bigvee_{i=0}^m \phi_i)$ are
 435 equivalently written as $\bigwedge_{i=0}^m (G\phi_i)$ and $\bigvee_{i=0}^m (F\phi_i)$, respec-
 436 tively, to promote deeper decomposition on boolean operators.
 437 Had the formulas been retained in their earlier format, then
 438 the formulas would not be decomposed any further since the
 439 outermost operator is temporal. This optimization generates
 440 smaller subformulas.

441 Second, *Lisa2* already identifies few subformula duplica-
 442 tions (using *Spot*'s inexpensive methods to determine formula
 443 equivalence) during the unrolling of the formula's AST. As a
 444 result, the outcome of the unrolling may already be a DAG as
 445 opposed to an AST. We do this as we observed that in some
 446 cases, the AST obtained from unrolling on boolean operators
 447 could become very large. Combining the unrolling with a
 448 shallow duplication-removal curb the growth in the AST.

449 We observe that in practice most DAG/AST nodes do not
 450 possess a common grandchild, making the node ineligible for
 451 the semantics-preserving transformation. Instead, it is more
 452 likely that several nodes possess a *popular* grandchild that may
 453 be a child of most but not all children of the node. In these
 454 cases, we perform the transformation only with the children
 455 that share the popular grandchild.

456 VII. EXPERIMENTAL ANALYSIS

457 A. Design and Setup for Empirical Evaluation

458 *Baselines and Fair Comparisons:* We compare *Lisa2* to
 459 the three state-of-the-art baselines: *Lydia*, *Lisa*, and *Lisa-*
 460 *Explicit*. All three tools are based on compositional algorithms.
 461 They differ in the depth of unrolling, few algorithmic details,
 462 and the underlying DFA datastructure. *Lydia* unrolls to
 463 the literals whereas *Lisa* and *Lisa-Explicit* unroll on the
 464 outermost conjunction only. In terms of DFA data structure,
 465 *Lydia* uses *Mona*'s ShMTBDD while *Lisa* and *Lisa-*
 466 *Explicit* use *Spot*'s labeled-graphs and ROBDDs. Since tool
 467 performance is known to be impacted by the DFA datastruc-
 468 ture, we establish the following fair comparisons:

- 469 • *Lisa2-Mona* vs. *Lydia*
- 470 • *Lisa2-Spot* vs. *Lisa* and *Lisa-Explicit*

471 All tools accept inputs in *Spot*-parsable format, ensuring
 472 consistency among tools in input format.

²<https://github.com/suguman-lab/lisa2>

473 *Benchmarks:* We use benchmarks from the LTLf-track at
 474 SYNTCOMP 2023³. We evaluate on 490 benchmarks of which
 475 400 formulas are generated randomly and the remaining 90
 476 are structured formulas derived from the "two-player games"
 477 category. Among the structured benchmarks, we use 20, 10,
 478 and 60 benchmarks from the single counter, double counter,
 479 and nim benchmark classes, respectively.

480 *Set-up:* All experiments were conducted on a single node
 481 of a high-performance cluster (<https://pace.gatech.edu/>). Each
 482 node consists of four quad-core Intel-Xeon processors running
 483 at 2.6 GHz with 4hrs timeout and 16GB of RAM each.

484 B. Performance-Related Observations

485 We begin by examining the performance of Lisa2 against its
 486 counterparts w.r.t. runtime and number of benchmarks solved.
 487 Overall, Lisa2 not only solves more benchmarks than all other
 488 counterparts, it also improves the runtime significantly. Most
 489 remarkable is its performance on the structured benchmarks
 490 where Lisa2 solves $\sim 50\%$ more benchmarks and displays
 491 runtime improvements in the range of $2x-8000x$. We describe
 492 our observations and inferences in detail below.

493 *Lisa2 demonstrates the best overall performance:* The
 494 cactus plots of the performance of all tools in Figure 4a
 495 (cactus plot on all benchmarks) and Figure 4b (cactus plot
 496 on structured benchmarks only) demonstrate that variants of
 497 Lisa2 solve the most number of benchmarks in both cases.
 498 Recall the fair comparisons from the previous section. We
 499 observe that on all benchmarks, Lisa2 Mona outperforms
 500 Lydia and Lisa2-Spot is comparable to/better than its fair
 501 counterparts.


502 Lisa2 comprehensively outperforms its fair counterparts
 503 on structured benchmarks. Lisa2-Spot solves almost twice
 504 as many benchmarks that its fair counterparts while Lisa2-
 505 Mona solves $\sim 40\%$ more benchmarks than Lydia. This high-
 506 lights the benefits of our optimizations on realistic bench-
 507 marks. More broadly, it reflects the merits of identifying and
 508 leveraging patterns appearing in structured (realistic) formulas.

509 Next, we examine each structured benchmark class in detail.

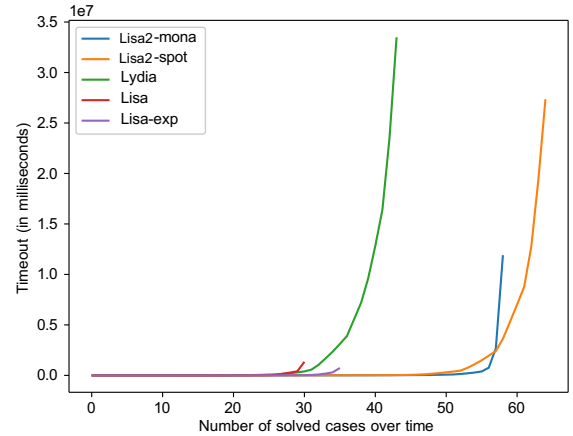
510 *Lisa2 performs remarkably on the nim benchmarks.:*
 511 Both variants of Lisa2 outperform their fair counterparts by
 512 a large margin in both, the number of benchmarks solved
 513 and runtime. Lisa2 solves almost twice as many benchmarks
 514 as their counterparts (Figure 5). Furthermore, the runtime
 515 speedup ranges between $2x-8000x$ with most benchmarks
 516 displaying greater than $500x$ speedup on Mona's ShMTBDD
 517 data structure; and on average $5x$ speedup on Spot's labeled-
 518 graph and ROBDD data structure (Table I).

519 This outcome is impressive as the nim-benchmarks had
 520 proven to be challenging for prior compositional approaches.
 521 This occurs since on these benchmarks the intermediate (min-
 522 imal) DFAs tend to be very large even though final (mininal)

³SYNTCOMP: <https://www.syntcomp.org>. Benchmarks were taken from
<https://github.com/whitemech/finite-synthesis-datasets/tree/main>. We chose
 benchmarks from the whitemech repository because (a). All SYNTCOMP23
 benchmarks in LTLf track were obtained by converting the whitemech
 benchmarks to TLSF format, (b). All baseline tools natively support the format
 used in whitemech as opposed to the TLSF format used by SYNTCOMP.



(a) Cactus plot: All benchmarks. Timeout 4hrs



(b) Cactus plot: Structured benchmarks. Timeout 4hrs

Fig. 4: Overall Performance

DFA is quite small. Hence, it is not uncommon for Lisa or
 Lydia to fail at an intermediate stage due to memory or time
 shortage. We attribute Lisa2's success to our optimizations
 in reducing the number of compositional steps. For most
 benchmarks in the nim-class, after duplication removal and
 semantic transformations, the resulting DAG comprised of 5-
 20% fewer compositions steps than the AST obtaining from
 unrolling on boolean operators only. In another class of nim-
 benchmarks derived from [20], this reduction even ranged
 between 20-50%, resulting in better performance gain.

These experiments clearly demonstrate the advantage of
 reducing the composition steps.

Performance on counter benchmarks: On the single- and
 double-counter classes of benchmarks, Lisa2-Mona demon-
 strates $1.4x-100x$ runtime improvement over Lydia. Whereas,
 Lisa2-Spot displays $1.5x-100x$ runtime improvement over
 Lisa but is sometimes slower than Lisa-Explicit, as demon-
 strated in Table II.

We attribute the performance of Lisa2 on the counter

523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541




Fig. 5: Number of benchmarks solved on the nim class. x - and y -axes denote benchmark class (total instances in brackets) and num. of benchmarks solved, respectively.

Cases	Lisa2-Mona	Lydia	Lisa2-Spot	Lisa	Lisa-exp
nim_01.1	10	21	10	0	0
nim_01.2	30	44	20	0	0
nim_01.3	30	89	40	10	10
nim_01.4	50	277	50	20	10
nim_01.5	50	1087	80	50	40
nim_01.6	80	4413	100	120	100
nim_01.7	90	16929	110	210	220
nim_01.8	120	64115	160	-	-
nim_01.9	190	181994	200	-	-
nim_01.10	230	780581	240	-	-
nim_01.11	380	1652381	330	-	-
nim_01.12	420	3203658	360	-	-
nim_01.13	620	-	560	-	-
nim_01.14	820	-	780	-	-
nim_01.15	2220	-	910	-	-
nim_01.16	9780	-	810	-	-
nim_01.17	49100	-	810	-	-
nim_01.18	-	-	1300	-	-
nim_01.19	108890	-	1490	-	-
nim_01.20	-	-	1820	-	-
nim_02.1	20	69	30	20	10
nim_02.2	40	483	80	60	80
nim_02.3	80	5979	180	380	400
nim_02.4	130	92548	380	2040	1760
nim_02.5	200	650682	820	-	-
nim_02.6	350	3574672	2430	-	-
nim_02.7	480	-	8890	-	-
nim_02.8	1150	-	76890	-	-
nim_02.9	2540	-	348840	-	-
nim_02.10	3950	-	466940	-	-
nim_02.11	-	-	-	-	-
nim_02.12	17690	-	452700	-	-
nim_02.13	-	-	-	-	-
nim_02.14	-	-	-	-	-
nim_02.15	-	-	3970670	-	-
nim_03.1	60	465	110	160	160
nim_03.2	160	52220	810	3050	3260
nim_03.3	1110	1653251	27080	-	-
nim_03.4	2110	-	91610	-	-
nim_03.5	4000	-	396780	-	-
nim_03.6	7270	-	1215350	-	-
nim_04.1	180	23813	1050	3920	4520
nim_04.2	2540	7083597	75430	-	-
nim_04.3	-	-	1670340	-	-
nim_04.4	-	-	6634310	-	-
nim_05.1	1310	755274	23500	112100	123220
nim_05.2	-	-	1718010	-	-

TABLE I: Runtime in milliseconds for nim. Timeout 4hrs

Cases	Lisa2-Mona	Lydia	Lisa2-Spot	Lisa	Lisa-exp
counter_1	10	7	10	10	10
counter_2	10	17	40	60	60
counter_3	10	33	310	560	540
counter_4	10	57	20	10	10
counter_5	20	82	20	30	10
counter_6	50	156	50	60	30
counter_7	200	371	170	300	100
counter_8	810	1111	690	9560	290
counter_9	3520	4212	2870	125350	990
counter_10	15190	16611	14770	113410	3700
counter_11	66390	76345	61690	-	15130
counter_12	366250	474631	277750	-	90000
counter_13	1910580	2419211	1762580	-	430210
counter_14	9241030	10013917	7970190	-	-
counter_15	-	-	-	-	550
counters_1	10	17	50	80	70
counters_2	10	55	10	10	10
counters_3	20	150	30	80	30
counters_4	80	1103	130	1220	170
counters_5	600	25784	1000	29540	1060
counters_6	7850	660391	8090	941940	7380
counters_7	74590	-	60060	-	47080

TABLE II: Runtime in milliseconds for counters. Timeout 4hrs.

benchmarks to the unrolling depth. This is because for most of these benchmarks, duplication removal and semantic transformation did not result in any significant reduction in composition steps as the benchmarks exhibit neither multiple occurrences of a subformula nor are their patterns amenable to the syntactic transformation. We observed that Lydia would fail because of the accumulation in number and size of intermediate DFAs in its AST that unrolls till the literals. This is aggravated by the ShMTDD datastructure to represent DFAs as they can become memory extensive. On the other hand, on these benchmarks, Lisa and Lisa-Explicit benefit from the shallowest unrolling. Lisa2-Spot unrolls the formula deeper than Lisa and Lisa-Explicit, resulting in many more composition steps. The runtime of Lisa2-Spot compared to Lisa-Explicit is further affected as the underlying datastructure of Spot's labeled graphs and ROBDDs are known to result in slower compositions.

A closer examination of this class of benchmark revealed a potential avenue for improvement. While the formulas did not have duplicates, they had several *symmetric subformulas* upto propositional isomorphism. Further optimizations based on leveraging such similarities within such formulas could further improve the performance of LTLf-to-DFA conversion tools, including ours.

Lisa2-Spot vs. Lisa2-Mona.: We compare the performance of Lisa2-Spot and Lisa2-Mona. Note that here the underlying algorithm is identical. The only difference between the two is the choice of datastructure for DFA representations. As a result, we expect this experiment to highlight the impact of datastructure on a tool's performance.

Our observations confirm that the datastructure plays a vital role in a tool's performance, as we observe that the tools Lisa2-Mona and Lisa2-Spot display the same differences displayed by the underlying datastructure, i.e. the observed trend is that Lisa2-Mona consumes more memory but is faster while Lisa2-Spot may be slower but it consumes lesser

Cases	Lisa2-Spot	Lisa2-Spot (Dup. Rem. only)	Lisa2-Mona	Lisa2-Mona (Dup. Rem. only)
nim_01_01	10	20	10	10
nim_01_02	20	20	30	20
nim_01_03	40	40	30	40
nim_01_04	50	60	50	50
nim_01_05	80	70	50	70
nim_01_06	100	100	80	90
nim_01_07	110	120	90	110
nim_01_08	160	150	120	140
nim_01_09	200	170	190	230
nim_01_10	240	210	230	210
nim_01_11	330	290	380	280
nim_01_12	360	350	420	440
nim_01_13	560	500	620	640
nim_01_14	780	570	820	870
nim_01_15	910	690	2220	1750
nim_01_16	810	810	9780	4270
nim_01_17	810	780	49100	50760
nim_01_18	1300	1170	-	-
nim_01_19	1490	1400	108890	81370
nim_01_20	1820	1700	-	-
nim_02_01	30	40	20	30
nim_02_02	80	80	40	40
nim_02_03	180	180	80	90
nim_02_04	380	400	130	140
nim_02_05	820	770	200	220
nim_02_06	2430	2240	350	330
nim_02_07	8890	4870	480	510
nim_02_08	76890	66990	1150	1150
nim_02_09	348840	276800	2540	2380
nim_02_10	466940	4576480	3950	5000
nim_02_11	-	2953960	-	6930
nim_02_12	452700	299940	17690	18400
nim_02_13	-	13147500	-	-
nim_02_14	-	-	-	-
nim_02_15	3970670	4717260	-	-
nim_03_01	110	100	60	60
nim_03_02	810	700	160	170
nim_03_03	27080	8580	1110	640
nim_03_04	91610	173470	2110	1800
nim_03_05	396780	314470	4000	5810
nim_03_06	1215350	9473690	7270	-
nim_04_01	1050	890	180	210
nim_04_02	75430	104110	2540	2120
nim_04_03	1670340	1093050	-	-
nim_04_04	6634310	14091500	-	-
nim_05_01	23500	17230	1310	1070
nim_05_02	1718010	5659390	-	-

TABLE III: Ablation Study: Runtime in millisecs for nim benchmarks on Lisa2 and its version with the duplicate removal optimization (i.e. no semantic transformation) only. The **lower runtime** is in bold. Timeout 4hrs.

memory, hence is capable to solve more benchmarks.

These observations further highlight the need for fair comparisons in LTLf-to-DFA conversions that we raised earlier, hence reflects on the importance of tools supporting both datastructures for DFA representation.

C. Ablation Study

Finally, we perform an ablation study to examine the effect of each optimization individually. Together the optimizations of duplicate removal and syntactic transformation reduce the number of composition operations. We are interested in studying their effects individually. For this, we compare the

performance of Lisa2 (under each DFA datastructure choice) against its own version in which the syntactic transformation has been disabled, i.e., they only applied duplicate removal.

Figure 5 demonstrates the performance of Lisa2-Spot and Lisa2-Mona against their variants that perform duplicate removal only. Apriori, one would imagine that compounding reduction in composition steps through duplicate removal and syntactic transformation would result in improved performance (both in number of benchmarks solved and runtime). However, Figure 5 demonstrates that in some cases (nim_02) the variant that only performed duplicate removal solved more benchmarks. This surprising result led us to further examine the runtime of these tools, shown in Table III. This reveals that there are a significant number of cases where only performing duplicate removal performed better than compounding both optimizations and there are equally many cases where compounding optimizations displayed the stronger runtime performance. In either case, the overall runtime performance is still an improvement over prior state-of-the-art tools.

In order to understand this behavior, we first ascertained that each optimization consumes such a negligible amount of time that it cannot be considered to be the reason behind runtime decline. Similarly, we ascertained that each optimization contributed to reducing the number of composition steps.

We conclude that the unpredictability, despite the reduction in number of composition steps, arises due to the creation of new nodes (new subformulas) after syntactic transformation. To elaborate further, syntactic transformations may result in creating subformulas that were not originally present in the input formula. It is possible that the new formulas are such that even though their DFA construction may require fewer composition steps, these steps may be more expensive as an intermediate DFA may be difficult to create. This results in unpredictability in performance when both optimizations are compounded. In contrast, duplicate removal never creates any new node (new subformula). It only reduces the number of times some nodes may be computed. Hence, duplicate removal will always reduce the overall runtime.

VIII. CONCLUDING REMARKS

This work presents Lisa2 which incorporates a series of simple-yet-effective optimizations for compositional approaches for LTLf-to-DFA conversion. Empirical evaluations of this tool displays significant performance improvement, especially on structured benchmarks derived from real-world scenarios: Lisa2 solves $\sim 50\%$ more benchmarks and shows runtime improvement in the range of 1.5x-8000x.

Our optimizations are based on reducing the number of composition steps required to construct the desired DFA. Despite the remarkable performance of Lisa2, our experiments reveal that simply reducing the number of composition steps may not be sufficient, especially if the reduction is accompanied with the creation of new subformulas for which DFA construction may be hard to generate.

We also emphasize on the need for fair comparison to compare algorithmic advances. This is crucial for LTLf-to-

644 DFA conversion as the choice of datastructure for DFAs have
645 a significant impact on a tool's performance.

646 *Acknowledgements:* We thank Marco Favorito and Kuldeep
647 Meel for their help in setting up baseline tools. We thank
648 the anonymous reviewers for their valuable feedback. This
649 work has been supported by the EPSRC through grant
650 EP/X021513/1 and Georgia Institute of Technology's Presi-
651 dential Undergraduate Research Award for Fall 2023.

652 REFERENCES

- 653 [1] G. De Giacomo and M. Y. Vardi, "Linear temporal logic and linear
654 dynamic logic on finite traces," in *IJCAI*. AAAI Press, 2013, pp. 854–
655 860.
- 656 [2] A. Pnueli, "The temporal logic of programs," in *FOCS*. IEEE, 1977,
657 pp. 46–57.
- 658 [3] A. Camacho, E. Triantafillou, C. Muise, J. Baier, and S. McIlraith, "Non-
659 deterministic planning with temporally extended goals: Ltl over finite
660 and infinite traces," in *Proceedings of the AAAI conference on artificial
661 intelligence*, vol. 31, no. 1, 2017.
- 662 [4] G. De Giacomo, F. M. Maggi, A. Marrella, and F. Patrizi, "On the
663 disruptive effectiveness of automated planning for ltlf-based trace align-
664 ment," in *Proceedings of the AAAI Conference on Artificial Intelligence*,
665 vol. 31, no. 1, 2017.
- 666 [5] J. A. Baier and S. McIlraith, "Planning with temporally extended goals
667 using heuristic search," in *ICAPS*. AAAI Press, 2006, pp. 342–345.
- 668 [6] M. Lahijanian, S. Almagor, D. Fried, L. E. Kavrakı, and M. Y. Vardi,
669 "This time the robot settles for a cost: A quantitative approach to
670 temporal logic planning with partial satisfaction." in *AAAI*. AAAI
671 Press, 2015, pp. 3664–3671.
- 672 [7] R. Brafman, G. De Giacomo, and F. Patrizi, "LTLf/LDLf non-markovian
673 rewards," in *AAAI*, vol. 32, no. 1, 2018.
- 674 [8] A. Camacho, R. T. Icarte, T. Q. Klassen, R. A. Valenzano, and S. A.
675 McIlraith, "LTL and beyond: Formal languages for reward function
676 specification in reinforcement learning." in *IJCAI*, vol. 19, 2019, pp.
677 6065–6073.
- 678 [9] K. Jothimurugan, S. Bansal, O. Bastani, and R. Alur, "Compositional
679 reinforcement learning from logical specifications," *Advances in Neural
680 Information Processing Systems*, vol. 34, pp. 10 026–10 039, 2021.
- 681 [10] M. Pestic, D. Bosnacki, and W. M. P. van der Aalst, "Enacting declarative
682 languages using LTL: avoiding errors and improving performance," in
683 *SPIN*. Springer, 2010, pp. 146–161.
- 684 [11] S. Bansal, Y. Li, L. M. Tabajara, M. Y. Vardi, and A. Wells, "Model
685 checking strategies from synthesis over finite traces," in *International
686 Symposium on Automated Technology for Verification and Analysis*.
687 Springer, 2023, pp. 227–247.
- 688 [12] O. Kupferman and M. Y. Vardi, "Model checking of safety properties,"
689 in *CAV*. Springer, 1999, pp. 172–183.
- 690 [13] S. Bansal, Y. Li, L. Tabajara, and M. Vardi, "Hybrid compositional
691 reasoning for reactive synthesis from finite-horizon specifications," in
692 *AAAI*, vol. 34, no. 06, 2020, pp. 9766–9774.
- 693 [14] G. De Giacomo and M. Favorito, "Compositional approach to translate
694 LTLf/LDLf into deterministic finite automata," in *Proceedings of the In-
695 ternational Conference on Automated Planning and Scheduling*, vol. 31,
696 2021, pp. 122–130.
- 697 [15] A. Duret-Lutz, E. Renault, M. Colange, F. Renkin, A. G. Aisse,
698 P. Schlehuber-Caissier, T. Medioni, A. Martin, J. Dubois, C. Gillard,
699 and H. Lauko, "From spot 2.0 to spot 2.10: What's new?" in *Computer
700 Aided Verification - 34th International Conference, CAV 2022, Haifa,
701 Israel, August 7-10, 2022, Proceedings, Part II*, ser. Lecture Notes in
702 Computer Science, S. Shoham and Y. VizeI, Eds., vol. 13372. Springer,
703 2022, pp. 174–187.
- 704 [16] J. G. Henriksen, J. Jensen, M. Jørgensen, N. Klarlund, R. Paige,
705 T. Rauhe, and A. Sandholm, "Mona: Monadic second-order logic in
706 practice," in *TACAS*. Springer, 1995, pp. 89–110.
- 707 [17] R. E. Bryant, "Graph-based algorithms for boolean function manipula-
708 tion," *IEEE Transactions on Computers*, vol. 100, no. 8, pp. 677–691,
709 1986.
- 710 [18] R. Gerth, D. Peled, M. Y. Vardi, and P. Wolper, "Simple on-the-fly
711 automatic verification of linear temporal logic," in *PSTV*. Springer,
712 1995, pp. 3–18.
- 713 [19] N. Klarlund and A. Møller, *MONA Version 1.4: User Manual*, Jan 2001.

[20] L. M. Tabajara and M. Y. Vardi, "Partitioning techniques in LTLf
714 synthesis," in *IJCAI*. AAAI Press, 2019, pp. 5599–5606. 715