

Engineering an LTL_f Synthesis Tool

Alexandre Duret-Lutz¹ , Shufang Zhu² , Nir Piterman³ ,
Giuseppe De Giacomo⁴ , and Moshe Y. Vardi⁵ 

¹ LRE, EPITA, Le Kremlin-Bicêtre, France

² University of Liverpool, Liverpool, UK

³ University of Gothenburg and Chalmers University of Technology, Gothenburg, Sweden

⁴ Sapienza University of Rome, Rome, Italy

⁵ Rice University, Houston, Texas, USA

Abstract. The problem of LTL_f reactive synthesis is to build a transducer, whose output is based on a history of inputs, such that, for every infinite sequence of inputs, the conjoint evolution of the inputs and outputs has a prefix that satisfies a given LTL_f specification.

We describe the implementation of an LTL_f synthesizer that outperforms existing tools on our benchmark suite. This is based on a new, direct translation from LTL_f to a DFA represented as an array of Binary Decision Diagrams (MTBDDs) sharing their nodes. This MTBDD-based representation can be interpreted directly as a reachability game that is solved on-the-fly during its construction.

1 Introduction

Reactive synthesis is concerned with synthesizing programs (a.k.a. strategies) for reactive computations (e.g., processes, protocols, controllers, robots) in active environments [47,30,26], typically, from temporal logic specifications. In AI, Reactive Synthesis, which is related to (strong) planning for temporally extended goals in fully observable nondeterministic domains [16,3,4,14,6,34,21,15], has been studied with a focus on logics on finite traces such as LTL_f [33,7,22,23]. In fact, LTL_f synthesis [23] is one of the two main success stories of reactive synthesis so far (the other being the GR(1) fragment of LTL [46]), and has brought about impressive advances in scalability [56,8,18,20].

Reactive synthesis for LTL_f involves the following steps [23]: (1) distinguishing uncontrollable input ($\mathcal{I}$) and controllable output ($\mathcal{O}$) variables in an LTL_f specification φ of the desired system behavior; (2) constructing a DFA accepting the behaviors satisfying φ ; (3) interpreting this DFA as a two-player reachability game, and finding a controller winning strategy. Step (2) has two main bottlenecks: the DFA is worst-case doubly-exponential and its propositional alphabet $\Sigma = 2^{\mathcal{I} \cup \mathcal{O}}$ is exponential. The first only happens in the worst case, while the second blow-up – which we call *alphabet explosion* – always happens.

Mona [39] addresses the alphabet-explosion problem, which happens also in MSO, by representing a DFA with Multi-Terminal Binary Decision Diagrams (MTBDDs) [36]. MTBDDs are a variant of BDDs [12] with arbitrary terminal

values. If terminal values encode destination states, an MTBDD can compactly represent all outgoing transitions of a single DFA state. A DFA is represented, through its transition function, as an array of MTBDDs sharing their nodes.

The first LTL_f synthesizer, Syft [56], converted LTL_f into first-order logic in order to build a MTBDD-encoded DFA with Mona. Syft then converted this DFA into a BDD representation to solve the reachability game using a symbolic fixpoint computation. Syft demonstrated that DFA construction is the main bottleneck in LTL_f synthesis, motivating several follow-up efforts.

One approach to effective DFA construction uses compositional techniques, decomposing the input LTL_f formula into smaller subformulas whose DFAs can be minimized before being recombined. Lisa [8] decomposes top-level conjunctions, while Lydia [19] and LydiaSyft [29] decompose every operator.

Compositional methods construct the full DFA before synthesis can proceed, limiting their scalability. On-the-fly approaches [53] construct the DFA incrementally, while simultaneously solving the game, allowing strategies to be found before the complete DFA is built. The DFA construction may use various techniques. Cynthia [20] uses Sentential Decision Diagrams (SDDs) [17] to generate all outgoing transitions of a state at once. Alternatively, Nike [28] and MoGuSer [54] use a SAT-based method to construct one successor at a time. The game is solved by forward exploration with suitable backpropagation.

Contributions and Outline In Section 3, we propose a direct and efficient translation from LTL_f to MTBDD-encoded DFA (henceforth called MTDFA). In Section 4, we show that given an appropriate ordering of BDD variables, LTL_f realizability can be solved by interpreting the MTBDD nodes of the MTDFA as the vertices of a reachability game, known to be solvable in linear time by backpropagation of the vertices that are winning for the *output* player. We give a linear-time implementation for solving the game on-the-fly while it is constructed. For more opportunities to abort the on-the-fly construction earlier, we additionally backpropagate vertices that are known to be winning by the *input* player. We implemented these techniques in two tools ([ltlf2dfa](#) and [ltlfsynt](#)) that compare favorably with other existing tools in benchmarks from the LTL_f -Synthesis Competition. To meet space limits, Section 5 only reports on the LTL_f realizability benchmark, and we refer readers to our artifact for the other results [24].

2 Preliminaries

2.1 Words over Assignments

A *word over* σ of length n over an alphabet Σ is a function $\sigma : \{0, 1, \dots, n-1\} \rightarrow \Sigma$. We use Σ^n (resp. Σ^* and Σ^+) to denote the set of words of length n (resp. any length $n \geq 0$ and $n > 0$). We use $|\sigma|$ to represent the length of a word σ . For $\sigma \in \Sigma^n$ and $0 \leq i < n$, $\sigma(..i)$ denotes the prefix of σ of length $i + 1$.

Let $\mathcal{P}$ be a finite set of Boolean variables (a.k.a. *atomic propositions*). We use $\mathbb{B}^{\mathcal{P}}$ to denote the set of all assignments, i.e., functions $\mathcal{P} \rightarrow \mathbb{B}$ mapping variables to values in $\mathbb{B} = \{\perp, \top\}$.

Given two disjoint sets of variables $\mathcal{P}_1$ and $\mathcal{P}_2$, and two assignments $w_1 \in \mathbb{B}^{\mathcal{P}_1}$ and $w_2 \in \mathbb{B}^{\mathcal{P}_2}$, we use $w_1 \sqcup w_2 : (\mathcal{P}_1 \cup \mathcal{P}_2) \rightarrow \mathbb{B}$ to denote their combination.

In a system modeled using discrete Boolean signals that evolve synchronously, we assign a variable to each signal, and use a word $\sigma \in (\mathbb{B}^{\mathcal{P}})^+$ over assignments of $\mathcal{P}$ to represent the conjoint evolution of all signals over time.

We extend $\sqcup$ to such words. For two words $\sigma_1 \in (\mathbb{B}^{\mathcal{P}_1})^n$, $\sigma_2 \in (\mathbb{B}^{\mathcal{P}_2})^n$ of length n over assignments that use disjoint sets of variables, we use $\sigma_1 \sqcup \sigma_2 \in (\mathbb{B}^{\mathcal{P}_1 \cup \mathcal{P}_2})^n$ to denote a word such that $(\sigma_1 \sqcup \sigma_2)(i) = \sigma_1(i) \sqcup \sigma_2(i)$ for $0 \leq i < n$.

2.2 Linear Temporal Logic over Finite, Nonempty Words.

We use classical LTL_f semantics over nonempty finite words [22].

Definition 1 (LTL_f formulas). An LTL_f formula φ is built from a set $\mathcal{P}$ of variables, using the following grammar where $p \in \mathcal{P}$, and $\odot \in \{\wedge, \vee, \rightarrow, \leftrightarrow, \dots\}$ is any Boolean operator: $\varphi ::= tt \mid ff \mid p \mid \neg\varphi \mid \varphi \odot \varphi \mid X\varphi \mid X^!\varphi \mid \varphi U \varphi \mid \varphi R \varphi \mid G\varphi \mid F\varphi$.

Symbols *tt* and *ff* represent the true and false LTL_f formulas. Temporal operators are *X* (weak next), *X[!]* (strong next), *U* (until), *R* (release), *G* (globally), and *F* (finally). $\text{LTL}_f(\mathcal{P})$ denotes the set of formulas produced by the above grammar. We use $\text{sf}(\varphi)$ to denote the set of subformulas for φ . A maximal temporal subformula of φ is a subformula whose primary operator is temporal and that is not strictly contained within any other temporal subformula of φ .

The satisfaction of a formula $\varphi \in \text{LTL}_f(\mathcal{P})$ by word $\sigma \in (\mathbb{B}^{\mathcal{P}})^+$ of length $n > 0$ at position $0 \leq i < n$, denoted $\sigma, i \models \varphi$, is defined as follows.

$$\begin{aligned} \sigma, i \models tt &\iff i < n & \sigma, i \models X\varphi &\iff (i + 1 = n) \vee (\sigma, i + 1 \models \varphi) \\ \sigma, i \models ff &\iff i = n & \sigma, i \models X^!\varphi &\iff (i + 1 < n) \wedge (\sigma, i + 1 \models \varphi) \\ \sigma, i \models p &\iff p \in \sigma(i) & \sigma, i \models F\varphi &\iff \exists j \in [i, n), \sigma, j \models \varphi \\ \sigma, i \models \neg\varphi &\iff \neg(\sigma, i \models \varphi) & \sigma, i \models G\varphi &\iff \forall j \in [i, n), \sigma, j \models \varphi \\ \sigma, i \models \varphi_1 \odot \varphi_2 &\iff (\sigma, i \models \varphi_1) \odot (\sigma, i \models \varphi_2) \\ \sigma, i \models \varphi_1 U \varphi_2 &\iff \exists j \in [i, n), (\sigma, j \models \varphi_2) \wedge (\forall k \in [i, j), \sigma, k \models \varphi_1) \\ \sigma, i \models \varphi_1 R \varphi_2 &\iff \forall j \in [i, n), (\sigma, j \models \varphi_2) \vee (\exists k \in [i, j), \sigma, k \models \varphi_1) \end{aligned}$$

The set of words that satisfy $\varphi \in \text{LTL}_f(\mathcal{P})$ is $\mathcal{L}(\varphi) = \{\sigma \in (\mathbb{B}^{\mathcal{P}})^+ \mid \sigma, 0 \models \varphi\}$.

Example 1. Consider the following LTL_f formulas over $\mathcal{P} = \{i_0, i_1, i_2, o_1, o_2\}$: $\Psi_1 = G((i_0 \rightarrow (o_1 \leftrightarrow i_1)) \wedge ((\neg i_0) \rightarrow (o_1 \leftrightarrow i_2)))$, and $\Psi_2 = (GFo_2) \leftrightarrow (Fi_0)$. If we interpret i_0, i_1, i_2 as input signals, and o_1, o_2 as output signals, formula Ψ_1 specifies a 1-bit multiplexer: the value of the signal o_1 should be equal to the value of either i_1 or i_2 depending on the setting of i_0 . Formula Ψ_2 specifies that the last value of o_2 should be $\top$ if and only if i_0 was $\top$ at some instant.

Definition 2 (Propositional Equivalence [27]). For $\varphi \in \text{LTL}_f(\mathcal{P})$, let φ_P be the Boolean formula obtained from φ by replacing every maximal temporal subformula ψ by a Boolean variable x_ψ . Two formulas $\alpha, \beta \in \text{LTL}_f(\mathcal{P})$ are propositionally equivalent, denoted $\alpha \equiv \beta$, if α_P and β_P are equivalent Boolean formulas.

Example 2. Formulas $\alpha = (Gb) \vee ((Fa) \wedge (Gb))$ and $\beta = Gb$ are propositionally equivalent. Indeed, $\alpha_P = x_{Gb} \vee (x_{Fa} \wedge x_{Gb}) = x_{Gb} = \beta_P$.

Note that $\alpha \equiv \beta$ implies $\mathcal{L}(\alpha) = \mathcal{L}(\beta)$, but the converse is not true in general. Since $\equiv$ is an equivalence relation, we use $[\alpha]_{\equiv} \in \text{LTL}_f(\mathcal{P})$ to denote some unique representative of the equivalence class of α with respect to $\equiv$.

2.3 LTL_f Realizability

Our goal is to build a tool that decides whether an LTL_f formula is *realizable*.

Definition 3 ([23,37]). *Given two disjoint sets of variables $\mathcal{I}$ (inputs) and $\mathcal{O}$ (outputs), a controller is a function $\rho : \mathcal{I}^* \rightarrow \mathcal{O}$, that produces an assignment of output variables given a history of assignments of input variables.*

Given a word of n input assignments $\sigma \in (\mathbb{B}^{\mathcal{I}})^n$, the controller can be used to generate a word of n output assignments $\sigma_\rho \in (\mathbb{B}^{\mathcal{O}})^n$. The definition of σ_ρ may use two semantics depending on whether we want the controller to have access to the current input assignment to decide the output assignment:

Mealy semantics: $\sigma_\rho(i) = \rho(\sigma(..i))$ for all $0 \leq i < n$.

Moore semantics: $\sigma_\rho(i) = \rho(\sigma(..i-1))$ for all $0 \leq i < n$.

A formula $\varphi \in \text{LTL}_f(\mathcal{I} \cup \mathcal{O})$ is said to be Mealy-realizable or Moore-realizable if there exists a controller ρ such that for any word $\sigma \in (\mathbb{B}^{\mathcal{I}})^\omega$ there exists a position k such that $(\sigma \sqcup \sigma_\rho)(..k) \in \mathcal{L}(\varphi)$ using the desired semantics.

Example 3. Formula Ψ_1 (from Example 1) is Mealy-realizable but not Moore-realizable. Formula Ψ_2 is both Mealy and Moore-realizable.

2.4 Multi-Terminal BDDs

Let $\mathcal{S}$ be a finite set. Given a finite set of variables $\mathcal{P} = \{p_0, p_1, \dots, p_{n-1}\}$ (that are implicitly ordered by their index) we use $f : \mathbb{B}^{\mathcal{P}} \rightarrow \mathcal{S}$ to denote a function that maps an assignment of all those variables to an element of $\mathcal{S}$. Given a variable $p \in \mathcal{P}$ and a Boolean $b \in \mathbb{B}$, the function $f_{p=b} : \mathbb{B}^{\mathcal{P} \setminus \{p\}} \rightarrow \mathcal{S}$ represents a generalized co-factor obtained by replacing p by b in f . When $\mathcal{S} = \mathbb{B}$, a function $f : \mathbb{B}^{\mathcal{P}} \rightarrow \mathbb{B}$ can be encoded into a Binary Decision Diagram (BDD) [11]. Multi-Terminal Binary Decision Diagrams (MTBDDs) [44,45,32,39], also called Algebraic Decision Diagrams (ADDs) [5,51], generalize BDDs by allowing arbitrary values on the leaves of the graph.

A *Multi-Terminal BDD* encodes any function $f : \mathbb{B}^{\mathcal{P}} \rightarrow \mathcal{S}$ as a rooted, directed acyclic graph. We use the term *nodes* to refer to the vertices of this graph. All nodes in an MTBDD are represented by triples of the form (p, ℓ, h) . In an internal node, $p \in \mathcal{P}$ and ℓ, h point to successors MTBDD nodes called the **low** and **high** links. The intent is that if (p, ℓ, h) is the root of the MTBDD representing the function f , then ℓ and h are the roots of the MTBDDs representing the functions $f_{p=\perp}$ and $f_{p=\top}$, respectively. Leaves of the graph, called *terminals*, hold values in $\mathcal{S}$. For consistency with internal nodes, we represent terminals with a triple of the form (∞, s, ∞) where $s \in \mathcal{S}$. When comparing the first elements of different triplets, we assume that ∞ is greater than all variables. We

use $\text{MTBDD}(\mathcal{P}, \mathcal{S})$ to denote the set of MTBDD nodes that can appear in the representation of an arbitrary function $\mathbb{B}^{\mathcal{P}} \rightarrow \mathcal{S}$.

Following the classical implementations of BDD packages [11,1], we assume that MTBDDs are *ordered* (variables of $\mathcal{P}$ are ordered and visited in increasing order by all branches of the MTBDD) and *reduced* (isomorphic subgraphs are merged by representing each triplet only once, and internal nodes with identical low and high links are skipped over). Doing so ensures that each function $f : \mathbb{B}^{\mathcal{P}} \rightarrow \mathcal{S}$ has a unique MTBDD representation for a given order of variables.

Given $m \in \text{MTBDD}(\mathcal{P}, \mathcal{S})$ and an assignment $w \in \mathbb{B}^{\mathcal{P}}$, we note $m(w)$ the element of $\mathcal{S}$ stored on the terminal of m that is reached after following the assignment w in the structure of m . We use $|m|$ to denote the number of MTBDD nodes that can be reached from m .

Cf. App. A.1

Let $m_1 \in \text{MTBDD}(\mathcal{P}, \mathcal{S}_1)$ and $m_2 \in \text{MTBDD}(\mathcal{P}, \mathcal{S}_2)$ be two MTBDD nodes representing functions $f_i : \mathbb{B}^{\mathcal{P}} \rightarrow \mathcal{S}_i$, and let $\odot : \mathcal{S}_1 \times \mathcal{S}_2 \rightarrow \mathcal{S}_3$, be a binary operation. One can easily construct $m_3 \in \text{MTBDD}(\mathcal{P}, \mathcal{S}_3)$ representing the function $f_3(p_0, \dots, p_{n-1}) = f_1(p_0, \dots, p_{n-1}) \odot f_2(p_0, \dots, p_{n-1})$, by generalizing the `apply2` function typically found in BDD libraries [32]. We use $m_1 \odot m_2$ to denote the MTBDD that results from this construction.

Cf. App. A.2

For $m \in \text{MTBDD}(\mathcal{P}, \mathcal{S})$ we use $\text{leaves}(m) \subseteq \mathcal{S}$ to denote the elements of $\mathcal{S}$ that label terminals reachable from m . This set can be computed in $\Theta(|m|)$.

Cf. App. A.3.

2.5 MTBDD-Based Deterministic Finite Automata

We now define an MTBDD-based representation of a DFA with a propositional alphabet, inspired by Mona's DFA representation [36,39].

Definition 4 (MTDFA). An MTDFA is a tuple $\mathcal{A} = \langle \mathcal{Q}, \mathcal{P}, \iota, \Delta \rangle$, where $\mathcal{Q}$ is a finite set of states, $\mathcal{P}$ is a finite (and ordered) set of variables, $\iota \in \mathcal{Q}$ is the initial state, $\Delta : \mathcal{Q} \rightarrow \text{MTBDD}(\mathcal{P}, \mathcal{Q} \times \mathbb{B})$ represents the set of outgoing transitions of each state. For a word $\sigma \in (\mathbb{B}^{\mathcal{P}})^*$ of length n , let $(q_i, b_i)_{0 \leq i \leq n}$ be a sequence of pairs defined recursively as follows: $(q_0, b_0) = (\iota, \perp)$, and for $0 < i \leq |\sigma|$, $(q_i, b_i) = \Delta(q_{i-1})(\sigma(i-1))$ is the pair reached by evaluating assignment $\sigma(i-1)$ on $\Delta(q_{i-1})$. The word σ is accepted by $\mathcal{A}$ iff $b_n = \top$. The language of $\mathcal{A}$, denoted $\mathcal{L}(\mathcal{A})$, is the set of words accepted by $\mathcal{A}$.

Example 4. Figure 1 shows an MTDFA where $\mathcal{Q} \subseteq \text{LTL}_f(\{i_0, i_1, i_2, o_1, o_2\})$. The set of states $\mathcal{Q}$ are the dashed rectangles on the left. For each such a state $q \in \mathcal{Q}$, the dashed arrow points to the MTBDD node representing $\Delta(q)$. The MTBDD nodes are shared between all states. If, starting from the initial state ι at the top-left, we read the assignment $w = (i_0 \rightarrow \top, i_1 \rightarrow \top, i_2 \rightarrow \top, o_1 \rightarrow \top, o_2 \rightarrow \top)$, we should follow only the high links (plain arrows) and we reach the $\Psi_1 \wedge (\text{GF}o_2)$ accepting terminal. If we read this assignment a second-time, starting this time from state $\Psi_1 \wedge (\text{GF}o_2)$ on the left, we reach the same accepting terminal. Therefore, non-empty words of the form $www \dots w$ are accepted by this automaton.

An MTDFA can be regarded as a semi-symbolic representation of a DFA over propositional alphabet. From a state q and reading the assignment w , the

Cf. App. C




Fig. 1. An MTDFA where $\mathcal{P} = \{i_0, i_1, i_2, o_0, o_1\}$ and $\mathcal{Q} \subseteq \text{LTL}_f(\mathcal{P})$. Following classical BDD representations a BDD node (p, ℓ, h) is represented by $(p) \xrightarrow{\ell}_h$. A terminal $(\infty, (\alpha, b), \infty)$ is represented by $\boxed{\alpha}$ if $b = \perp$, or $\boxed{\alpha}$ if $b = \top$. Finally, MTBDD $m = \Delta(\alpha)$ representing the successors of state α is indicated with $\boxed{\alpha} \rightarrow m$. Subformula Ψ_1 abbreviates $\text{G}((i_0 \rightarrow (o_1 \leftrightarrow i_1)) \wedge ((\neg i_0) \rightarrow (o_1 \leftrightarrow i_2)))$.

automaton jumps to the state q' that is the result of computing $(q', b) = \Delta(q)(w)$. The value of b indicates whether that assignment is allowed to be the last one of the word being read. By definition, an MTDFA cannot accept the empty word.

MTDFAs are compact representations of DFAs, because the MTBDD representation of the successors of each state can share their common nodes. Boolean operations can be implemented over MTDFAs, with the expected semantics, i.e., $\mathcal{L}(\mathcal{A}_1 \odot \mathcal{A}_2) = \{\sigma \in (\mathbb{B}^{\mathcal{P}})^+ \mid (\sigma \in \mathcal{L}(\mathcal{A}_1)) \odot (\sigma \in \mathcal{L}(\mathcal{A}_2))\}$.

Cf. App. B

3 Translating LTL_f to MTBDD and MTDFA

This section shows how to directly transform a formula $\varphi \in \text{LTL}_f(\mathcal{P})$ into an MTDFA $\mathcal{A}_\varphi = \langle \mathcal{Q}, \mathcal{P}, \varphi, \Delta \rangle$ such that $\mathcal{L}(\varphi) = \mathcal{L}(\mathcal{A}_\varphi)$. The translation is reminiscent of other translations of LTL_f to DFA [22,20], but it leverages the fact that MTBDDs can provide a normal form for LTL_f formulas.

The construction maps states to LTL_f formulas, i.e., $\mathcal{Q} \subseteq \text{LTL}_f(\mathcal{P})$. Terminals appearing in the MTBDDs of $\mathcal{A}_\varphi$ will be labeled by pairs $(\alpha, b) \in \text{LTL}_f(\mathcal{P}) \times \mathbb{B}$, so we use $\text{term}(\alpha, b) = (\infty, (\alpha, b), \infty)$ to shorten the notation from Section 2.4.

The conversion from φ to $\mathcal{A}_\varphi$ is based on the function $\text{tr} : \text{LTL}_f(\mathcal{P}) \rightarrow \text{MTBDD}(\mathcal{P}, \text{LTL}_f(\mathcal{P}) \times \mathbb{B})$ defined inductively as follows:

$$\begin{aligned}
 \text{tr}(ff) &= \text{term}(ff, \perp) & \text{tr}(X\alpha) &= \text{term}(\alpha, \top) \\
 \text{tr}(tt) &= \text{term}(tt, \top) & \text{tr}(X^!\alpha) &= \text{term}(\alpha, \perp) \\
 \text{tr}(p) &= (p, \text{term}(ff, \perp), \text{term}(tt, \top)) \text{ for } p \in \mathcal{P} & \text{tr}(\neg\alpha) &= \neg\text{tr}(\alpha) \\
 \text{tr}(\alpha \odot \beta) &= \text{tr}(\alpha) \odot \text{tr}(\beta) \text{ for any } \odot \in \{\wedge, \vee, \rightarrow, \leftrightarrow, \oplus\} \\
 \text{tr}(\alpha \cup \beta) &= \text{tr}(\beta) \vee (\text{tr}(\alpha) \wedge \text{term}(\alpha \cup \beta, \perp)) & \text{tr}(F\alpha) &= \text{tr}(\alpha) \vee \text{term}(F\alpha, \perp) \\
 \text{tr}(\alpha \text{ R } \beta) &= \text{tr}(\beta) \wedge (\text{tr}(\alpha) \vee \text{term}(\alpha \text{ R } \beta, \top)) & \text{tr}(G\alpha) &= \text{tr}(\alpha) \wedge \text{term}(G\alpha, \top)
 \end{aligned}$$

Boolean operators that appear to the right of the equal sign are applied on MTBDDs as discussed in Section 2.4. Terminals in $\text{LTL}_f(\mathcal{P}) \times \mathbb{B}$ are combined with: $(\alpha_1, b_1) \odot (\alpha_2, b_2) = ([\alpha_1 \odot \alpha_2]_{\equiv}, b_1 \odot b_2)$ and $\neg(\alpha, b) = ([\neg\alpha]_{\equiv}, \neg b)$.

Theorem 1. For $\varphi \in \text{LTL}_f(\mathcal{P})$, let $\mathcal{A}_\varphi = \langle \mathcal{Q}, \mathcal{P}, \iota, \Delta \rangle$ be the MTDFA obtained by setting $\iota = [\varphi]_\equiv$, $\Delta = \text{tr}$, and letting $\mathcal{Q}$ be the smallest subset of $\text{LTL}_f(\mathcal{P})$ such that $\iota \in \mathcal{Q}$, and such that for any $q \in \mathcal{Q}$ and for any $(\alpha, b) \in \text{leaves}(\Delta(q))$, then $\alpha \in \mathcal{Q}$. With this construction, $|\mathcal{Q}|$ is finite and $\mathcal{L}(\varphi) = \mathcal{L}(\mathcal{A}_\varphi)$.

Proof. (sketch) By definition of tr , $\mathcal{Q}$ contains only Boolean combinations of subformulas of φ . Propositional equivalence implies that the number of such combinations is finite: $|\mathcal{Q}| \leq 2^{|\text{sf}(\varphi)|}$. The language equivalence follows from the definition of LTL_f , and from some classical LTL_f equivalences. For instance the rule for $\text{tr}(\alpha \cup \beta)$ is based on the equivalence $\mathcal{L}(\alpha \cup \beta) = \mathcal{L}(\beta \vee (\alpha \wedge X^1(\alpha \cup \beta)))$.

Example 5. Figure 1 is the MTDFA for formula $\Psi_1 \wedge \Psi_2$, presented in Example 1. Many more examples can be found in the associated artifact [24].

Also App. D

The definition of $\text{tr}(\cdot)$ as an MTBDD representation of the set of successors of a state can be thought as a symbolic representation of Antimirov’s linear forms [2] for DFA with propositional alphabets. Antimirov presented linear forms as an efficient way to construct all (partial) derivatives at once, without having to iterate over the alphabet. For LTL_f , *formula progressions* [20] are the equivalent of Brozowski derivatives [13]. Here, $\text{tr}(\cdot)$ computes all formulas progressions at once, without having to iterate over an exponential number of assignments.

Finally, note that while this construction works with any order for $\mathcal{P}$, different orders might produce a different number of MTBDD nodes.

Optimizations The previous definitions can be improved in several ways.

Our implementation of MTBDD actually supports terminals that are the Boolean terminals of standard BDDs as well as the terminals used so far. So we are actually using $\text{MTBDD}(\mathcal{P}, (\text{LTL}_f(\mathcal{P}) \times \mathbb{B}) \cup \mathbb{B})$, and we encode $\text{term}(ff, \perp)$ and $\text{term}(tt, \top)$ directly as $\perp$ and $\top$ respectively. With those changes, `apply2` may be modified to shortcut the recursion depending on the values of m_1 , m_2 , and $\odot$. For instance if $\odot = \wedge$ and $m_1 = \top$, then m_2 can be returned immediately. Such shortcuts may be implemented for $\text{MTBDD}(\mathcal{P}, \mathcal{S} \cup \mathbb{B})$ regardless of the nature of $\mathcal{S}$, so our implementation of MTBDD operations is independent of LTL_f .

Cf. App A.4

When combining terminals during the computation of tr , one has to compute the representative formula $[\alpha_1 \odot \alpha_2]_\equiv$. This can be done by converting α_{1P} and α_{2P} into BDDs, keeping track of such conversions in a hash table. Two propositionally equivalent formulas will have the same BDD representation. While we are looking for a representative formula, we can also use the opportunity to simplify the formula at hand. We use the following very simple rewritings, for patterns that occur naturally in the output of tr :

$$(\alpha \cup \beta) \vee \beta \rightsquigarrow \alpha \cup \beta, \quad (\alpha \mathbf{R} \beta) \wedge \beta \rightsquigarrow \alpha \mathbf{R} \beta, \quad (\mathbf{F}\beta) \vee \beta \rightsquigarrow \mathbf{F}\beta, \quad (\mathbf{G}\beta) \wedge \beta \rightsquigarrow \mathbf{G}\beta.$$

Once $\mathcal{A}_\varphi$ has been built, two states $q, q' \in \mathcal{Q}$ such that $\Delta(q) = \Delta(q')$ can be merged by replacing all occurrences of q' by q in the leaves of Δ .

Example 6. The automaton from Figure 1 has two pairs of states that can be merged. However, if the rule $(\mathbf{G}\beta) \wedge \beta \rightsquigarrow \mathbf{G}\beta$ is applied during the construction, then the occurrence of $(\mathbf{G}\mathbf{F}o_2) \wedge (\mathbf{F}o_2)$ will already be replaced by $\mathbf{G}\mathbf{F}o_2$, producing the simplified automaton without requiring any merging.

Cf. App. C & D




Fig. 2. Interpretation of the MTBDD of Figure 1 as a game with $\mathcal{I} = \{i_0, i_1, i_2\}$, $\mathcal{O} = \{o_1, o_2\}$. Each MTBDD *node* of the MTBDD is viewed as a *vertex* of the game, with terminal of the form $(\alpha, \perp)$ looping back to $\Delta(\alpha)$. Player O decides where to go from diamond and rectangular vertices and wants to reach the **green vertices** corresponding to accepting terminals. Player I decides where to go from round vertices and wants to reach **ff** or avoid green vertices.

4 Deciding LTL_f Realizability

LTL_f realizability (Def. 3) is solved by reducing the problem to a two-player reachability game where one player decides the input assignments and the other player decides the output assignments [23]. Section 4.1 presents reachability games and how to interpret the MTBDD as a reachability game, and Section 4.2 shows how we can solve the game on-the-fly while constructing it.

4.1 Reachability Games & Backpropagation

Definition 5 (Reachability Game). A *Reachability Game* is $\mathcal{G} = \langle \mathcal{V} = \mathcal{V}_o \uplus \mathcal{V}_i, \mathcal{E}, \mathcal{F}_o \rangle$, where $\mathcal{V}$ is a finite set of vertices partitioned to player output (abbreviated O) and player input (abbreviated I), $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is a finite set of edges, and $\mathcal{F}_o \subseteq \mathcal{V}$ is the set of target states. Let $\mathcal{E}(v) = \{(v, v') \mid (v, v') \in \mathcal{E}\}$. This graph is also referred to as the *game arena*.

A *strategy* for player O is a cycle-free subgraph $\langle W, \sigma \rangle \subseteq \langle \mathcal{V}, \mathcal{E} \rangle$ such that (a) for every $v \in W$ we have $v \in \mathcal{F}_o$ or $\mathcal{E}(v) \cap \sigma \neq \emptyset$ and (b) if $v \in W \cap \mathcal{V}_i$ then $\mathcal{E}(v) \subseteq \sigma$. A vertex v is *winning* for O if $v \in W$ for some strategy $\langle W, \sigma \rangle$.

Such a reachability game can be solved by backpropagation identifying the maximal set W in a strategy. Namely, start from $W = \mathcal{F}_o$. Then W is iteratively augmented with every vertex in $\mathcal{V}_o$ that has some edge to W , and every (non dead-end) vertex in $\mathcal{V}_i$ whose edges all lead to W . At the end of this backpropagation, which can be performed in linear time [35, Theorem 3.1.2], every vertex in W is winning for O, and every vertex outside W is losing for O. Notice that every dead-end that is not in $\mathcal{F}_o$ cannot be winning. It follows that we can identify some (but not necessarily all) vertices that are losing by setting L as the set of all dead-ends and adding to L every $\mathcal{V}_i$ vertex that has some edge to L and every $\mathcal{V}_o$ vertex whose edges all lead to L .

Let $\mathcal{A}_\varphi = \langle \mathcal{Q}, \mathcal{I} \uplus \mathcal{O}, \iota, \Delta \rangle$ be a translation of $\varphi \in \text{LTL}_f(\mathcal{I} \uplus \mathcal{O})$ (per Th. 1) such that variables of $\mathcal{I}$ appear before $\mathcal{O}$ in the MTBDD encoding of Δ .

Definition 6 (Realizability Game). We define the reachability game $\mathcal{G}_\varphi = \langle \mathcal{V} = \mathcal{V}_I \uplus \mathcal{V}_O, \mathcal{E}, \mathcal{F}_O \rangle$ in which $\mathcal{V} \subseteq \text{MTBDD}(\mathcal{I} \uplus \mathcal{O})$ corresponds the set of nodes that appear in the MTBDD encoding of Δ . $\mathcal{V}_O$ contains all nodes (p, ℓ, h) such that $p \in \mathcal{O}$ or $p = \infty$ (terminals), and $\mathcal{V}_I$ contains those with $p \in \mathcal{I}$. The edges $\mathcal{E}$ follows the structure of Δ , i.e., if $\mathcal{A}_\varphi$ has a node $r = (p, \ell, h)$, then $\{(r, \ell), (r, h)\} \subseteq \mathcal{E}$. Additionally, for any terminal $t = (\infty, (\alpha, \perp), \infty)$ such that $\alpha \neq \text{ff}$, $\mathcal{E}$ contains the edge $(t, \Delta(\alpha))$. Finally, $\mathcal{F}_O$ is the set of accepting terminals, i.e., nodes of the form $(\infty, (\alpha, \top), \infty)$.

Theorem 2. Vertex $\Delta(\iota)$ is winning for $\mathcal{O}$ in $\mathcal{G}_\varphi$ iff φ is Mealy-realizable.

Moore realizability can be checked similarly by changing the order of $\mathcal{I}$ and $\mathcal{O}$ in the MTBDD encoding of Δ .

Example 7. Figure 2 shows how to interpret the MTDFA of Figure 1 as a game, by turning each MTBDD node into a game vertex. The player owning each vertex is chosen according to the variable that labels it. Vertices corresponding to accepting terminals become winning targets for the output player, so the game stops once they are reached. Solving this game will find every internal node as winning for $\mathcal{O}$, so the corresponding formula is Mealy-realizable.

The difference with DFA games [23,20,28,54] is that instead of having player $\mathcal{I}$ select all input signals at once, and then player $\mathcal{O}$ select all output signals at once, our game proceeds by selecting one signal at a time. Sharing nodes that represent identical partial assignments contributes to the scalability of our approach.

4.2 Solving Realizability On-the-fly

We now show how to construct and solve $\mathcal{G}_\varphi$ on-the-fly, for better efficiency. The construction is easier to study in two parts: (1) the on-the-fly solving of reachability games, based on backpropagation, and (2) the incremental construction of $\mathcal{G}_\varphi$, done with a forward exploration of a subset of the MTDFA for φ .

Algorithm 1 presents the first part: a set of functions for constructing a game arena incrementally, while performing the linear-time backpropagation algorithm on-the-fly. At all points during this construction, the winning status of a vertex ($\text{winner}[x]$) will be one of $\mathcal{O}$ (player $\mathcal{O}$ can force the play to reach $\mathcal{F}_O$, i.e., the vertex belongs to W), $\mathcal{I}$ (player $\mathcal{I}$ can force the play to avoid $\mathcal{F}_O$, i.e., the vertex belongs to L), or $\mathcal{U}$ (undetermined yet), and the algorithm will backpropagate both $\mathcal{O}$ and $\mathcal{I}$. At the end of the construction, all vertices with status $\mathcal{U}$ will be considered as winning for $\mathcal{I}$. Like in the standard algorithm for solving reachability games [35, Th. 3.1.2] each state uses a counter (*count*, lines 7,15) to track the number of its undetermined successors. When a vertex x is marked as winning for player w by calling `set_winner(x,w)`, an undetermined predecessor p has its counter decreased (line 15), and p can be marked as winning for w (line 16) if either vertex p is owned by w (player w can choose to go to x) or the counter dropped to 0 (meaning that all choices at p were winning for w).

To solve the game while it is constructed, we *freeze* vertices. A vertex should be frozen after all its successors have been introduced with `new_edge`. The

Cf. App. D

```

var: owner[]; // map each vertex to one of {0,1}
var: pred[]; // map vertices to sets of predecessor vertices
var: count[]; // map vertices to # of undetermined successors
var: winner[]; // map vertices to one of {0,1,U}
var: frozen[]; // map vertices to their frozen status (a Boolean)
1 Function new_vertex( $x \in \mathcal{V}$ ,  $own \in \{0,1\}$ ) // new vertex owned by own
2   | owner[x]  $\leftarrow$  own; pred[x]  $\leftarrow$   $\emptyset$ ; count[x]  $\leftarrow$  0; frozen[x]  $\leftarrow$   $\perp$ ;
3   | winner[x]  $\leftarrow$  U; // undetermined winner
4 Function new_edge( $src \in \mathcal{V}$ ,  $dst \in \mathcal{V}$ )
5   | assert (frozen[src] =  $\perp$ );
6   | if winner[dst] = U then
7     | count[src]  $\leftarrow$  count[src] + 1; pred[dst]  $\leftarrow$  pred[dst]  $\cup$  {src};
8   | else if winner[dst] = owner[src] then set_winner(src, owner[src]);
   | // ignore the edge otherwise, it will never be used
9 Function freeze_vertex( $x \in \mathcal{V}$ ) // promise not to add more successors
10  | frozen[x]  $\leftarrow$   $\top$ ; // next line, we assume  $\neg 1 = 0$  and  $\neg 0 = 1$ 
11  | if winner[x] = U  $\wedge$  count[n] = 0 then set_winner(x,  $\neg$ owner[x]);
12 Function set_winner( $x \in \mathcal{V}$ ,  $w \in \{0,1\}$ ) // with linear backprop.
13  | assert (winner[x] = U); winner[x]  $\leftarrow$  w;
14  | foreach  $p \in pred[x]$  such that winner[p] = U do
15    | count[p]  $\leftarrow$  count[p] - 1;
16  | if owner[p] = w  $\vee$  (count[p] = 0  $\wedge$  frozen[p]) then set_winner(p, w);

```

Algorithm 1: API for solving a reachability game on-the-fly. Construct the game arena with `new_vertex` and `new_edge`. Once all successors of a vertex have been connected, call `freeze_vertex`. Call `set_winner` at any point to designate vertices winning for one player.

counter dropping to 0 is only checked on frozen vertices (lines 11, 16) since it is only meaningful if all successors of a vertex are known.

Algorithm 2 is the second part. It shows how to build $\mathcal{G}_\varphi$ incrementally. It translates the states α of the corresponding MTDFA one at a time, and uses the functions of Algorithm 1 to turn each node of $\text{tr}(\alpha)$ into a vertex of the game. Since the functions of Algorithm 1 update the winning status of the states as soon as possible, Algorithm 2 can use that to cut parts of the exploration.

Instead of using $\Delta(\varphi) = \text{tr}(\varphi)$ as initial vertex of the game, as in Theorem 2, we consider $\text{init} = \text{term}(\varphi, \perp)$ as initial vertex (line 3): this makes no theoretical difference, since $\text{term}(\varphi, \perp)$ has $\text{tr}(\varphi)$ as unique successor. Lines 5,9–11,13, and 32 implements the exploration of all the LTL_f formulas α that would label the states of the MTDFA for φ (as needed to implement Theorem 1). The actual order in which formulas are removed from *todo* on line 11 is free. (We found out that handling *todo* as a queue to implement a BFS exploration worked marginally better than using it as a stack to do a DFS-like exploration, so we use a BFS in practice.)

Each α is translated into an MTBDD $\text{tr}(\alpha)$ representing its possible successors. The constructed game should have one vertex per MTBDD node in $\text{tr}(\alpha)$. Those vertices are created in the inner **while** loop (lines 19–31). Function `declare_vertex` is used to assign the correct owner to each new node according to its decision variable (as in Def. 6) as well as adding those nodes to the

```

1 Function realizability( $\varphi \in \text{LTL}_f(\mathcal{I} \uplus \mathcal{O})$ )
2   configure the MTBDD library to put variables in  $\mathcal{I}$  before those in  $\mathcal{O}$ ;
3    $init \leftarrow \text{term}(\varphi, \perp)$ ;  $\text{new\_vertex}(init, 1)$ ;
4    $\mathcal{V} \leftarrow \{init\}$ ; // nodes created as game vertices
5    $\mathcal{Q} \leftarrow \emptyset$ ; // LTLf formulas processed by main loop on line 10
6   Function declare_vertex( $r \in \text{MTBDD}(\mathcal{I} \uplus \mathcal{O}, \text{LTL}_f(\mathcal{I} \uplus \mathcal{O}) \times \mathbb{B})$ )
7      $(p, \ell, h) \leftarrow r$ ; if  $p = \infty \vee p \in \mathcal{I}$  then  $own \leftarrow 1$  else  $own \leftarrow 0$ ;
8      $\text{new\_vertex}(r, own)$ ;  $\mathcal{V} \leftarrow \mathcal{V} \cup \{r\}$ ;  $to\_encode \leftarrow to\_encode \cup \{r\}$ ;
9    $todo \leftarrow \{\varphi\}$ ;
10  while  $todo \neq \emptyset \wedge \text{winner}[init] = \mathcal{U}$  do
11     $\alpha \leftarrow todo.\text{pop\_any}()$ ;  $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{\alpha\}$ ;
12    [optional: add one-step (un)realizability check here, see Sec. 5];
13     $a \leftarrow \text{term}(\alpha, \perp)$ ;  $m \leftarrow \text{tr}(\alpha)$ ;
14    if  $m \in \mathcal{V}$  then //  $m$  has already been encoded
15       $\text{new\_edge}(a, m)$ ;  $\text{freeze\_vertex}(a)$ ; continue to line 10;
16     $to\_encode \leftarrow \emptyset$ ;  $leaves \leftarrow \emptyset$ ;
17     $\text{declare\_vertex}(m)$ ;  $\text{new\_edge}(a, m)$ ;  $\text{freeze\_vertex}(a)$ ;
18    if  $\text{winner}[a] \neq \mathcal{U}$  then continue to line 10;
19    while  $to\_encode \neq \emptyset$  do
20       $r \leftarrow to\_encode.\text{pop\_any}()$ ;
21       $(p, \ell, h) \leftarrow r$ ;
22      if  $p = \infty$  then // this is a terminal labeled by  $\ell$ 
23         $(\beta, b) \leftarrow \ell$ ;
24        if  $b$  then  $\text{set\_winner}(r, 0)$ ;
25        else if  $\beta = \text{ff}$  then  $\text{set\_winner}(r, 1)$ ;
26        else if  $\beta \notin \mathcal{Q}$  then  $leaves \leftarrow leaves \cup \{\beta\}$ ;
27      else
28        if  $\ell \notin \mathcal{V}$  then  $\text{declare\_vertex}(\ell)$ ;
29        if  $h \notin \mathcal{V}$  then  $\text{declare\_vertex}(h)$ ;
30         $\text{new\_edge}(r, \ell)$ ;  $\text{new\_edge}(r, h)$ ;  $\text{freeze\_vertex}(r)$ ;
31        if  $\text{winner}[a] \neq \mathcal{U}$  then continue to line 10;
32     $todo \leftarrow todo \cup leaves$ ;
33  return  $\text{winner}[init] = 0$ ;

```

Algorithm 2: On-the-fly realizability check with Mealy semantics (for Moore semantics, swap the order of $\mathcal{I}$ and $\mathcal{O}$ on the first line).

to_encode set processed by this inner loop. Terminal nodes are either marked as winning for one of the players (lines 24–25) or stored in $leaves$ (line 26).

Since connecting game vertices may backpropagate their winning status, the encoding loop can terminate early whenever the vertex associated to $\text{term}(\alpha, \perp)$ becomes determined (lines 18 and 31). If that vertex is not determined, the $leaves$ of α are added to $todo$ (line 32) for further exploration.

The entire construction can also stop as soon as the initial vertex is determined (line 10). However, if the algorithm terminates with $\text{winning}[init] = \mathcal{U}$, it still means that 0 cannot reach its targets. Therefore, as tested by line 33, formula φ is realizable iff $\text{winning}[init] = 0$ in the end.

Theorem 3. *Algorithm 2 returns tt iff φ is Mealy-realizable.*

5 Implementation and Evaluation

Cf. App. F

Our algorithms have been implemented in Spot [25], after extending its fork of BuDDy [42] to support MTBDDs. The release of Spot 2.14 distributes two new command-line tools: [ltlf2dfa](#) and [ltlfsynt](#), implementing translation from LTL_f to MTDFA, and solving LTL_f synthesis. We describe and evaluate [ltlfsynt](#) in the following.

Cf. App. G

Preprocessing Before executing Algorithm 2, we use a few preprocessing techniques to simplify the problem. We remove variables that always have the same polarity in the specification (a simplification used also by Strix [50]), and we decompose the specifications into output-disjoint sub-specifications that can be solved independently [31]. A specification such as $\Psi_1 \wedge \Psi_2$, from Example 1, is not solved directly as demonstrated here, but split into two output-disjoint specifications Ψ_1 and Ψ_2 that are solved separately. Finally, we also simplify LTL_f formulas using very simple rewriting rules such as $X(\alpha) \wedge X(\beta) \rightsquigarrow X(\alpha \wedge \beta)$ that reduce the number of MTBDD operations required during translation.

Cf. App. H

One-step (un)realizability checks An additional optimization consists in performing one-step realizability and one-step unrealizability checks in Algorithm 2. The principle is to transform the formula α into two smaller Boolean formulas α_r and α_u , such that if α_r is realizable it implies that α is realizable, and if α_u is unrealizable it implies that α is unrealizable [53, Theorems 2–3]. Those Boolean formulas can be translated to BDDs for which realizability can be checked by quantification. On success, it avoids the translation of the larger formula α . The simple formula $\Psi_1 \wedge \Psi_2$ of our running example is actually one-step realizable.

Synthesis After deciding realizability, [ltlfsynt](#) is able to extract a strategy from the solved game in the form of a Mealy machine, and encode that into an And-Invert Graph (AIG) [10]: the expected output of the Synthesis Competition for the LTL_f synthesis tracks. The conversion from Mealy to AIG reuses prior work [48,49] developed for Spot’s LTL (not LTL_f) synthesis tool. We do not detail nor evaluate these extra steps here due to lack of space.

Evaluation We evaluated the task of deciding LTL_f reachability over specifications from the Synthesis Competition [38]. We took all [tlsf-fin](#) specifications from [SyntComp’s repository](#), excluded some duplicate specifications as well as some specifications that were too large to be solved by any tool, and converted the specifications from TLSF v1.2 [37] to LTL_f using [syfco](#) [37].

We used BenchExec 3.22 [9] to track time and memory usage of each tool. Tasks were run on a Core i7-3770 with *Turbo Boost* disabled, and frequency scaled down to 1.6GHz to prevent CPU throttling. The computer has 4 physical cores and 16GB of memory. BenchExec was configured to run up to 3 tasks in parallel with a memory limit of 4GB per task, and a time limit of 15 minutes.

More in App. I.

Figure 3 compares five configurations of [ltlfsynt](#) against seven other tools. We verified that all tools were in agreement. Lydia 0.1.3 [19], SyftMax (or Syft 2.0) [55] and LydiaSyft 0.1.0-alpha [29] are all using Mona to construct a DFA by composition; they then solve the resulting game symbolically after encoding




Fig. 3. Cactus plots comparing time and memory usage of different configurations.

it using BDDs. Lisa [8] uses a hybrid compositional construction, mixing explicit compositions (using Spot), with symbolic compositions (using BuDDy), solving the game symbolically in the end. Cynthia 0.1.0 [20], Nike 0.1.0 [28], and MoGuSer [54] all use an on-the-fly construction of a DFA game that they solve via forward exploration with backpropagation, but they do not implement backpropagation in linear time, as we do. Yet, the costly part of synthesis is game generation, not solving. Cynthia uses SDDs [17] to compute successors and represent states, while Nike and MoGuSer use SAT-based techniques to compute successors and BDDs to represent states. Nike, Lisa, and LydiaSyft were the top-3 contenders of the LTL_f track of SyntComp in 2023 and 2024.

Configuration `ltlf synt:bfs-nopre` corresponds to Algorithm 2 where *todo* is a queue: it already solves more cases than all other tested tools. Suffix `-nopre` indicates that preprocessings of the specification are disabled (this makes comparison fairer, since other tools have no such preprocessings). The version with preprocessings enabled is simply called `ltlf synt:bfs`. Variants with “-os” adds the one-step (un)realizability checks that LydiaSyft, Cynthia, and Nike also perform.

We also include a configuration `ltlfsynt:full-N` that corresponds to first translating the specification into a MTDFA using Theorem 1, and then solving the game by linear propagation. The difference between `ltlfsynt:full` and `ltlfsynt:bfs` shows the gain obtained with the on-the-fly translation: although that look small in the cactus plot, it is important in some specifications.

Tab. 2–3 in App. I.

Data Availability Statement Implementation, supporting scripts, detailed analysis of this benchmark, and additional examples are archived on Zenodo [24].

6 Conclusion

We have presented the implementation of `ltlfsynt`, and evaluated it to be faster at deciding LTL_f realizability than seven existing tools, including the winners of SyntComp’24. The implementation uses a direct and efficient translation from LTL_f to DFA represented by MTBDDs, which can then be solved as a game played directly on the structure of the MTBDDs. The two constructions (translation and game solving) are performed together on-the-fly, to allow early termination.

Although `ltlfsynt` also includes a preliminary implementation of LTL_f synthesis of And-Inverter graphs, we leave it as future work to document it and ensure its correctness.

Finally, the need for solving a reachability game while it is discovered also occurs in other equivalent contexts such as HornSAT, where linear algorithms that do not use “counters” and “predecessors” (unlike ours) have been developed [43]. Using such algorithms might improve our solution by saving memory.

References

1. Andersen, H.R.: An introduction to binary decision diagrams. Lecture notes for Efficient Algorithms and Programs, Fall 1999 (1999), <https://web.archive.org/web/20090530154634/http://www.itu.dk:80/people/hra/bdd-eap.pdf>
2. Antimirov, V.: Partial derivatives of regular expressions and finite automaton constructions. *Theoretical Computer Science* **155**(2), 291–319 (Mar 1996). [https://doi.org/10.1016/0304-3975\(95\)00182-4](https://doi.org/10.1016/0304-3975(95)00182-4)
3. Bacchus, F., Kabanza, F.: Planning for temporally extended goals. *Annals of Mathematics and Artificial Intelligence* **22**, 5–27 (1998). <https://doi.org/10.1023/A:1018985923441>
4. Bacchus, F., Kabanza, F.: Using temporal logics to express search control knowledge for planning. *Artificial Intelligence* **116**(1–2), 123–191 (2000). [https://doi.org/10.1016/S0004-3702\(99\)00071-5](https://doi.org/10.1016/S0004-3702(99)00071-5)
5. Bahar, R.I., Frohm, E.A., Gaona, C.M., Hachtel, G.D., Macii, E., Pardo, A., Somenzi, F.: Algebraic decision diagrams and their applications. In: *Proceedings of 1993 International Conference on Computer Aided Design (ICCAD’93)*. pp. 188–191. IEEE Computer Society Press (Nov 1993). <https://doi.org/10.1109/ICCAD.1993.580054>
6. Baier, J.A., Fritz, C., McIlraith, S.A.: Exploiting procedural domain control knowledge in state-of-the-art planners. In: *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS’07)*. pp. 26–33. AAAI (2007), <https://aaai.org/papers/icaps-07-004>

7. Baier, J.A., McIlraith, S.A.: Planning with first-order temporally extended goals using heuristic search. In: Proceedings of the 21st national conference on Artificial intelligence (AAAI'06). pp. 788–795. AAAI Press (2006). <https://doi.org/10.5555/1597538.1597664>
8. Bansal, S., Li, Y., Tabajara, L.M., Vardi, M.Y.: Hybrid compositional reasoning for reactive synthesis from finite-horizon specifications. In: Proceedings of the 34th national conference on Artificial intelligence (AAAI'20). pp. 9766–9774. AAAI Press (2020). <https://doi.org/10.1609/AAAI.V34I06.6528>
9. Beyer, D., Löwe, S., Wendler, P.: Reliable benchmarking: requirements and solutions. *International Journal on Software Tools for Technology Transfer* **21**, 1–29 (Feb 2019). <https://doi.org/10.1007/s10009-017-0469-y>
10. Biere, A., Heljanko, K., Wieringa, S.: AIGER 1.9 and beyond. Tech. Rep. 11/2, Institute for Formal Models and Verification, Johannes Kepler University, Altenbergerstr. 69, 4040 Linz, Austria (2011), <https://fmv.jku.at/aiger/>
11. Bryant, R.E.: Graph-based algorithms for boolean function manipulation. *IEEE Transactions on Computers* **35**(8), 677–691 (Aug 1986). <https://doi.org/10.1109/TC.1986.1676819>
12. Bryant, R.E.: Symbolic boolean manipulation with ordered binary-decision diagrams. *ACM Comput. Surv.* **24**(3), 293–318 (Sep 1992). <https://doi.org/10.1145/136035.136043>
13. Brzozowski, J.A.: Derivatives of regular expressions. *Journal of the ACM* **11**(4), 481–494 (Oct 1964). <https://doi.org/10.1145/321239.321249>
14. Calvanese, D., De Giacomo, G., Vardi, M.Y.: Reasoning about actions and planning in LTL action theories. In: Proceedings of the Eight International Conference on Principles of Knowledge Representation and Reasoning (KR'02). pp. 593–602. Morgan Kaufmann (2002). <https://doi.org/10.5555/3087093.3087142>
15. Camacho, A., Bienvenu, M., McIlraith, S.A.: Towards a unified view of AI planning and reactive synthesis. In: Proceedings of the 29th International Conference on Automated Planning and Scheduling (ICAPS'19). pp. 58–67. AAAI Press (2019). <https://doi.org/10.1609/icaps.v29i1.3460>
16. Cimatti, A., Pistore, M., Roveri, M., Traverso, P.: Weak, strong, and strong cyclic planning via symbolic model checking. *Artificial Intelligence* **147**(1–2), 35–84 (2003). [https://doi.org/10.1016/S0004-3702\(02\)00374-0](https://doi.org/10.1016/S0004-3702(02)00374-0)
17. Darwiche, A.: SDD: A new canonical representation of propositional knowledge bases. In: Proceedings of the 22nd International Joint Conference on Artificial Intelligence. pp. 819–826. AAAI Press (2011). <https://doi.org/10.5591/978-1-57735-516-8/IJCAI11-143>
18. De Giacomo, G., Favorito, M.: Compositional approach to translate LTL_f/LDL_f into deterministic finite automata. In: Proceedings of the 31st International Conference on Automated Planning and Scheduling (ICAPS'21). pp. 122–130 (2021). <https://doi.org/10.1609/icaps.v31i1.15954>
19. De Giacomo, G., Favorito, M.: Compositional approach to translate LTL_f/LDL_f into deterministic finite automata. In: Biundo, S., Do, M., Goldman, R., Katz, M., Yang, Q., Zhuo, H.H. (eds.) Proceedings of the 31'st International Conference on Automated Planning and Scheduling (ICAPS'21). pp. 122–130. AAAI Press (Aug 2021). <https://doi.org/10.1609/icaps.v31i1.15954>
20. De Giacomo, G., Favorito, M., Li, J., Vardi, M.Y., Xiao, S., Zhu, S.: LTL_f synthesis as AND-OR graph search: Knowledge compilation at work. In: Raedt, L.D. (ed.) Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI'22). pp. 2591–2598. International Joint Conferences on Artificial Intelligence Organization (Jul 2022). <https://doi.org/10.24963/ijcai.2022/359>

21. De Giacomo, G., Rubin, S.: Automata-theoretic foundations of fond planning for LTL_f/LDL_f goals. In: Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI'18). pp. 4729–4735 (2018). <https://doi.org/10.24963/ijcai.2018/657>
22. De Giacomo, G., Vardi, M.Y.: Linear temporal logic and linear dynamic logic on finite traces. In: Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI'13). pp. 854–860. IJCAI'13, AAAI Press (Aug 2013). <https://doi.org/10.5555/2540128.2540252>
23. De Giacomo, G., Vardi, M.Y.: Synthesis for LTL and LDL on finite traces. In: Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI'15). pp. 1558–1564. AAAI Press (2015). <https://doi.org/10.5555/2832415.2832466>
24. Duret-Lutz, A.: Supporting material for "Engineering an LTLf Synthesizer Tool" (2025). <https://doi.org/10.5281/zenodo.15752968>
25. Duret-Lutz, A., Renault, E., Colange, M., Renkin, F., Aisse, A.G., Schlehuber-Caissier, P., Medioni, T., Martin, A., Dubois, J., Gillard, C., Lauko, H.: From Spot 2.0 to Spot 2.10: What's new? In: Proceedings of the 34th International Conference on Computer Aided Verification (CAV'22). Lecture Notes in Computer Science, vol. 13372, pp. 174–187. Springer (Aug 2022). https://doi.org/10.1007/978-3-031-13188-2_9
26. Ehlers, R., Lafortune, S., Tripakis, S., Vardi, M.Y.: Supervisory control and reactive synthesis: a comparative introduction. *Discrete Event Dynamic Systems* **27**(2), 209–260 (2017). <https://doi.org/10.1007/s10626-015-0223-0>
27. Esparza, J., Křetínský, J., Sickert, S.: One theorem to rule them all: A unified translation of LTL into ω -automata. In: Dawar, A., Grädel, E. (eds.) Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS'18). pp. 384–393. ACM (2018). <https://doi.org/10.1145/3209108.3209161>
28. Favorito, M.: Forward LTLf synthesis: DPLL at work. In: Benedictis, R.D., Castiglioni, M., Ferraioli, D., Malvone, V., Maratea, M., Scala, E., Serafini, L., Serina, I., Tosello, E., Umbrico, A., Vallati, M. (eds.) Proceedings of the 30th Workshop on Experimental evaluation of algorithms for solving problems with combinatorial explosion (RCRA'23). CEUR Workshop Proceedings, vol. 3585 (2023). https://ceur-ws.org/Vol-3585/paper7_RCRA4.pdf
29. Favorito, M., Zhu, S.: LydiaSyft: A compositional symbolic synthesis framework for LTLf specifications. In: Proceedings of the 31st International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'25). Lecture Notes in Computer Science, vol. 15696, pp. 295–302. Springer (May 2025). https://doi.org/10.1007/978-3-031-90643-5_15
30. Finkbeiner, B.: Synthesis of reactive systems. In: Javier Esparza, Orna Grumberg, S.S. (ed.) Dependable Software Systems Engineering, NATO Science for Peace and Security Series — D: Information and Communication Security, vol. 45, pp. 72–98. IOS Press (2016). <https://doi.org/10.3233/978-1-61499-627-9-72>
31. Finkbeiner, B., Geier, G., Passing, N.: Specification decomposition for reactive synthesis. In: Proceedings for the 13th NASA Formal Methods Symposium (NFM'21). Lecture Notes in Computer Science, vol. 12673, pp. 113–130. Springer (2021). https://doi.org/10.1007/978-3-030-76384-8_8
32. Fujita, M., McGeer, P.C., Yang, J.C.: Multi-terminal binary decision diagrams: An efficient data structure for matrix representation. *Formal Methods in System Design* **10**(2/3), 149–169 (1997). <https://doi.org/10.1023/A:1008647823331>

33. Gabbay, D., Pnueli, A., Shelah, S., Stavi, J.: On the temporal analysis of fairness. In: Proceedings of the 7th ACM SIGPLAN-SIGACT symposium on Principles of programming languages (POPL'80). pp. 163–173. Association for Computing Machinery (1980). <https://doi.org/10.1145/567446.5674>
34. Gerevini, A., Haslum, P., Long, D., Saetti, A., Dimopoulos, Y.: Deterministic planning in the fifth international planning competition: PDDL3 and experimental evaluation of the planners. *Artificial Intelligence* **173**(5–6), 619–668 (2009). <https://doi.org/10.1016/j.artint.2008.10.012>
35. Grädel, E.: Finite model theory and descriptive complexity. In: Finite Model Theory and Its Applications, chap. 3, pp. 125–230. Texts in Theoretical Computer Science an EATCS Series, Springer Berlin Heidelberg, Berlin, Heidelberg (2007). https://doi.org/10.1007/3-540-68804-8_3
36. Henriksen, J.G., Jensen, J., Jørgensen, M., Klarlund, N., Paige, R., Rauhe, T., Sandholm, A.: Mona: Monadic second-order logic in practice. In: Brinksma, E., Cleaveland, W.R., Larsen, K.G., Margaria, T., Steffen, B. (eds.) First International Workshop on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'95). pp. 89–110. Springer Berlin Heidelberg (1995). https://doi.org/10.1007/3-540-60630-0_5
37. Jacobs, S., Perez, G.A., Schlehuber-Caissier, P.: The temporal logic synthesis format TLSF v1.2. arXiv (2023). <https://doi.org/10.48550/arXiv.2303.03839>
38. Jacobs, S., Perez, G.A., Abraham, R., Bruyère, V., Cadilhac, M., Colange, M., Delfosse, C., van Dijk, T., Duret-Lutz, A., Faymonville, P., Finkbeiner, B., Khalimov, A., Klein, F., Luttenberger, M., Meyer, K.J., Michaud, T., Pommellet, A., Renkin, F., Schlehuber-Caissier, P., Sakr, M., Sickert, S., Staquet, G., Tamines, C., Tentrup, L., Walker, A.: The reactive synthesis competition (SYNTCOMP): 2018–2021. arXiv (Jun 2022). <https://doi.org/10.48550/ARXIV.2206.00251>
39. Klarlund, N., Møller, A.: MONA version 1.4, user manual. Tech. rep., BRICS (Jul 2001), <https://www.brics.dk/mona/mona14.pdf>
40. Klarlund, N., Rauhe, T.: BDD algorithms and cache misses. Tech. Rep. BR-96-26, BRICS (Jul 1996), <https://www.brics.dk/mona/papers/bdd-alg-cache-miss/article.pdf>
41. Kluyver, T., Ragan-Kelley, B., Pérez, F., Granger, B., Bussonnier, M., Frederic, J., Kelley, K., Hamrick, J., Grout, J., Corlay, S., Ivanov, P., Avila, D., Abdalla, S., Willing, C., development team, J.: Jupyter notebooks — a publishing format for reproducible computational workflows. In: Loizides, F., Schmidt, B. (eds.) Proceedings of 20th International Conference on Electronic Publishing: Positioning and Power in Academic Publishing: Players, Agents and Agendas (ELPUB'16). pp. 87–90. IOS Press (2016). <https://doi.org/10.3233/978-1-61499-649-1-87>
42. Lind-Nielsen, J.: BuDDy: A binary decision diagram package. User's manual. (1999), <https://web.archive.org/web/20040402015529/http://www.itu.dk/research/buddy/>
43. Liu, X., Smolka, S.A.: Simple linear-time algorithms for minimal fixed points. In: Larsen, K.G., Skyum, S., Winskel, G. (eds.) Proceedings of the 25th International Colloquium on Automata, Languages and Programming (ICALP'98). pp. 53–66. Springer Berlin Heidelberg (1998). <https://doi.org/10.1007/BFb0055035>
44. Long, D.: BDD library. source archive, <https://www.cs.cmu.edu/~modelcheck/bdd.html>
45. Minato, S.i.: Representation of Multi-Valued Functions, pp. 39–47. Springer US, Boston, MA (1996). https://doi.org/10.1007/978-1-4613-1303-8_4

46. Piterman, N., Pnueli, A., Sa'ar, Y.: Synthesis of reactive(1) designs. In: Proceedings of the 7th international conference on Verification, Model Checking, and Abstract Interpretation (VMCAI'06). Lecture Notes in Computer Science, vol. 3855, pp. 364–380. Springer (2006). https://doi.org/10.1007/11609773_24
47. Pnueli, A., Rosner, R.: On the synthesis of a reactive module. In: Proceedings of the 16th ACM SIGPLAN-SIGACT symposium on Principles of Programming Languages (POPL'89). Association for Computing Machinery (1989). <https://doi.org/10.1145/75277.75293>
48. Renkin, F., Schlehuber-Caissier, P., Duret-Lutz, A., Pommellet, A.: Effective reductions of Mealy machines. In: Proceedings of the 42nd International Conference on Formal Techniques for Distributed Objects, Components, and Systems (FORTE'22). Lecture Notes in Computer Science, vol. 13273, pp. 170–187. Springer (Jun 2022). https://doi.org/10.1007/978-3-031-08679-3_8
49. Renkin, F., Schlehuber-Caissier, P., Duret-Lutz, A., Pommellet, A.: Dissecting `ltlsynt`. Formal Methods in System Design (2023). <https://doi.org/10.1007/s10703-022-00407-6>
50. Sickert, S., Meyer, P.: Modernizing strix (2021), <https://www7.in.tum.de/~sickert/publications/MeyerS21.pdf>
51. Somenzi, F.: CUDD: CU Decision Diagram package release 3.0.0 (Dec 2015), <https://web.archive.org/web/20171208230728/http://vlsi.colorado.edu/~fabio/CUDD/cudd.pdf>
52. Tabajara, L.M., Vardi, M.Y.: Partitioning techniques in LTLf synthesis. In: Kraus, S. (ed.) Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI'19). pp. 5599–5606. ijcai.org (Aug 2019). <https://doi.org/10.24963/IJCAI.2019/777>
53. Xiao, S., Li, J., Zhu, S., Shi, Y., Pu, G., Vardi, M.: On-the-fly synthesis for LTL over finite traces. In: Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI'21, Technical Track 7). pp. 6530–6537 (May 2021). <https://doi.org/10.1609/aaai.v35i7.16809>
54. Xiao, S., Li, Y., Huang, X., Xu, Y., Li, J., Pu, G., Strichman, O., Vardi, M.Y.: Model-guided synthesis for LTL over finite traces. In: Proceedings of the 25th International Conference on Verification, Model Checking, and Abstract Interpretation. Lecture Notes in Computer Science, vol. 14499, pp. 186–207. Springer (2024). https://doi.org/10.1007/978-3-031-50524-9_9
55. Zhu, S., De Giacomo, G.: Synthesis of maximally permissive strategies for LTLf specifications. In: Raedt, L.D. (ed.) Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI'22). pp. 2783–2789. ijcai.org (Jul 2022). <https://doi.org/10.24963/IJCAI.2022/386>
56. Zhu, S., Tabajara, L.M., Li, J., Pu, G., Vardi, M.Y.: Symbolic LTLf synthesis. In: Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI'17). pp. 1362–1369 (2017). <https://doi.org/10.24963/ijcai.2017/189>

These appendices and the margin notes that point to them were part of the submission for interested reviewers, but they have not been peer-reviewed, and are not part of the CIAA'25 proceedings.

A MTBDD operations

This section details some of the MTBDD operations described in section 2.4. We believe those functions should appear straightforward to any reader familiar with BDD implementations. We show them for the sake of being comprehensive.

A.1 Evaluating an MTBDD using an assignment

For $m \in \text{MTBDD}(\mathcal{P}, \mathcal{S})$, and $w \in \mathbb{B}^{\mathcal{P}}$, Algorithm 3 shows how to compute $m(w)$ by descending the structure of m according to w .

```

Function eval( $m, w$ )
  input   :  $m \in \text{MTBDD}(\mathcal{P}, \mathcal{S})$ ,  $w \in \mathbb{B}^{\mathcal{P}}$ 
  output :  $m(w) \in \mathcal{S}$ 
  ( $p, \ell, h$ )  $\leftarrow m$ ;
  while  $p \neq \infty$  do
    if  $w(v)$  then ( $p, \ell, h$ )  $\leftarrow h$  else ( $p, \ell, h$ )  $\leftarrow \ell$ ;
  return  $\ell$ ;

```

Algorithm 3: Evaluating an MTBDD using an assignment.

A.2 Binary and Unary Operations on MTBDDs

Algorithm 4 shows how the implementation of `apply2` follows a classical recursive definition typically found in BDD packages [11,32,1]. The function `makebdd` is in charge of ensuring the *reduced* property of the MTBDD: for any triplet of the form (p, r, r) where the *low* and *high* links are equal, `makebdd` returns r to skip over the node. For other triplets, `makebdd` will look up and possibly update a global hash table to ensure that each triplet is represented only once. The hash table H is used for memoization; assuming lossless caching (i.e., no dropped entry on hash collision), this ensures that the number of recursive calls performed is at most $|m_1| \cdot |m_2|$. Our implementation, as discussed in Section F, uses a lossy cache, therefore the complexity might be higher.

An `apply1` function can be written along the same lines for unary operators.

A.3 Leaves of an MTBDD

Function `leaves( $m$ )`, shown by Algorithm 5 is a straightforward way to collect the leaves that appear in an MTBDD m .

```

Function apply2( $m_1, m_2, \odot, H$ )
  input :  $m_1 \in \text{MTBDD}(\mathcal{P}, \mathcal{S}_1)$ ,  $m_2 \in \text{MTBDD}(\mathcal{P}, \mathcal{S}_2)$ ,  $\odot : \mathcal{S}_1 \times \mathcal{S}_2 \rightarrow \mathcal{S}_3$ ,
          $H : \text{hashmap}$ 
  output :  $m_1 \odot m_2 \in \text{MTBDD}(\mathcal{P}, \mathcal{S}_3)$ 
  if ( $m_1, m_2, \odot$ )  $\in H$  then
     $\lfloor$  return  $H[(m_1, m_2, \odot)]$  ;
    ( $p_1, \ell_1, h_1$ )  $\leftarrow m_1$ ;
    ( $p_2, \ell_2, h_2$ )  $\leftarrow m_2$ ;
    if  $p_1 < p_2$  then
       $r \leftarrow \text{makebdd}(p_1, \text{apply2}(\ell_1, m_2, \odot, H), \text{apply2}(h_1, m_2, \odot, H))$ ;
    else if  $p_2 < p_1$  then
       $r \leftarrow \text{makebdd}(p_2, \text{apply2}(m_1, \ell_2, \odot, H), \text{apply2}(m_1, h_2, \odot, H))$ ;
    else if  $p_1 < \infty$  then //  $p_1 = p_2$ 
       $r \leftarrow \text{makebdd}(p_1, \text{apply2}(\ell_1, \ell_2, \odot, H), \text{apply2}(h_1, h_2, \odot, H))$ ;
    else //  $p_1 = p_2 = \infty$ , we have terminals holding values  $\ell_1$  and  $\ell_2$ 
       $r \leftarrow \text{makebdd}(\infty, \ell_1 \odot \ell_2, \infty)$ ;
     $H[(m_1, m_2, \odot)] \leftarrow r$ ;
  return  $r$ ;

```

Algorithm 4: Composing two MTBDDs by applying a binary operator to their terminals.

```

Function leaves( $m$ )
  input :  $m \in \text{MTBDD}(\mathcal{P}, \mathcal{S})$ 
  output : the subset of  $\mathcal{S}$  that appears on leaves of  $m$ 
   $seen \leftarrow \{m\}$ ;
   $todo \leftarrow \{m\}$ ;
   $res \leftarrow \emptyset$ ;
  while  $todo \neq \emptyset$  do
     $m \leftarrow todo.\text{pop\_any}()$ ;
    ( $p, \ell, h$ )  $\leftarrow m$ ;
    if  $p = \infty$  then // We reached a leaf labeled by  $\ell$ 
       $res \leftarrow res \cup \{\ell\}$ ;
    else
       $todo \leftarrow todo \cup (\{\ell, r\} \setminus seen)$ ;
       $seen \leftarrow seen \cup \{\ell, r\}$ ;
  return  $res$ ;

```

Algorithm 5: Gathering the leaves of an MTBDD can be done with a simple linear traversal of an MTBDD.

A.4 Boolean Operations with Shortcuts

Algorithm 6 shows how to implement Boolean operations on MTBDDs with terminals in $\mathcal{S} \cup \mathbb{B}$, shortcutting the recursion when one of the operands is a terminal labeled by a value in $\mathbb{B}$.

```

Function apply2sc( $m_1, m_2, \odot, H$ )
  input :  $m_1 \in \text{MTBDD}(\mathcal{P}, \mathcal{S}_1 \cup \mathbb{B}), m_2 \in \text{MTBDD}(\mathcal{P}, \mathcal{S}_2 \cup \mathbb{B}),$ 
            $\odot : \mathcal{S}_1 \times \mathcal{S}_2 \rightarrow \mathcal{S}_3, H : \text{hashmap}$ 
  output :  $m_1 \odot m_2 \in \text{MTBDD}(\mathcal{P}, \mathcal{S}_3)$ 

  if ( $m_1, m_2, \odot$ )  $\in H$  then
    return  $H[(m_1, m_2, \odot)]$ ;
  ( $p_1, \ell_1, h_1$ )  $\leftarrow m_1$ ;
  ( $p_2, \ell_2, h_2$ )  $\leftarrow m_2$ ;
  if ( $p_1 = \infty \wedge \ell_1 \in \mathbb{B}$ )  $\vee$  ( $p_2 = \infty \wedge \ell_2 \in \mathbb{B}$ ) then
    switch  $\odot$  do
      case  $\wedge$  do
        if  $\perp \in \{\ell_1, \ell_2\}$  then return  $(\infty, \perp, \infty)$ ;
        if  $\ell_1 = \top$  then return  $m_2$ ;
        if  $\ell_2 = \top$  then return  $m_1$ ;
      case  $\vee$  do
        if  $\top \in \{\ell_1, \ell_2\}$  then return  $(\infty, \top, \infty)$ ;
        if  $\ell_1 = \perp$  then return  $m_2$ ;
        if  $\ell_2 = \perp$  then return  $m_1$ ;
      case ... do ...;
  if  $p_1 < p_2$  then
     $r \leftarrow \text{makebdd}(p_1, \text{apply2sc}(\ell_1, m_2, \odot, H),$ 
       $\text{apply2sc}(h_1, m_2, \odot, H))$ ;
  else if  $p_2 < p_1$  then
     $r \leftarrow \text{makebdd}(p_2, \text{apply2sc}(m_1, \ell_2, \odot, H),$ 
       $\text{apply2sc}(m_1, h_2, \odot, H))$ ;
  else if  $p_1 < \infty$  then //  $p_1 = p_2$ 
     $r \leftarrow \text{makebdd}(p_1, \text{apply2sc}(\ell_1, \ell_2, \odot, H), \text{apply2sc}(h_1, h_2, \odot, H))$ ;
  else //  $p_1 = p_2 = \infty$ , we have terminals holding values  $\ell_1$  and  $\ell_2$ 
     $r \leftarrow \text{makebdd}(\infty, \ell_1 \odot \ell_2, \infty)$ ;
     $H[(m_1, m_2, \odot)] \leftarrow r$ ;
  return  $r$ ;

```

Algorithm 6: Variant of `apply2` that implements shortcuts when one of the argument is a Boolean leaf.

B Boolean Operations on MTDFAs

Although it is not necessary for the approach we presented, our implementation supports all Boolean operations over MTDFAs.

Since $\Delta(q) \in \text{MTBDD}(\mathcal{P}, \mathcal{Q} \times \mathbb{B})$ has terminals labeled by pairs of the form $(q, b) \in \mathcal{Q} \times \mathbb{B}$, let us extend any Boolean operator $\odot : \mathbb{B} \times \mathbb{B} \rightarrow \mathbb{B}$ so that it can work on such pairs. More formally, for $(q_1, b_1) \in \mathcal{Q}_1 \times \mathbb{B}$ and $(q_2, b_2) \in \mathcal{Q}_2 \times \mathbb{B}$ we define $(q_1, b_1) \odot (q_2, b_2)$ to be equal to $((q_1, q_2), (b_1 \odot b_2)) \in ((\mathcal{Q}_1 \times \mathcal{Q}_2) \times \mathbb{B})$. Using Algorithm 4 to apply $\odot$ elements of $\text{MTBDD}(\mathcal{P}, \mathcal{Q} \times \mathbb{B})$ gives us a very simple way to combine MTDFAs, as shown by the following definition.

Definition 7 (Composition of two MTDFAs). Let $\mathcal{A}_1 = \langle \mathcal{Q}_1, \mathcal{P}, \iota_1, \Delta_1 \rangle$ and $\mathcal{A}_2 = \langle \mathcal{Q}_2, \mathcal{P}, \iota_2, \Delta_2 \rangle$ be two MTDFAs over the same variables $\mathcal{P}$, and let $\odot \in \{\wedge, \vee, \rightarrow, \leftrightarrow, \dots\}$ be any Boolean binary operator.

Then, let $\mathcal{A}_1 \odot \mathcal{A}_2$ denote the composition of $\mathcal{A}_1$ and $\mathcal{A}_2$ defined as the MTDFA $\langle \mathcal{Q}_1 \times \mathcal{Q}_2, \mathcal{P}, (\iota_1, \iota_2), \Delta \rangle$ where for any $(q_1, q_2) \in \mathcal{Q}_1 \times \mathcal{Q}_2$ we have $\Delta_3((q_1, q_2)) = \Delta_1(q_1) \odot \Delta_2(q_2)$.

Property 1. With the notations from Definition 7, $\mathcal{L}(\mathcal{A}_1 \odot \mathcal{A}_2) = \{\sigma \in (\mathbb{B}^{\mathcal{P}})^+ \mid (\sigma \in \mathcal{L}(\mathcal{A}_1)) \odot (\sigma \in \mathcal{L}(\mathcal{A}_2))\}$. In particular $\mathcal{L}(\mathcal{A}_1 \wedge \mathcal{A}_2) = \mathcal{L}(\mathcal{A}_1) \cap \mathcal{L}(\mathcal{A}_2)$ and $\mathcal{L}(\mathcal{A}_1 \vee \mathcal{A}_2) = \mathcal{L}(\mathcal{A}_1) \cup \mathcal{L}(\mathcal{A}_2)$. If $\oplus$ designates the *exclusive or* operator, testing the equivalence of two automata $\mathcal{L}(\mathcal{A}_1) = \mathcal{L}(\mathcal{A}_2)$ amounts to testing whether $\mathcal{L}(\mathcal{A}_1 \oplus \mathcal{A}_2) = \emptyset$.

The complementation of an MTDFA (with respect to $(\mathbb{B}^{\mathcal{P}})^+$ not $(\mathbb{B}^{\mathcal{P}})^*$) can be defined using the unary Boolean negation similarly.

Such compositional operations are at the heart of the compositional LTL_f translations used by Lisa [8], Lydia [19] and LydiaSyft [29]. This is efficient as it allows minimizing intermediate automata before combining them. Our translator tool [ltlf2dfa](#) uses such a compositional approach by default. For LTL_f synthesis, our tool [ltlfsynt](#) also has the option to build the automaton by composition, but this is not enabled by default: using an on-the-fly construction as presented in Algorithm 2 is more efficient. We refer the reader to the artifact [24] for benchmark comparisons involving our own implementation of the compositional approach.

C Simplified MTDFA

Figure 4 shows a simplified version of the MTDFA from Figure 1 that can be obtained by any one of two optimizations discussed in Section 3:

- merge states with identical MTBDD representations, or
- apply the $(G\beta) \wedge \beta \rightsquigarrow G\beta$ simplification during construction.

The second optimization is faster, as it does not require computing $\text{tr}(q)$ on some state q only to later find that the result is identical to some previous $\text{tr}(q')$.

This simplified automaton may also help understand the “transition-based” nature of those MTDFAs. Here we have pairs of terminal with identical formula labels, but different acceptance: words are allowed to finish on one, but not the other. If they continue, they continue from the state specified by the formula. Figure 5 shows an equivalent “transition-based DFA” using notations that should be more readable by readers familiar with finite automata.

D Try it Online!

The [Spot Sandbox](#) website offers online access to the development version of Spot (which includes the work described here) via Jupyter notebooks [41] or shell terminals.




Fig. 4. The MTDFA from Figure 1, simplified by merging all states that have an identical MTBDD successor and adjusting the terminals.




Fig. 5. Transition-based DFA interpretation of the MTDFA of Figure 4. A word $\sigma \in (\mathcal{P}^{\mathbb{B}})^+$ is accepted if there is a run of the automaton such that each assignment $\sigma(i)$ is compatible with the Boolean formula labeling the transition, and if the last transition visited was accepting (double line). The sink state corresponding to ff has been trimmed for clarity.

In order to try the [ltlf2dfa](#) and [ltlfsynt](#) command-line tools, simply connect to [Spot Sandbox](#), hit the “New” button, and start a “Terminal”.

The example directory contains two Jupyter notebooks directly related to this submission:

- [backprop.ipynb](#) illustrates Algorithm 1. There, players 0 and 1 are called **True** and **False** respectively.
- [ltlf2dfa.ipynb](#) illustrates the translation of Section 3 with the optimizations discussed in Section 3 (page 3), the MTDFA operations mentioned in Appendix B, and some other game solving techniques not discussed here.

An HTML version of these two notebooks can also be found in directory `more-examples/` of the associated artifact [24].

E Backpropagation of Losing Vertices

The example of Figure 2 does not make it very clear how marking ff as a losing vertex (i.e., winning for 1) may improve the on-the-fly game solving: it does not help in that example.




Fig. 6. If this game is created on-the-fly, it is useful to mark the ff terminal as losing. When $\Delta(\iota)$ is encoded into a game, the fact that ff is winning for player I will cause the two round vertices above it to be immediately marked as winning as well. Now, since the initial vertex is known to be winning for I, hence losing for O, the exploration may stop without having to encode $\Delta(\alpha)$.

Figure 6 shows a scenario where marking states as losing and propagating this information is useful to avoid some unnecessary exploration of a large part of the automaton. Algorithm 2, described in Section 4, translates one state of the MTDFA at a time, starting from ι , and encodes that state into a game by calling `new_vertex`, `new_edge`, etc. In the example of Figure 6, after the MTBDD for $\Delta(\iota)$ has been encoded (the top five nodes of Figure 6), the initial node will be marked as winning for I already (because I can select the appropriate value of i_0 and i_1 to reach ff), therefore, the algorithm can stop immediately. Had we decided to backpropagate only the states winning for player O, the algorithm would have to continue encoding $\Delta(\alpha)$ into the game and probably many other states reachable from there. At the end of the backpropagation, the initial node would still be undetermined, and we would also conclude that O cannot win.

Such an interruption of the on-the-fly exploration is used, does not only occur when the initial state is determined for the initial state, but at every search: if during the encoding of $\Delta(\alpha)$ we find that the winning status of the root node of $\Delta(\alpha)$ is determined (line 31 of Algorithm 2), then it is unnecessary to explore the rejecting leaves of $\Delta(\alpha)$.

F Implementation Details: MTBDDs in BuDDy

BuDDy [42] is a BDD library created by Jørn Lind-Nielsen for his Ph.D. project. Maintenance was passed to someone else in 2004. The Spot developer has contributed a few changes and fixes to the “original” project, but it soon became apparent that some of the changes motivated by Spot’s needs could not be merged upstream (e.g., because they would break other projects for the sake of efficiency). Nowadays, Spot is distributed with its own fork of BuDDy that includes several extra functions, a more compact representation of the BDD nodes (16 bytes per node instead of 20), a “derecursed” implementation of the most common BDD operations. Moving away from BuDDy, to another BDD library

would be very challenging. Therefore, for this work, we modified BuDDy to add support for MTBDDs with `int`-valued terminals (our MTBDD implementation knows nothing about LTL_f).

Our implementation differs from Mona’s MTBDDs or CUDD’s ADDs in several ways. First, BuDDy is designed around a global unicity table, which stores reference counted BDDs. There is no notion of “BDD manager” as in Mona or CUDD that allows building independent BDDs. We introduced support for MTBDD directly into this table, by reserving the highest possible variable number to indicate a terminal (storing the terminal’s value in the `low` link, as suggested by our notation in this paper), and adding an extra `if` in the garbage collector so it correctly deals with those nodes. This change allows to mix MTBDD terminals with regular BDD terminals (false and true). Existing BDD function wills work as they have always done when a BDD does not use the new terminals. If multi-terminals are used, a new set of functions should be used.

In CUDD’s ADD implementation, the set of operations that can be passed to the equivalent of the `apply2` function (see Algorithm 4) is restricted to a fixed set of algebraic operations that have well defined semantics. In Mona and in our implementation, the user may pass an arbitrary function in order to interpret the terminals (which can only store an integer) and combine them. For instance, to implement the presented algorithm where terminal are supposed to be labeled by pairs $(\alpha, b) \in \text{LTL}_f(\mathcal{P}) \times \mathbb{B}$, we store b in the lower bit of the terminal’s value, and use the other bits as an index in an array that stores α . If we create a new formula while combining two terminals, we add the new formula to that array, and build the value of the newly formed terminal from the corresponding index in that array.

One issue with implementing MTBDD operations is how to implement the operation cache (the H argument of Algorithm 4) when the function to apply on the leaves is supplied by the user. Since the supplied function may depend on global variables, it is important that this operation cache can be reset by the user.

We implement those user-controlled operation caches using lossy hash tables similar to what are used internally by BuDDy for classical BDD operations. Algorithm 4, the line $H[(m_1, m_2, \odot)] \leftarrow r$ that saves the result of the last operation may actually erase the result of a previous operation that would have been hashed to the same index. Therefore, the efficiency of our MTBDD algorithms will depend on how many collisions they generate, and this in turn depends on the size allocated for this hash table: ideally H should have a size of the same order as the number of BDD nodes used in the MTBDD resulting from the operation. We use two empirical heuristics to estimate a size for H . For unary operations on MTDFAs, we set $|H| = |\mathcal{P}| \cdot |\mathcal{Q}|/2$, and for binary operations on MTDFAs (e.g., Def. 7), we set $|H| = |\mathcal{P}_1 \cup \mathcal{P}_2| \cdot |\mathcal{Q}_1| \cdot |\mathcal{Q}_2|/4$. For operations performed during the translation of LTL_f formulas to MTDFAs (Th. 1), we use a hash table that is 20% of the total number of nodes allocated by BuDDy, but we share it for all MTBDD operations performed during the translation.

Mona handles those caches differently: it also estimates an initial size for those caches (with different formulas [40]), but by default it will handle any collision by chaining, growing an overflow table to store collisions as needed. This difference probably contributes to the additional “out-of-memory” errors that Mona-based tools tend to show in our benchmarks.

G Simple Rewriting Rules

We use a specification decomposition technique based on [31]. We try to rewrite the input specification φ into a conjunction $\varphi = \bigwedge_i \varphi_i$, where each φ_i uses non-overlapping sets of outputs. Formula $\Psi = \Psi_1 \wedge \Psi_2$ from Example 1 is already in this form. However, in general, the specification may be more complex, like $G(\xi_0) \rightarrow \bigwedge_i \xi_i$. In such a case, we rewrite the formula as $\bigwedge_i (G(\xi_0) \rightarrow \xi_i)$ before partitioning the terms of this conjunction into groups that use overlapping sets of output variables. Such a rewriting, necessary to an effective decomposition, may introduce a lot of redundancy in the formula (in this example $G(\xi_0)$ is duplicated several times).

For this reason, we apply simple language-preserving rewritings on LTL_f formulas before attempting to translate them into an MTDFA. These rewritings undo some of the changes that had to be done earlier to look for possible decompositions. They are also performed when decomposition is disabled. More generally, the goal is to reduce the number of temporal operators, in order to reduce the number of MTBDD operations that need to be performed.

$$(\alpha \rightarrow \beta) \wedge (\alpha \rightarrow \gamma) \rightsquigarrow \alpha \rightarrow (\beta \wedge \gamma) \quad (1)$$

$$(\alpha \rightarrow \beta) \vee (\gamma \rightarrow \delta) \rightsquigarrow (\neg\alpha) \vee \beta \vee (\neg\gamma) \vee \delta \quad (2)$$

$$\bigwedge_i G(\alpha_i) \wedge \bigwedge_j GF(\beta_j) \rightsquigarrow G(\bigwedge_i \alpha_i \wedge F(\bigwedge_j \beta_j)) \quad (3)$$

$$\bigvee_i F(\alpha_i) \vee \bigvee_j FG(\beta_j) \rightsquigarrow F(\bigvee_i \alpha_i \vee G(\bigvee_j \beta_j)) \quad (4)$$

$$X\alpha \wedge X\beta \rightsquigarrow X(\alpha \wedge \beta) \quad (5)$$

$$X\alpha \vee X\beta \rightsquigarrow X(\alpha \vee \beta) \quad (6)$$

$$X^!\alpha \wedge X^!\beta \rightsquigarrow X^!(\alpha \wedge \beta) \quad (7)$$

$$X^!\alpha \vee X^!\beta \rightsquigarrow X^!(\alpha \vee \beta) \quad (8)$$

$$GF(\alpha) \rightsquigarrow GF(\alpha_r) \quad (9)$$

$$FG(\alpha) \rightsquigarrow GF(\alpha_r) \quad (10)$$

Equation (2) is the only equation that does not reduce the number of operators. However, our implementation automatically removes duplicate operands for n -ary operators such as $\wedge$ or $\vee$, so this is more likely to occur after this rewriting.

In LTL_f, formulas $\text{GF}(\alpha)$ and $\text{FG}(\alpha)$ are equivalent, and specify that α should hold on the last position of the word. Therefore, in (9)–(10), any temporal operators in α can be removed using the same rules as in Theorem 4 in Appendix H.

H One-step (Un)Realizability Checks

To test if an LTL_f formula φ is realizable or unrealizable in one-step, we can rewrite the formula into Boolean formulas φ_r or φ_u using one of the following theorems that follow from the LTL_f semantics.

Then testing whether a Boolean formula is (un)realizable can be achieved by representing that formula as a BDD, and then removing input/output variables by universal/existential quantification, in the order required by the selected semantics (Moore or Mealy).

Theorem 4 (One-step realizability [53, Th. 2]). *For $\varphi \in \text{LTL}_f(\mathcal{P})$, define φ_r inductively using the following rules:*

$$\begin{array}{llll} ff_r = ff & (X^! \alpha)_r = ff & (G\alpha)_r = \alpha_r & (\alpha R \beta)_r = \beta_r \\ tt_r = tt & (X\alpha)_r = tt & (F\alpha)_r = \alpha_r & (\alpha U \beta)_r = \beta_r \\ p_r = p & \text{for } p \in \mathcal{P} & (\neg\alpha)_r = \neg(\alpha_r) & (\alpha \odot \beta)_r = \alpha_r \odot \beta_r \end{array}$$

Where $\odot \in \{\wedge, \vee, \rightarrow, \leftrightarrow, \oplus\}$.

If the Boolean formula φ_r is realizable, then φ is realizable too.

Theorem 5 (One-step unrealizability [53, Th. 3]). *Consider a formula $\varphi \in \text{LTL}_f(\mathcal{P})$. To simplify the definition, we assume φ to be in negative normal form (i.e., negations have been pushed down the syntactic tree, and may only occur in front of variables, and operators $\rightarrow$, $\leftrightarrow$, $\oplus$ have been rewritten away). We define φ_u inductively as follows:*

$$\begin{array}{llllll} ff_u = ff & (X^! \alpha)_u = tt & (G\alpha)_u = \alpha_u & (\alpha R \beta)_u = \alpha_u \wedge \beta_u & (\alpha \wedge \beta)_u = \alpha_u \wedge \beta_u \\ tt_u = tt & (X\alpha)_u = tt & (F\alpha)_u = \alpha_u & (\alpha U \beta)_u = \alpha_u \vee \beta_u & (\alpha \vee \beta)_u = \alpha_u \vee \beta_u \end{array}$$

For any variable $p \in \mathcal{P}$, we have $p_u = p$ and $(\neg p)_u = \neg p$.

If the Boolean formula φ_u is not realizable, then φ is not realizable.

I More Benchmark Results

The SyntComp benchmarks contain specifications that can be partitioned in three groups:

game Those specifications describe two-players games. They have three sub-families [52]: *single counter*, *double counters*, and *nim*.

pattern Those specifications are scalable patterns built either from nesting U operators, or by making conjunctions of terms such as $G(v_i)$ or $F(v_j)$. [53]

random Those specifications are random conjunctions of LTL_f specifications [56].

Of these three sets, the *games* are the most challenging to solve. The *patterns* use each variable only once, so they can all be reduced to *tt* or *ff* by the preprocessing technique discussed in Section 5, or by the one-step (un)realizability checks. Since *random* specifications are built as a conjunction of subspecifications that often have nonintersecting variable sets, they can very often be decomposed into output-disjoint specifications that can be solved separately [31].

Table 1 shows how the different tools succeed in these different benchmarks.

Figure 7 compare the best configuration of `ltlfsynt` against Nike: there are no cases where Nike is faster. If we disable preprocessings and one-step (un)realizability in `ltlfsynt`, the comparison is more balanced, as shown in Figure 8. Note that we have kept one-step (un)realizability enabled in this comparison, because Nike uses it too. This is the reason why *pattern* benchmarks are solved instantaneously by both tools.

Tables 2, 3, and 4 look at the runtime of the tools on *game* benchmarks. Values highlighted in yellow are within 5% of the minimum value of each line.

Table 2 shows a family of specifications where preprocessings are useless, and using one-step (un)realizability slows things down.

Table 3 shows a family of specifications where one-step (un)realizability is what allows `ltlfsynt` to solve many more instance than other tools (even tools like Nike or LydiaSyft who also implement one-step (un)realizability). The suspicious behavior of Lydia/LidyaSyft/SyftMax cycling between timeouts, segmentation faults, and out-of-memory has been double-checked: this is really how they terminated.

Finally, Table 4 shows very impressive results by `ltlfsynt` on the challenging Nim family of benchmarks: the highest configuration that third-party tools are able to solve is `nim_04_01`, but `ltlfsynt` solves this instantaneously in all configurations, and can handle much larger instances.

A more detailed analysis of the benchmark results can be found in directory `ltlfsynt-analysis/` of the associated artifact [24].

Table 1. Count of different tool outputs grouped by different types of benchmarks. Status *false* and *true* indicate how many times a tool successfully managed to decide unrealizability (*false*) or realizability (*true*). The other status are error conditions: TIMEOUT (over 15 minutes), OOMEM (over 4GB), ABORT (aborted), SEGV (segmentation violation). The later two errors are likely caused by some incorrect handling of out-of-memory conditions.

type	tool	status						all
		false	true	TIMEOUT	OOMEM	ABORT	SEGV	
game	ltlfsynt:full-N	40	45	19	16			120
	ltlfsynt:bfs-nopre	40	47	4	29			120
	ltlfsynt:bfs-os-nopre	40	57	3	20			120
	ltlfsynt:bfs	42	47	4	27			120
	ltlfsynt:bfs-os	42	57	3	18			120
	SyftMax	7	29	23	45	1	15	120
	Lydia	7	28	24	45	1	15	120
	LydiaSyft	7	29	20	45	1	18	120
	Lisa	12	22	71	11		4	120
	MoGuSer	3	16	5	96			120
	Cynthia	14	26	35	45			120
	Nike	5	15	30	70			120
pattern	ltlfsynt:full-N	21	19					40
	ltlfsynt:bfs-nopre	21	19					40
	ltlfsynt:bfs-os-nopre	21	19					40
	ltlfsynt:bfs	21	19					40
	ltlfsynt:bfs-os	21	19					40
	SyftMax	20	17	1	2			40
	Lydia	21	17		2			40
	LydiaSyft	21	19					40
	Lisa	20	17	1	2			40
	MoGuSer	21	19					40
	Cynthia	21	19					40
	Nike	21	19					40
random	ltlfsynt:full-N	957	278	1	1			1237
	ltlfsynt:bfs-nopre	873	278	52	34			1237
	ltlfsynt:bfs-os-nopre	872	278	58	29			1237
	ltlfsynt:bfs	957	278	1	1			1237
	ltlfsynt:bfs-os	957	278	1	1			1237
	SyftMax	790	259	15	173			1237
	Lydia	803	266		103		65	1237
	LydiaSyft	790	273	12	162			1237
	Lisa	872	265	100				1237
	MoGuSer	829	257	118	33			1237
	Cynthia	192	239	247	559			1237
	Nike	898	278	61				1237
all	ltlfsynt:full-N	1018	342	20	17			1397
	ltlfsynt:bfs-nopre	934	344	56	63			1397
	ltlfsynt:bfs-os-nopre	933	354	61	49			1397
	ltlfsynt:bfs	1020	344	5	28			1397
	ltlfsynt:bfs-os	1020	354	4	19			1397
	SyftMax	817	305	39	220	1	15	1397
	Lydia	831	311	24	150	1	80	1397
	LydiaSyft	818	321	32	207	1	18	1397
	Lisa	904	304	172	13		4	1397
	MoGuSer	853	292	123	129			1397
	Cynthia	227	284	282	604			1397
	Nike	924	312	91	70			1397




Fig. 7. Scatter plots comparing time and memory usage of Nike against `ltlf synt`'s best configuration. Dots on the lines marked as T and M on the side represent timeouts or out-of-memory cases.




Fig. 8. Scatter plots comparing time and memory usage of Nike against `ltlf synt`'s on-the-fly construction but without preprocessing.

Table 2. Runtime of the different configurations on the Single Counter benchmark.

file	ltlfsynt					Lydia	SyftMax	LydiaSyft	Lisa	MoGuSer	Cynthia	Nike
	full-N	bfs	bfs-os	bfs-os-nopre	bfs-nopre							
counter_01	0.018	0.020	0.021	0.025	0.024	0.044	0.028	0.131	0.042	0.016	0.036	0.035
counter_02	0.018	0.022	0.022	0.026	0.025	0.075	0.033	0.155	0.125	0.041	0.498	0.058
counter_03	0.019	0.023	0.021	0.027	0.025	0.121	0.039	0.180	1.157	0.198	1.530	0.428
counter_04	0.021	0.023	0.025	0.029	0.027	0.214	0.053	0.211	0.086	0.757	6.514	10.745
counter_05	0.028	0.028	0.033	0.034	0.031	0.359	0.093	0.281	0.314	5.278	T.O.	450.869
counter_06	0.036	0.034	0.046	0.047	0.038	0.763	0.232	0.435	1.755	32.487	OOM	T.O.
counter_07	0.053	0.045	0.074	0.074	0.049	2.150	0.733	0.961	14.779	186.882	OOM	OOM
counter_08	0.085	0.073	0.126	0.128	0.078	7.658	2.735	2.840	624.710	OOM	T.O.	OOM
counter_09	0.167	0.134	0.256	0.258	0.138	30.908	11.346	11.016	T.O.	OOM	OOM	OOM
counter_10	0.359	0.277	0.556	0.549	0.257	129.377	53.268	54.602	T.O.	OOM	OOM	OOM
counter_11	0.801	0.609	1.238	1.231	0.562	587.784	248.118	246.556	T.O.	OOM	OOM	OOM
counter_12	2.016	1.338	2.876	2.791	1.298	T.O.	T.O.	T.O.	T.O.	OOM	OOM	OOM
counter_13	4.840	3.088	6.406	6.239	2.919	T.O.	T.O.	T.O.	T.O.	OOM	OOM	OOM
counter_14	12.697	7.220	14.581	13.881	6.706	T.O.	T.O.	T.O.	T.O.	OOM	OOM	OOM
counter_15	35.908	16.667	32.601	31.341	15.909	T.O.	T.O.	T.O.	T.O.	OOM	OOM	OOM
counter_16	118.579	42.600	78.935	76.412	40.700	T.O.	T.O.	T.O.	OOM	OOM	OOM	OOM
counter_17	468.023	121.415	199.486	198.288	115.179	T.O.	T.O.	T.O.	OOM	OOM	OOM	OOM
counter_18	OOM	380.706	OOM	OOM	363.738	SEGV	SEGV	SEGV	OOM	OOM	OOM	OOM

Table 3. Runtime of the different configurations on the Double Counters benchmark.

file	ltlfsynt					Lydia	SyftMax	LydiaSyft	Lisa	MoGuSer	Cynthia	Nike
	full-N	bfs	bfs-os	bfs-os-nopre	bfs-nopre							
counters_01	0.017	0.021	0.021	0.025	0.025	0.081	0.036	0.156	0.156	0.019	0.065	0.034
counters_02	0.021	0.022	0.023	0.025	0.027	0.221	0.063	0.218	0.077	0.040	0.224	0.051
counters_03	0.034	0.032	0.022	0.024	0.039	0.699	0.222	0.404	0.292	30.434	OOM	0.746
counters_04	0.085	0.067	0.023	0.028	0.069	4.098	1.626	1.755	2.908	0.459	OOM	45.052
counters_05	0.435	0.234	0.026	0.029	0.242	52.224	35.268	34.598	65.050	T.O.	OOM	OOM
counters_06	2.482	1.172	0.030	0.034	1.096	T.O.	869.787	855.488	T.O.	T.O.	OOM	OOM
counters_07	18.462	6.080	0.037	0.042	5.943	T.O.	T.O.	T.O.	T.O.	78.484	OOM	OOM
counters_08	240.522	32.519	0.051	0.058	32.070	SEGV	SEGV	SEGV	OOM	OOM	OOM	OOM
counters_09	OOM	200.552	0.086	0.101	196.024	SEGV	SEGV	SEGV	T.O.	OOM	OOM	OOM
counters_10	OOM	OOM	0.160	0.183	OOM	OOM	OOM	OOM	SEGV	OOM	OOM	OOM
counters_11	OOM	OOM	0.329	0.370	OOM	OOM	OOM	OOM	OOM	OOM	OOM	OOM
counters_12	OOM	OOM	0.732	0.781	OOM	T.O.	T.O.	T.O.	SEGV	OOM	OOM	OOM
counters_13	OOM	OOM	1.659	1.800	OOM	T.O.	T.O.	T.O.	OOM	OOM	OOM	OOM
counters_14	OOM	OOM	3.680	4.097	OOM	T.O.	T.O.	T.O.	T.O.	OOM	OOM	OOM
counters_15	OOM	OOM	8.427	9.236	OOM	SEGV	SEGV	SEGV	SEGV	OOM	OOM	OOM
counters_16	OOM	OOM	19.592	21.190	OOM	SEGV	SEGV	SEGV	T.O.	OOM	OOM	OOM
counters_17	OOM	OOM	44.924	48.138	OOM	SEGV	SEGV	SEGV	T.O.	OOM	OOM	OOM
counters_18	OOM	OOM	100.748	108.932	OOM	OOM	OOM	OOM	T.O.	OOM	OOM	OOM
counters_19	OOM	OOM	235.041	252.037	OOM	OOM	OOM	OOM	T.O.	OOM	OOM	OOM
counters_20	OOM	OOM	584.165	630.793	OOM	ABORT	ABORT	ABORT	T.O.	OOM	OOM	OOM

Table 4. Runtime of the different configurations on the Nim benchmark.

	lflsynt					Lydia	SyftMax	LydiaSyft	Lisa	MoGuSer	Cynthia	Nike	
file	full-N	bfs	bfs-os	bfs-os-nopre	bfs-nopre								
nim_01.01	0.017	0.021	0.020		0.025	0.025	0.090	0.036	0.205	0.037	0.036	0.026	0.044
nim_01.02	0.017	0.022	0.022		0.026	0.026	0.163	0.047	0.249	0.058	0.253	0.072	0.058
nim_01.03	0.019	0.022	0.023		0.029	0.029	0.294	0.087	0.344	0.091	3.223	0.097	0.754
nim_01.04	0.021	0.022	0.024		0.026	0.031	0.501	0.204	0.522	0.150	34.811	0.115	4.713
nim_01.05	0.022	0.024	0.025		0.029	0.028	2.666	2.076	2.506	0.256	T.O.	0.192	34.726
nim_01.06	0.023	0.025	0.028		0.029	0.029	11.469	10.559	11.116	0.440	OOM	0.252	214.023
nim_01.07	0.025	0.028	0.029		0.031	0.030	40.465	39.790	39.501	0.771	T.O.	0.503	T.O.
nim_01.08	0.028	0.028	0.029		0.032	0.029	125.995	118.861	118.659	1.351	OOM	0.769	T.O.
nim_01.09	0.034	0.028	0.030		0.034	0.032	423.505	424.027	389.467	2.237	OOM	0.920	T.O.
nim_01.10	0.031	0.030	0.032		0.033	0.032	57.491	53.651	56.621	T.O.	OOM	0.804	T.O.
nim_01.11	0.033	0.029	0.033		0.034	0.033	221.574	215.627	209.561	T.O.	OOM	1.573	T.O.
nim_01.12	0.035	0.031	0.035		0.033	0.033	408.379	419.961	420.114	T.O.	OOM	1.515	T.O.
nim_01.13	0.035	0.035	0.035		0.035	0.035	T.O.	T.O.	T.O.	T.O.	OOM	2.544	OOM
nim_01.14	0.038	0.034	0.037		0.038	0.035	T.O.	T.O.	T.O.	T.O.	OOM	3.293	OOM
nim_01.15	0.042	0.037	0.039		0.039	0.036	T.O.	T.O.	T.O.	T.O.	OOM	4.375	OOM
nim_01.16	0.046	0.035	0.040		0.041	0.036	T.O.	T.O.	T.O.	T.O.	OOM	5.513	OOM
nim_01.17	0.050	0.036	0.043		0.044	0.038	T.O.	T.O.	T.O.	T.O.	OOM	6.621	OOM
nim_01.18	0.049	0.040	0.043		0.043	0.040	T.O.	T.O.	SEGV	T.O.	OOM	8.320	OOM
nim_01.19	0.055	0.042	0.047		0.046	0.041	SEGV	SEGV	SEGV	T.O.	OOM	9.579	OOM
nim_01.20	0.061	0.043	0.050		0.049	0.042	OOM	OOM	OOM	T.O.	OOM	7.498	OOM
nim_02.01	0.018	0.021	0.023		0.026	0.026	0.248	0.082	0.301	0.086	0.929	0.252	0.148
nim_02.02	0.026	0.025	0.032		0.033	0.031	0.719	0.427	0.725	0.244	225.577	1.518	11.794
nim_02.03	0.029	0.031	0.037		0.041	0.035	9.951	9.302	9.760	0.868	T.O.	2.797	295.478
nim_02.04	0.039	0.037	0.051		0.053	0.040	61.611	60.511	61.550	3.399	OOM	16.925	T.O.
nim_02.05	0.053	0.047	0.073		0.072	0.055	T.O.	T.O.	T.O.	11.476	OOM	55.084	T.O.
nim_02.06	0.076	0.061	0.107		0.103	0.073	T.O.	T.O.	T.O.	32.721	OOM	23.316	T.O.
nim_02.07	0.119	0.095	0.164		0.164	0.105	T.O.	T.O.	SEGV	T.O.	OOM	95.698	T.O.
nim_02.08	0.157	0.141	0.263		0.230	0.128	SEGV	SEGV	SEGV	T.O.	OOM	723.682	T.O.
nim_02.09	0.258	0.216	0.393		0.344	0.182	SEGV	SEGV	SEGV	T.O.	OOM	427.258	T.O.
nim_02.10	0.433	0.318	0.546		0.516	0.271	SEGV	SEGV	SEGV	T.O.	OOM	T.O.	T.O.
nim_02.11	0.598	0.500	0.817		0.803	0.429	SEGV	SEGV	SEGV	T.O.	OOM	T.O.	OOM
nim_02.12	0.932	0.642	1.186		1.103	0.602	OOM	OOM	OOM	T.O.	OOM	696.837	OOM
nim_02.13	1.146	0.888	1.653		1.652	0.902	OOM	OOM	OOM	T.O.	OOM	T.O.	OOM
nim_02.14	1.573	1.242	2.300		2.131	1.268	OOM	OOM	OOM	T.O.	OOM	T.O.	OOM
nim_02.15	2.237	1.696	3.079		2.825	1.945	OOM	OOM	OOM	T.O.	OOM	T.O.	OOM
nim_02.16	3.254	2.077	3.838		3.595	2.146	OOM	OOM	OOM	T.O.	OOM	T.O.	OOM
nim_02.17	4.299	3.151	4.908		4.547	2.996	OOM	OOM	OOM	T.O.	OOM	T.O.	OOM
nim_02.18	5.569	3.519	6.250		5.811	3.551	OOM	OOM	OOM	T.O.	OOM	T.O.	OOM
nim_02.19	7.423	4.922	8.054		7.496	5.053	OOM	OOM	OOM	T.O.	OOM	T.O.	OOM
nim_02.20	9.513	6.039	9.816		9.084	6.064	OOM	OOM	OOM	T.O.	OOM	T.O.	OOM
nim_03.01	0.023	0.028	0.030		0.034	0.036	1.038	0.770	1.070	0.395	380.086	1.438	10.598
nim_03.02	0.061	0.059	0.067		0.072	0.056	70.922	70.258	68.449	5.258	OOM	32.446	T.O.
nim_03.03	0.190	0.161	0.204		0.201	0.140	T.O.	T.O.	T.O.	52.381	OOM	T.O.	T.O.
nim_03.04	0.736	0.515	0.595		0.621	0.476	T.O.	T.O.	SEGV	OOM	OOM	T.O.	T.O.
nim_03.05	1.951	1.480	1.801		1.685	1.410	SEGV	SEGV	SEGV	T.O.	OOM	T.O.	T.O.
nim_03.06	5.778	4.089	4.636		4.584	4.012	OOM	OOM	OOM	T.O.	OOM	T.O.	OOM
nim_03.07	16.955	9.713	10.808		10.850	9.725	OOM	OOM	OOM	T.O.	OOM	T.O.	OOM
nim_03.08	46.546	25.346	26.402		26.030	24.768	OOM	OOM	OOM	T.O.	OOM	T.O.	OOM
nim_03.09	126.784	64.669	64.223		63.547	66.692	OOM	OOM	OOM	T.O.	OOM	T.O.	OOM
nim_03.10	348.848	169.718	168.065		167.389	173.443	OOM	OOM	OOM	T.O.	OOM	T.O.	OOM
nim_03.11	T.O.	423.880	413.958		411.177	426.671	OOM	OOM	OOM	T.O.	OOM	OOM	OOM
nim_04.01	0.038	0.042	0.043		0.084	0.072	53.638	54.005	51.963	6.454	OOM	29.272	152.136
nim_04.02	0.941	0.712	0.791		0.783	0.727	T.O.	T.O.	T.O.	224.130	OOM	T.O.	T.O.
nim_04.03	9.519	6.601	7.119		7.015	6.694	SEGV	SEGV	SEGV	OOM	OOM	T.O.	T.O.
nim_04.04	104.648	58.148	60.060		60.705	58.543	OOM	OOM	OOM	T.O.	OOM	OOM	T.O.
nim_04.05	T.O.	547.200	554.162		555.808	549.473	OOM	OOM	OOM	T.O.	OOM	OOM	OOM
nim_05.01	0.082	0.079	0.094		0.678	0.666	T.O.	T.O.	T.O.	153.449	OOM	455.075	T.O.
nim_05.02	46.786	27.694	28.546		28.013	27.628	SEGV	SEGV	SEGV	OOM	OOM	T.O.	T.O.
nim_06.01	0.274	0.301	0.305		9.038	8.997	SEGV	SEGV	SEGV	OOM	OOM	T.O.	T.O.
nim_07.01	1.433	1.487	1.521		408.231	402.286	OOM	OOM	OOM	OOM	OOM	OOM	T.O.
nim_08.01	7.288	6.986	7.326		OOM	OOM	OOM	OOM	OOM	T.O.	OOM	OOM	T.O.
nim_09.01	38.163	35.055	36.017		OOM	OOM	OOM	OOM	OOM	T.O.	OOM	OOM	T.O.