# Satisfiability checking for Mission-time LTL (MLTL)

Jianwen Li [a,*], Moshe Y. Vardi [b], Kristin Y. Rozier [c,*]

[a] *Shanghai Key Laboratory of Trustworthy Computing, East China Normal University,* [1] *Shanghai, China*
[b] *Rice University, Houston, TX, United States of America*
[c] *Iowa State University, Ames, IA, United States of America*

A R T I C L E   I N F O

A B S T R A C T

Mission-time Linear Temporal Logic (LTL), abbreviated as MLTL, is a bounded variant of Metric Temporal Logic (MTL) over naturals designed to generically specify requirements for mission-based system operation common to aircraft, spacecraft, vehicles, and robots. Despite the utility of MLTL as a specification logic, major gaps remain in analyzing MLTL, e.g., for specification debugging or model checking, centering on the absence of any complete MLTL satisfiability checker. In this paper, we explore both the theoretical and algorithmic problems of MLTL satisfiability checking. We prove that the MLTL satisfiability checking problem is NEXPTIME-complete and that satisfiability checking $MLTL_0$, the variant of MLTL where all intervals start at 0, is PSPACE-complete. To explore the best algorithmic solution for MLTL satisfiability checking, we reduce this problem to LTL satisfiability checking, $LTL_f$ (LTL over finite traces) satisfiability checking, and model checking respectively, thus conducting translations for MLTL-to-LTL, MLTL-to-$LTL_f$, and MLTL-to-SMV. Moreover, we propose a new SMT-based solution for MLTL satisfiability checking and create a translation for MLTL-to-SMT. Our extensive experimental evaluation shows that while the MLTL-to-SMV translation with NuXmv model checker performs best on the benchmarks whose interval ranges are small (than 100), the MLTL-to-SMT translation with the Z3 SMT solver offers the most scalable performance.

© 2022 Elsevier Inc. All rights reserved.

## 1. Introduction

Mission-time Linear Temporal Logic (LTL), abbreviated as MLTL [1], has the syntax of Linear Temporal Logic with the option of integer bounds on the temporal operators. It was created as a generalization of the variations [2–4] on finitely-bounded linear temporal logic, ideal for specification of missions carried out by aircraft, spacecraft, rovers, and other vehicular or robotic systems. For example, $\Box_{[0,10]} p$ (Globally $p$) indicates that $p$ has to be true at each time point from 0 to 10, while $\Diamond_{[0,10]} p$ (Eventually $p$) means $p$ has to be true at some point from 0 to 10. MLTL provides the readability of LTL [5], while assuming, when a different duration is not specified, that all requirements must be upheld during the (a priori known) length of a given mission, such as during the half-hour battery life of an Unmanned Aerial System (UAS). Using integer bounds instead of real-number or real-time bounds leads to more generic specifications that are adaptable to model checking at different levels of abstraction, or runtime monitoring on different platforms (e.g., in software vs in hardware). Integer bounds should be read as generic time units, referring to the basic temporal resolution of the system, which can

---

* Corresponding authors.
  *E-mail addresses:* jwli@sei.ecnu.edu.cn (J. Li), vardi@cs.rice.edu (M.Y. Vardi), kyrozier@iastate.edu (K.Y. Rozier).
[1] Part of this work was done at Iowa State University.

generically be resolved to units such as clock ticks or seconds depending on the mission. Integer bounds also allow generic specification with respect to different granularities of time, e.g., to allow easy updates to model-checking models, and re-usable specifications for the same requirements on different embedded systems that may have different resource limits for storing runtime monitors. MLTL has been used in many industrial case studies [1,6–11], and was the official logic of the 2018 Runtime Verification Benchmark Competition [12]. Many specifications from other case studies, in logics such as MTL (Metric Temporal Logic) [2] and STL (Signal Temporal Logic) [3], can be represented in MLTL. We intuitively relate MLTL to LTL and MTL-over-naturals as follows: (1) MLTL formulas are LTL formulas with bounded intervals over temporal operators, and interpreted over finite traces. (2) MLTL formulas are MTL-over-naturals formulas without any unbounded intervals, and interpreted over finite traces.

Despite the practical utility of MLTL, no model checker currently accepts this logic as a specification language. The model checker nuXmv encodes a related logic for use in symbolic model checking, where the $\Box$ and $\Diamond$ operators of an LTLSPEC can have integer bounds [13], though bounds cannot be placed on the $\mathcal{U}$ or $\mathcal{V}$ (the Release operator of nuXmv) operators.

We also critically need an MLTL satisfiability checker to enable specification debugging. Specification is a major bottleneck to the formal verification of mission-based, especially autonomous, systems [14], with a key part of the problem being the availability of good tools for *specification debugging*. Satisfiability checking is an integral tool for specification debugging: [15,16] argued that for every requirement $\varphi$ we need to check $\varphi$ and $\neg\varphi$ for satisfiability; we also need to check the conjunction of all requirements to ensure that they can all be true of the same system at the same time. Specification debugging is essential to model checking [16–18] because a positive answer may not mean there is no bug and a negative answer may not mean there is a bug if the specification is valid/unsatisfiable, respectively. Specification debugging is critical for synthesis and runtime verification (RV) since in these cases there is no model; synthesis and RV are both entirely dependent on the specification. For synthesis, satisfiability checking is the best-available specification-debugging technique, since other techniques, such as vacuity checking (cf. [19,20]) reference a model in addition to the specification. While there are artifacts one can use in RV, specification debugging is still limited outside of satisfiability checking yet central to correct analysis. A false positive due to RV of an incorrect specification can have disastrous consequences, such as triggering an abort of an (otherwise successful) mission to Mars. Arguably, the biggest challenge to creating an RV algorithm or tool is the dearth of benchmarks for checking correctness or comparatively analyzing these [21], where a benchmark consists of some runtime trace, a temporal logic formula reasoning about that trace, and some verdict designating whether the trace at a given time satisfies the requirement formula. A MLTL satisfiability solver is useful for RV benchmark generation [22].

Despite the critical need for an MLTL satisfiability solver, no such tool currently exists. To the best of our knowledge, there is only one available solver (*zot* [23]) for checking the satisfiability of MTL-over-naturals formulas, interpreted over infinite traces. Since MLTL formulas are interpreted over finite traces and there is no trivial reduction from one to another, *zot* cannot be directly applied to MLTL satisfiability checking.

Our approach is inspired by satisfiability-checking algorithms from other logics. For LTL satisfiability solving, we observe that there are multiple efficient translations from LTL satisfiability to model checking, using nuXmv [17]; we therefore consider here translations to nuXmv model checking, both indirectly (as a translation to LTL), and directly using the new KLIVE [24] back-end and the BMC back-end, taking advantage of the bounded nature of MLTL. The bounded nature of MLTL enables us to also consider a direct encoding at the word-level, suitable as input to an SMT solver. Our contribution is both theoretic and experimental. We first consider the complexity of such translations. We prove that the MLTL satisfiability checking problem is NEXPTIME-complete and that satisfiability checking MLTL$_0$, the variant of MLTL where all intervals start at 0, is PSPACE-complete. Secondly, we introduce translation algorithms for MLTL-to-LTL$_f$ (LTL over finite traces [4]), MLTL-to-LTL, MLTL-to-SMV, and MLTL-to-SMT, thus creating four options for MLTL satisfiability checking. Our results show that the MLTL-to-SMT translation with the Z3 SMT solver offers the most scalable performance, though the MLTL-to-SMV translation with an SMV model checker can offer the best performance when the intervals in the MLTL formulas are restricted to small ranges less than 100.

In addition to including all missing proofs, this paper extends the conference version [25] by introducing more details of the MLTL-to-SMT translation, e.g., examplize the encoding and propose different SMT encodings for MLTL satisfiability checking, and showing more experimental results to support our previous conclusion as well as to evaluate the performance of different SMT encodings.

## 2. Preliminaries

A (closed) interval over naturals $I = [a, b]$ ($0 \le a \le b$ are natural numbers) is a set of naturals $\{i \mid a \le i \le b\}$. $I$ is called *bounded* iff $b < +\infty$; otherwise $I$ is *unbounded*. MLTL is defined using bounded intervals. Unlike Metric Temporal Logic (MTL) [26], it is not necessary to introduce open or half-open intervals over the natural domain, as every open or half-open bounded interval is reducible to an equivalent closed bounded interval, e.g., $(1, 2) = \emptyset$, $(1, 3) = [2, 2]$, $(1, 3] = [2, 3]$, etc. Let $\mathcal{AP}$ be a set of atomic propositions, then the syntax of a formula in MLTL is

$$\varphi ::= \text{true} \mid \text{false} \mid p \mid \neg\varphi \mid \varphi \land \psi \mid \varphi \lor \psi \mid \varphi\,\mathcal{U}_I\,\psi \mid \varphi\mathcal{R}_I\psi$$

where $I$ is a bounded interval, $p \in \mathcal{AP}$ is an *atom*, and $\varphi$ and $\psi$ are subformulas.

The semantics of MLTL formulas is interpreted over finite traces bounded by base-10 (decimal) intervals. Let $\pi$ be a finite trace in which every position $\pi[i]$ ($i \ge 0$) is over $2^{\mathcal{AP}}$, and $|\pi|$ denotes the length of $\pi$ ($|\pi| < +\infty$ when $\pi$ is a

finite trace). We use $\pi_i$ ($|\pi| > i \geq 0$) to represent the suffix of $\pi$ starting from position $i$ (including $i$). Let $a, b \in \mathbb{I}, a \leq b$; we define that $\pi$ models (satisfies) an MLTL formula $\varphi$, denoted as $\pi \models \varphi$, as follows:

- $\pi \models p$ iff $p \in \pi[0]$;
- $\pi \models \neg\varphi$ iff $\pi \not\models \varphi$;
- $\pi \models \varphi \wedge \psi$ iff $\pi \models \varphi$ and $\pi \models \psi$;
- $\pi \models \varphi \, \mathcal{U}_{[a,b]} \, \psi$ iff $|\pi| > a$ and, there exists $i \in [a, b]$, $i < |\pi|$ such that $\pi_i \models \psi$ and for every $j \in [a, b]$, $j < i$ it holds that $\pi_j \models \varphi$;

Given two MLTL formulas $\varphi, \psi$, we denote $\varphi = \psi$ iff they are *syntactically equivalent*, and $\varphi \equiv \psi$ iff they are *semantically equivalent*, i.e., $\pi \models \varphi$ iff $\pi \models \psi$ for a finite trace $\pi$. In MLTL semantics, we define false $\equiv \neg$true, $\varphi \vee \psi \equiv \neg(\neg\varphi \wedge \neg\psi)$, and $\neg(\varphi \, \mathcal{U}_I \, \psi) \equiv (\neg\varphi \mathcal{R}_I \neg\psi)$. MLTL keeps the standard operator equivalences from LTL, including $(\Diamond_I\varphi) \equiv (true \, \mathcal{U}_I\varphi)$ (Eventually), $(\Box_I\varphi) \equiv (false \, \mathcal{R}_I \, \varphi)$ (Globally), and $(\varphi \, \mathcal{R}_I \, \psi) \equiv (\neg(\neg\varphi \, \mathcal{U}_I \, \neg\psi))$. Notably, MLTL discards the neXt ($\mathcal{X}$) operator, which is essential in LTL [5], since $\mathcal{X}\varphi$ is semantically equivalent to $\Box_{[1,1]}\varphi$.

Compared to the traditional MTL-over-naturals[2] [27], the Until formula in MLTL is interpreted in a slightly different way. In MTL-over-naturals, the satisfaction of $\varphi \, \mathcal{U}_I \, \psi$ requires $\varphi$ to hold from position 0 to the position where $\psi$ holds (in $I$), while in MLTL $\varphi$ is only required to hold within the interval $I$, before the time $\psi$ holds. From the perspective of writing specifications, cf. [1,8], this adjustment is more user-friendly. It is not hard to see that MLTL is as expressive as the standard MTL-over-naturals: the formula $\varphi \, \mathcal{U}_{[a,b]} \, \psi$ in MTL-over-naturals can be represented as $(\Box_{[0,a-1]}\varphi) \wedge (\varphi\mathcal{U}_{[a,b]} \, \psi)$ in MLTL; $\varphi \, \mathcal{U}_{[a,b]} \, \psi$ in MLTL can be represented as $\Diamond_{[a,a]}(\varphi \, \mathcal{U}_{[0,b-a]} \, \psi)$ in MTL-over-naturals.

We say an MLTL formula is in *BNF* (Backus Naur Form) if the formula contains only $\neg$, $\wedge$ and $\mathcal{U}_I$ operators. It is trivial to see that every MLTL formula can be converted to its (semantically) equivalent BNF with a linear cost. Consider $\varphi = (\neg a) \vee ((\neg b)\mathcal{R}_I(\neg c))$ as an example. Its BNF form is $\neg(a \wedge (b \, \mathcal{U}_I \, c))$. Without explicit clarification, this paper assumes that every MLTL formula is in BNF.

The closure of an MLTL formula $\varphi$, denoted as $cl(\varphi)$, is a set of formulas such that: 1) $\varphi \in cl(\varphi)$; 2) $\varphi \in cl(\varphi)$ if $\neg\varphi \in cl(\varphi)$; 3) $\varphi, \psi \in cl(\varphi)$ if $\varphi \, op \, \psi \in cl(\varphi)$, where $op$ can be $\wedge$ or $\mathcal{U}_I$. Let $|cl(\varphi)|$ be the size of $cl(\varphi)$. Since the definition of $cl(\varphi)$ ignores the intervals in $\varphi$, $|cl(\varphi)|$ is linear in the number of operators in $\varphi$. We also define the closure(*) of an MLTL formula $\varphi$, denoted $cl^*(\varphi)$, as the set of formulas such that: 1) $cl(\varphi) \subseteq cl^*(\varphi)$; 2) if $\varphi \, \mathcal{U}_{[a,b]} \, \psi \in cl^*(\varphi)$ for $0 < a \leq b$, then $\varphi \, \mathcal{U}_{[a-1,b-1]} \, \psi$ is in $cl^*(\varphi)$; 3) if $\varphi \, \mathcal{U}_{[0,b]} \, \psi \in cl^*(\varphi)$ for $0 < b$, then $\varphi \, \mathcal{U}_{[0,b-1]} \, \psi$ is in $cl^*(\varphi)$. Let $|cl^*(\varphi)|$ be the size of $cl^*(\varphi)$ and $K$ be the maximal natural number in the intervals of $\varphi$. It is not hard to see that $|cl^*(\varphi)|$ is at most $K \cdot |cl(\varphi)|$.

We also consider a fragment of MLTL, namely MLTL$_0$, which is more frequently used in practice, cf. [6,1]. Informally speaking, MLTL$_0$ formulas are MLTL formulas in which all intervals start from 0. For example, $\Diamond_{[0,4]}a \wedge (a \, \mathcal{U}_{[0,1]} \, b)$ is a MLTL$_0$ formula, while $\Diamond_{[2,4]}a$ is not.

Given an MLTL formula $\varphi$, the *satisfiability problem* asks whether there is a finite trace $\pi$ such that $\pi \models \varphi$ holds. To solve this problem, we can reduce it to the satisfiability problem of the related logics LTL and LTL$_f$ (LTL over finite traces [4]), and leverage the off-the-shelf satisfiability checking solvers for these well-explored logics. We abbreviate MLTL, LTL, and LTL$_f$ satisfiability checking as MLTL-SAT, LTL-SAT, and LTL$_f$-SAT respectively.

**Linear Temporal Logic over finite traces.** We assume readers are familiar with LTL (over infinite traces) [5]. Linear Temporal Logic over finite traces, short for LTL$_f$ [4], is a variant of LTL that has the same syntax, except that for LTL$_f$, the dual operator of $\mathcal{X}$ is $\mathcal{N}$ (weak Next), which differs $\mathcal{X}$ in the last state of the finite trace. In the last state of a finite trace, $\mathcal{X}\psi$ can never be satisfied, while $\mathcal{N}\psi$ is satisfiable. Given an LTL$_f$ formula $\varphi$, there is an LTL formula $\psi$ such that $\varphi$ is satisfiable iff $\psi$ is satisfiable. In detail, $\psi = \Diamond Tail \wedge t(\varphi)$ where $Tail$ is a new atom identifying the end of the satisfying trace and $t(\varphi)$ is constructed as follows:

- $t(p) = p$ where $p$ is an atom;
- $t(\neg\psi) = \neg t(\psi)$;
- $t(\mathcal{X}\psi) = \neg Tail \wedge \mathcal{X}t(\psi)$;
- $t(\psi_1 \wedge \psi_2) = t(\psi_1) \wedge t(\psi_2)$;
- $t(\psi_1 U \psi_2) = t(\neg Tail \wedge \psi_1)\mathcal{U}t(\psi_2)$.

In the above reduction, $\varphi$ is in BNF. Since the reduction is linear in the size of the original LTL$_f$ formula and LTL-SAT is PSPACE-complete [28], LTL$_f$-SAT is also a PSPACE-complete problem [4].

## 3. Complexity of MLTL-SAT

It is known that the complexity of MITL (Metric Interval Temporal Logic) satisfiability is EXPSPACE-complete, and the satisfiability complexity of the fragment of MITL named MITL$_{0,\infty}$ is PSPACE-complete [29]. MLTL (resp. MLTL$_0$) can be

---

[2] In this paper, MTL-over-naturals is interpreted over finite traces.

viewed as a variant of MITL (resp. MITL$_{0,\infty}$) that is interpreted over the naturals. We show that MLTL satisfiability checking is NEXPTIME-complete, via a reduction from MLTL to LTL$_f$.

**Lemma 1.** *Let $\varphi$ be an MLTL formula, and $K$ be the maximal natural appearing in the intervals of $\varphi$ ($K$ is set to 1 if there are no intervals in $\varphi$). There is an LTL$_f$ formula $\theta$ that recognizes the same language as $\varphi$. Moreover, the size of $\theta$ is in $O(K \cdot |cl(\varphi)|)$.*

**Proof.** For an MLTL formula $\varphi$, we define the LTL$_f$ formula $f(\varphi)$ recursively as follows:

- If $\varphi = \text{true}, \text{false}$, or an atom $p$, $f(\varphi) = \varphi$;
- If $\varphi = \neg\psi$, $f(\varphi) = \neg f(\psi)$;
- If $\varphi = \xi \wedge \psi$, $f(\varphi) = f(\xi) \wedge f(\psi)$;
- If $\varphi = \xi \, \mathcal{U}_{[a,b]} \, \psi$,

$$f(\varphi) = \begin{cases} \mathcal{X}(f(\xi \, \mathcal{U}_{[a-1,b-1]} \, \psi)), & \text{if } 0 < a \leq b; \\ f(\psi) \vee (f(\xi) \wedge \mathcal{X}(f(\xi \mathcal{U}_{[a,b-1]}\psi))), & \text{if } a = 0 \text{ and } 0 < b; \\ f(\psi), & \text{if } a = 0 \text{ and } b = 0. \end{cases}$$

$\mathcal{X}$ represents the neXt operator in LTL$_f$. Based on the above translation, the size of $f(\varphi)$ is at most linear to $K \cdot |cl(\varphi)|$, i.e., in $O(K \cdot |cl(\varphi)|)$. Now we prove by induction over the type of $\varphi$ that $\pi \models \varphi$ iff $\pi \models f(\varphi)$ for a finite trace $\pi$, i.e. $\varphi$ and $f(\varphi)$ accept the same language. Obviously, $\pi \models \varphi$ iff $\pi \models f(\varphi)$ holds when $\varphi$ is true, false or an atom $p$. Inductively,

- if $\varphi = \neg\psi$, $f(\varphi) = \neg f(\psi)$. According to the assumption hypothesis, $\pi \models \psi$ iff $\psi \models f(\psi)$ holds for some finite trace $\pi$. As a result, $\pi \not\models \psi$ iff $\pi \not\models f(\psi)$ holds, which is equivalent to say $\pi \models \neg\psi$ iff $\pi \models f(\varphi)$ holds;
- if $\varphi = \xi \wedge \psi$, $f(\varphi) = f(\xi) \wedge f(\psi)$. According to the assumption hypothesis, $\pi_1 \models \xi$ iff $\pi_1 \models f(\xi)$ and $\pi_2 \models \psi$ iff $\pi_2 \models f(\psi)$ hold for two finite traces $\pi_1$ and $\pi_2$. As a result, for a finite trace $\pi$, it is true that $\pi \models \xi \wedge \psi$ iff $\pi \models f(\xi) \wedge f(\psi)$ holds, which is equivalent to say $\pi \models \xi \wedge \psi$ iff $\pi \models f(\xi) \wedge f(\psi)$;
- if $\varphi = \xi \, \mathcal{U}_{[a,b]} \, \psi$,
  - when $0 < a \leq b$, $f(\varphi)$ is $\mathcal{X}(f(\xi \, \mathcal{U}_{[a-1,b-1]} \, \psi))$. Based on the assumption hypothesis, $\pi' \models \xi \, \mathcal{U}_{[a-1,b-1]} \, \psi$ iff $\pi' \models f(\xi \, \mathcal{U}_{[a-1,b-1]} \, \psi)$ holds for a finite trace $\pi'$. Then according to the semantics of the $\mathcal{X}$ operator and the MLTL formulas, we have that $\varphi$ is semantically equivalent to $\mathcal{X}(\xi \, \mathcal{U}_{[a-1,b-1]} \, \psi)$. As a result, $\pi \models \varphi$ iff $\pi \models \mathcal{X}(f(\xi \, \mathcal{U}_{[a-1,b-1]} \, \psi))$ for every $\pi = \omega \cdot \pi'$ ($\omega \in 2^{\Sigma_\varphi}$);
  - when $0 = a < b$, $f(\varphi)$ is $f(\psi) \vee (f(\xi) \wedge \mathcal{X}(f(\xi \, \mathcal{U}_{[a,b-1]} \, \psi)))$. According to the semantics of the $\mathcal{X}$ operator and the MLTL formulas, we have that $\varphi$ is semantically equivalent to $\psi \vee (\xi \wedge \mathcal{X}(\xi \, \mathcal{U}_{[a,b-1]} \, \psi))$. Thus for some finite trace $\pi$, $\pi \models \varphi$ holds iff $\pi \models \psi$ or $\pi \models \xi \wedge \mathcal{X}(\xi \, \mathcal{U}_{[a,b-1]} \, \psi)$ holds. From the assumption hypothesis, we have that $\pi \models \psi$ iff $\pi \models f(\psi)$ holds, or $\pi \models \xi \wedge \mathcal{X}(\xi \, \mathcal{U}_{[a,b-1]} \, \psi)$ iff $\pi \models f(\xi) \wedge \mathcal{X}(f(\xi \, \mathcal{U}_{[a,b-1]} \, \psi))$ holds. That means, $\pi \models \varphi$ iff $\pi \models f(\psi)$ or $\pi \models f(\xi) \wedge \mathcal{X}(f(\xi \, \mathcal{U}_{[a,b-1]} \, \psi))$ holds;
  - when $0 = a = b$, $f(\varphi) = f(\psi)$. Based on the assumption hypothesis, $\pi \models \psi$ iff $\pi \models f(\psi)$ for some finite trace $\pi$. Also, according to the MLTL semantics, $\varphi$ is semantically equivalent to $\psi$. As a result, we have that $\pi \models \varphi$ iff $\pi \models f(\psi)$ holds, which means $\pi \models \varphi$ iff $\pi \models f(\varphi)$ holds.

Let $\theta = f(\varphi)$ and we can conclude that $\theta$ and $\varphi$ accepts the same language, and the size of $\theta$ is in $O(K \cdot |cl(\varphi)|)$. □

Notably, the construction of $f(\varphi)$ can always terminate since $f$ is defined recursively and every time it is invoked, the bounds of temporal operators decrease. Consider the special case that $a\mathcal{U}_{[0,0]}b \equiv b$, the temporal operators are erased as soon as the bound becomes $[0, 0]$. Therefore, $f(\varphi)$ can be terminated when there are no temporal operators left.

We use the construction shown in Lemma 1 to explore several useful properties of MLTL. For instance, the LTL$_f$ formula translated from an MLTL formula contains only the $\mathcal{X}$ temporal operator or its dual $\mathcal{N}$, which represents weak Next [30,31], and the number of these operators is strictly smaller than $K \cdot |cl(\varphi)|$. Every $\mathcal{X}$ or $\mathcal{N}$ subformula in the LTL$_f$ formula corresponds to some temporal formula in $cl^*(\varphi)$. Notably, because the natural-number intervals in $\varphi$ are written in base 10 (decimal) notation, the blow-up in the translation of Lemma 1 is exponential.

The next lower bound is reminiscent of the NEXPTIME-lower bound shown in [32] for a fragment of Metric Interval Temporal Logic (MITL), but is different in the details of the proof as the two logics are quite different.

**Theorem 1.** *The complexity of MLTL satisfiability checking is NEXPTIME-complete.*

**Proof.** By Lemma 1, there is an LTL$_f$ formula $\theta$ that accepts the same traces as MLTL formula $\varphi$, and the size of $\theta$ is in $O(K \cdot |cl(\varphi)|)$. The only temporal connectives used in $\theta$ are $\mathcal{X}$ and $\mathcal{N}$, since the translation to LTL$_f$ reduces all MLTL temporal connectives in $\varphi$ to nested $\mathcal{X}$'s or $\mathcal{N}$'s (produced by simplifying $\neg\mathcal{X}$). Thus, if $\theta$ is satisfiable, then it is satisfiable by a trace whose length is bounded by the length of $\theta$. Thus, we can just guess a trace $\pi$ of exponential length of $\theta$ and check that it satisfies $\varphi$. As a result, the upper bound for MLTL-SAT is NEXPTIME.

Before proving the NEXPTIME lower bound, recall the PSPACE-lower bound proof in [28] for LTL satisfiability. The proof reduces the acceptance problem for a linear-space bounded Turing machine $M$ to LTL satisfiability. Given a Turing machine $M$ and an integer $k$, we construct a formula $\varphi_M$ such that $\varphi_M$ is satisfiable iff $M$ accepts the empty tape using $k$ tape cells. The argument is that we can encode such a space-bounded computation of $M$ by a trace $\pi$ of length $c^k$ for some constant $c$, and then use $\varphi_M$ to force $\pi$ to encode an accepting computation of $M$. The formula $\varphi_M$ has to match corresponding points in successive configurations of $M$, which can be expressed using a $O(k)$-nested $\mathcal{X}$'s, since such points are $O(k)$ points apart.

To prove a NEXPTIME-lower bound for MLTL, we reduce the acceptance problem for exponentially bounded non-deterministic Turing machines to MLTL satisfiability. Given a non-deterministic Turing machine $M$ and an integer $k$, we construct an MLTL formula $\varphi_M$ of length $O(k)$ such that $\varphi_M$ is satisfiable iff $M$ accepts the empty tape in time $2^k$. Note that such a computation of a $2^k$-time bounded Turing machines consists of $2^k$ many configurations of length $2^k$ each, so the whole computation is of exponential length – $4^k$, and can be encoded by a trace $\pi$ of length $4^k$, where every point of $\pi$ encodes one cell in the computation of $M$. Unlike the reduction in [28], in the encoding here corresponding points in successive configurations are exponentially far ($2^k$) from each other, because each configuration has $2^k$ cells, so the relationship between such successive points cannot be expressed in LTL. Because, however, the constants in the intervals of MLTL are written in base-10 (decimal) notation, we can write formulas of size $O(k)$, e.g., formulas of the form $p\ \mathcal{U}_{[0,2^k]}\ q$, that relate points that are $2^k$ apart.

The key is to express the fact that one Turing machine configuration is a proper successor of another configuration using a formula of size $O(k)$. In the PSPACE-lower-bound proof of [28], LTL formulas of size $O(k)$ relate successive configurations of $k$-space-bounded machines. Here MLTL formulas of size $O(k)$ relate successive configurations of $2^k$-time-bounded machines. Thus, we can write a formula $\varphi_M$ of length $O(k)$ that forces trace $\pi$ to encode a computation of $M$ of length $2^k$. □

Now we consider $\text{MLTL}_0$ formulas, and prove that the complexity of checking the satisfiability of $\text{MLTL}_0$ formulas is PSPACE-complete. We first introduce the following lemma to show an inherent feature of $\text{MLTL}_0$ formulas.

**Lemma 2.** *The conjunction of identical $\text{MLTL}_0$ $\mathcal{U}$-rooted formulas is equivalent to the conjunct with the smallest interval range:* $(\xi\ \mathcal{U}_{[0,a]}\ \psi) \wedge (\xi\ \mathcal{U}_{[0,b]}\ \psi) \equiv (\xi\ \mathcal{U}_{[0,a]}\ \psi)$, *where $b > a$.*

**Proof.** We first prove that for $i \geq 0$, the equation $(\xi\ \mathcal{U}_{[0,i]}\ \psi) \wedge (\xi\ \mathcal{U}_{[0,i+1]}\ \psi) \equiv (\xi\ \mathcal{U}_{[0,i]}\ \psi)$ holds. When $i = 0$, we have $(\xi\ \mathcal{U}_{[0,0]}\ \psi) \equiv f(\psi)$ and $(\xi\ \mathcal{U}_{[0,1]}\ \psi) \equiv (f(\psi) \vee f(\xi) \wedge \mathcal{X}(f(\psi)))$. So $(\xi\ \mathcal{U}_{[0,0]}\ \psi) \wedge (\xi\ \mathcal{U}_{[0,1]}\ \psi) \equiv f(\psi) \equiv (\xi\ \mathcal{U}_{[0,0]}\ \psi)$ is true. Inductively, assume that $(\xi\ \mathcal{U}_{[0,k]}\ \psi) \wedge (\xi\ \mathcal{U}_{[0,k+1]}\ \psi) \equiv (\xi\ \mathcal{U}_{[0,k]}\ \psi)$ is true for $k \geq 0$. When $i = k + 1$, we have $(\xi\ \mathcal{U}_{[0,k+1]}\ \psi) \equiv (f(\psi) \vee f(\xi) \wedge \mathcal{X}(\xi\ \mathcal{U}_{[0,k]}\ \psi))$ and $(\xi\ \mathcal{U}_{[0,k+2]}\ \psi) \equiv (f(\psi) \vee f(\xi) \wedge \mathcal{X}(\xi\ \mathcal{U}_{[0,k+1]}\ \psi))$. By hypothesis assumption, $(\xi\ \mathcal{U}_{[0,k]}\ \psi) \wedge (\xi\ \mathcal{U}_{[0,k+1]}\ \psi) \equiv (\xi\ \mathcal{U}_{[0,k]}\ \psi)$ implies that the following equivalence is true:

$$(\xi\ \mathcal{U}_{[0,k+1]}\ \psi) \wedge (\xi\ \mathcal{U}_{[0,k+2]}\ \psi)$$
$$\equiv\ (f(\psi) \vee (f(\xi) \wedge \mathcal{X}(\xi\ \mathcal{U}_{[0,k]}\ \psi))) \wedge (f(\psi) \vee (f(\xi) \wedge \mathcal{X}(\xi\ \mathcal{U}_{[0,k+1]}\ \psi)))$$
$$\equiv\ f(\psi) \vee (f(\xi) \wedge \mathcal{X}(\xi\ \mathcal{U}_{[0,k]}\ \psi \wedge \xi\ \mathcal{U}_{[0,k+1]}\ \psi))$$
$$\equiv\ f(\psi) \vee (f(\xi) \wedge \mathcal{X}(\xi\ \mathcal{U}_{[0,k]}\ \psi))$$
$$\equiv\ (\xi\ \mathcal{U}_{[0,k+1]}\ \psi).$$

Since $(\xi\ \mathcal{U}_{[0,i]}\ \psi) \wedge (\xi\ \mathcal{U}_{[0,i+1]}\ \psi) \equiv (\xi\ \mathcal{U}_{[0,i]}\ \psi)$ is true, we can prove by induction that $(\xi\ \mathcal{U}_{[0,i]}\ \psi) \wedge (\xi\ \mathcal{U}_{[0,j]}\ \psi) \equiv (\xi\ \mathcal{U}_{[0,i]}\ \psi)$ is true, where $j > i$. Because $b > a$ is true, it directly implies that $(\xi\ \mathcal{U}_{[0,a]}\ \psi) \wedge (\xi\ \mathcal{U}_{[0,b]}\ \psi) \equiv (\xi\ \mathcal{U}_{[0,a]}\ \psi)$ is true. □

**Lemma 3.** *$\mathcal{X}$-free $\text{LTL}_f$-SAT is reducible to $\text{MLTL}_0$-SAT at a linear cost.*

**Proof.** According to [28], the satisfiability checking of $\mathcal{X}$-free LTL formulas is still PSPACE-complete. This also applies to the satisfiability checking of $\mathcal{X}$-free $\text{LTL}_f$ formulas. Given an $\mathcal{X}$-free $\text{LTL}_f$ formula $\varphi$, we construct the corresponding MLTL formula $m(\varphi)$ recursively as follows:

- $m(p) = p$ where $p$ is an atom;
- $m(\neg\xi) = \neg m(\xi)$;
- $m(\xi \wedge \psi) = m(\xi) \wedge m(\psi)$;
- $m(\xi\ \mathcal{U}\ \psi) = m(\xi)\ \mathcal{U}_{[0,2^{|\varphi|}]}\ m(\psi)$.

Notably for the Until $\text{LTL}_f$ formula, we bound it with the interval $[0, 2^{|\varphi|}]$, where $\varphi$ is the original $\mathcal{X}$-free $\text{LTL}_f$ formula, in the corresponding MLTL formula, which is motivated by the fact that every satisfiable $\text{LTL}_f$ formula has a finite model whose length is less than $2^{|\varphi|}$ [4]. The above translation has linear blow-up, because the integers in intervals use the decimal

notation. Now we prove by induction over the type of $\varphi$ that $\varphi$ is satisfiable iff $m(\varphi)$ is satisfiable. That is, we prove that $(\Rightarrow)$ $\pi \models \varphi$ implies $\pi \models m(\varphi)$ and $(\Leftarrow)$ $\pi \models m(\varphi)$ implies $\pi \models \varphi$, for some finite trace $\pi$.

We consider the Until formula $\eta = \xi \mathcal{U} \psi$ (noting that $\varphi$ is fixed to the original LTL$_f$ formula), and the proofs are trivial for other types. $(\Rightarrow)$ $\eta$ is satisfiable implies there is a finite trace $\pi$ such that $\pi \models \eta$ and $|\pi| \leq 2^{|\varphi|}$ [4]. Moreover, $\pi \models \eta$ holds iff there is $0 \leq i$ such that $\pi_i \models \psi$ and for every $0 \leq j < i$, $\pi_j \models \xi$ is true (from LTL$_f$ semantics). By the induction hypothesis, $\pi_i \models \psi$ implies $\pi_i \models m(\psi)$ and $\pi_j \models \xi$ implies $\pi_j \models m(\xi)$. Also, $i \leq 2^{|\varphi|}$ is true because of $|\pi| \leq 2^{|\varphi|}$. As a result, $\pi \models \eta$ implies that there is $0 \leq i \leq 2^{|\varphi|}$ such that $\pi_i \models m(\psi)$ and for every $0 \leq j < i$, $\pi_j \models m(\xi)$ is true. According to the MLTL semantics, $\pi \models m(\eta)$ is true. $(\Leftarrow)$ $m(\eta)$ is satisfiable implies there is a finite trace $\pi$ such that $\pi \models m(\eta)$. According to MLTL semantics, there is $0 \leq i \leq 2^{|\varphi|}$ such that $\pi_i \models m(\psi)$ and for every $0 \leq j < i$ it holds that $\pi_j \models m(\xi)$. By hypothesis assumption, $\pi_i \models m(\psi)$ implies $\pi_i \models \psi$ and $\pi_j \models m(\xi)$ implies $\pi_j \models \xi$. Also, $0 \leq i \leq 2^{|\varphi|}$ implies $0 \leq i$. As a result, $\pi \models m(\eta)$ implies that there is $0 \leq i$ such that $\pi_i \models \psi$ and for every $0 \leq j < i$ it holds that $\pi_j \models \xi$. From LTL$_f$ semantics, it is true that $\pi \models \eta$. $\square$

**Theorem 2.** *The complexity of checking the satisfiability of MLTL$_0$ is PSPACE-complete.*

**Proof.** Since Lemma 3 shows a linear reduction from $\mathcal{X}$-free LTL$_f$-SAT to MLTL$_0$-SAT and $\mathcal{X}$-free LTL$_f$-SAT is PSPACE-complete [4], it directly implies that the lower bound of MLTL$_0$-SAT is PSPACE-hard.

For the upper bound, recall from the proof of Theorem 1 that an MLTL formula $\varphi$ is translated to an LTL$_f$ formula $\theta$ of length $K \cdot |cl(\varphi)|$, which, as we commented, involved an exponential blow-up in the notation for $K$. Following the automata-theoretic approach for satisfiability, one would translate $\theta$ to an NFA and check its non-emptiness [4]. Normally, such a translation would involve another exponential blow-up. We show that this is not the case for MLTL$_0$. Recalling from the automaton construction in [4] that every state of the automaton is a set of subformulas of $\theta$, the size of a state is at most $K \cdot |cl(\varphi)|$. In the general case, if $\psi_1, \psi_2$ are two subformulas of $\theta$ corresponding to the MLTL formulas $\xi \mathcal{U}_{I_1} \psi$ and $\xi \mathcal{U}_{I_2} \psi$, $\psi_1$ and $\psi_2$ can be in the same state of the automaton, which implies that the size of the state can be at most $K \cdot |cl(\varphi)|$. When the formula $\varphi$ is restricted to MLTL$_0$, we show that the exponential blow-up can be avoided. Lemma 2 shows that either $\psi_1$ or $\psi_2$ in the state is enough, since assuming $I_1 \subseteq I_2$, then $(\psi_1 \wedge \psi_2) \equiv \psi_1$, by Lemma 2. So the size of the state in the automaton for a MLTL$_0$ formula $\varphi$ is at most $|cl(\varphi)|$. For each subformula in the state, there can be $K$ possible values (e.g., for $\Diamond_I \xi$ in the state, we can have $\Diamond_{[0,1]} \xi$, $\Diamond_{[0,2]} \xi$, etc.). Therefore the size of the automaton is in $O(2^{|cl(\varphi)|} \cdot K^{|cl(\varphi)|}) \approx 2^{O(|cl(\varphi)|)}$. Therefore, MLTL$_0$ satisfiability checking is a PSPACE-complete problem. $\square$

## 4. Implementation of MLTL-SAT

We first show how to reduce MLTL-SAT to the well-explored LTL$_f$-SAT and LTL-SAT. Then we introduce two new satisfiability-checking strategies based on the inherent properties of MLTL formulas, which are able to leverage the state-of-art model-checking and SMT-solving techniques.

### 4.1. MLTL-SAT via logic translation

For a formula $\varphi$ from one logic, and $\psi$ from another logic, we say $\varphi$ and $\psi$ are *equi-satisfiable* when $\varphi$ is satisfiable under its semantics iff $\psi$ is satisfiable under its semantics. Based on Lemma 1 and Theorem 1, we have the following corollary,

**Corollary 1** (*MLTL-SAT to LTL$_f$-SAT*). *MLTL-SAT can be reduced to LTL$_f$-SAT with an exponential blow-up.*

From Corollary 1, MLTL-SAT is reducible to LTL$_f$-SAT, enabling use of the off-the-shelf LTL$_f$ satisfiability solvers, cf. aaltaf [30]. It is also straightforward to consider MLTL-SAT via LTL-SAT; LTL-SAT has been studied for more than a decade, and many off-the-shelf LTL solvers are available, cf. [15,17,33].

**Theorem 3** (*MLTL to LTL*). *For an MLTL formula $\varphi$, there is an LTL formula $\theta$ such that $\varphi$ and $\theta$ are equi-satisfiable, and the size of $\theta$ is in $O(K \cdot |cl(\varphi)|)$, where $K$ is the maximal integer in $\varphi$.*

**Proof.** Lemma 1 provides a translation from the MLTL formula $\varphi$ to the equivalent LTL$_f$ formula $\varphi'$, with a blow-up of $O(K \cdot |cl(\varphi)|)$. As shown in Section 2, there is a linear translation from the LTL$_f$ formula $\varphi'$ to its equi-satisfiable LTL formula $\theta$ [4]. Therefore, the blow-up from $\varphi$ to $\theta$ is in $O(K \cdot |cl(\varphi)|)$. $\square$

**Corollary 2** (*MLTL-SAT to LTL-SAT*). *MLTL-SAT can be reduced to LTL-SAT with an exponential blow-up.*

Since MLTL-SAT is reducible to LTL-SAT, MLTL-SAT can also benefit from the power of LTL satisfiability solvers. Moreover, the reduction from MLTL-SAT to LTL-SAT enables leveraging modern model-checking techniques to solve the MLTL-SAT problem, due to the fact that LTL-SAT has been shown to be reducible to model checking with a linear blow-up [15,16].

**Corollary 3** *(MLTL-SAT to LTL-model-checking). MLTL-SAT can be reduced to LTL model checking with an exponential blow-up.*

In our implementation, we choose the model checker nuXmv [34] for LTL satisfiability checking, as it allows an LTL formula to be directly input as the temporal specification together with a universal model as described in [15,16].

### 4.2. Model generation

Using the LTL formula as the temporal specification in nuXmv has been shown, however, to not be the most efficient way to use model checking for satisfiability checking [17]. Consider the MLTL formula $\Diamond_{[0,10]}a \wedge \Diamond_{[1,11]}a$. The translated LTL$_f$ formula is $f(\Diamond_{[0,10]}a) \wedge \mathcal{X}(f(\Diamond_{[0,10]}a))$, where $f(\Diamond_{[0,10]}a)$ has to be constructed twice. To avoid such redundant construction, we follow [17] and encode directly the input MLTL formula as an SMV model (the input model of nuXmv) rather than treating the LTL formula, which is obtained from the input MLTL formula, as a specification.

An SMV [35] model consists of a Boolean transition system $Sys = (V, I, T)$, where $V$ is a set of Boolean variables, $I$ is a Boolean formula representing the initial states of $Sys$, and $T$ is the Boolean transition formula. Moreover, a specification to be verified against the system is also contained in the SMV model (here we focus on the LTL specification). Given the input MLTL formula $\varphi$, we construct the corresponding SMV model $M_\varphi$ as follows.

1. Introduce a Boolean variable for each atom in $\varphi$ as well as for "*Tail*" (new variable identifying the end of a finite trace).
2. Introduce a Boolean variable $\mathcal{X}\_\psi$ for each $\mathcal{U}$ formula $\psi$ in $cl^*(\varphi)$, which represents the temporal formula $\mathcal{X}\psi$.
3. Introduce a temporary Boolean variable[3] $T\_\psi$ for each $\mathcal{U}$ formula in $cl^*(\varphi)$.
4. A Boolean formula $e(\psi)$ is used to represent the formula $\psi$ in $cl^*(\varphi)$ in the SMV model, which is defined recursively as follows.
   (a) $e(\psi) = \psi$, if $\psi$ is an Boolean atom;
   (b) $e(\psi) = \neg e(\psi_1)$, if $\psi = \neg\psi_1$;
   (c) $e(\psi) = e(\psi_1) \wedge e(\psi_2)$, if $\psi = \psi_1 \wedge \psi_2$;
   (d) $e(\psi) = T\_\psi$, if $\psi$ is an $\mathcal{U}$ formula.
5. Let the initial Boolean formula $I$ of the system $Sys$ be $e(\varphi)$.
6. For each temporary variable $T\_\psi$, create a DEFINE statement according to the type and interval of $\psi$, as follows.

$$T_{\psi_1 \mathcal{U}_{[a,b]}\psi_2} = \begin{cases} \mathcal{X}\_(\psi_1 \mathcal{U}_{[a-1,b-1]}\psi_2), & \text{if } 0 < a \le b; \\ e(\psi_2) \vee (e(\psi_1) \wedge \mathcal{X}\_(\psi_1 \mathcal{U}_{[0,b-1]}\psi_2)), & \text{if } a = 0 \text{ and } 0 < b; \\ e(\psi_2), & \text{if } a = 0 \text{ and } b = 0. \end{cases}$$

7. Create the Boolean formula $(\mathcal{X}\_\psi \leftrightarrow (\neg Tail \wedge next(e(\psi))))$ for each $\mathcal{X}\_\psi$ in the VAR list (the set $V$ in $Sys$) of the SMV model. In the formula, *next* is a specific function in SMV language to define the values in the next period. For example, $next(e(\psi))$ means that $e(\psi)$ has to be true in the next period.
8. Finally, designate the LTL formula $\Box\neg Tail$ as the temporal specification of the SMV model $M_\varphi$ (which implies that a counterexample trace satisfies $\Diamond Tail$).

For $sys = (V, I, T)$, $V$ is defined to include all the above three kinds of Boolean variables created in Item 1-3; $I$ is defined in Item 5, and the transition formula $T$ is defined as the combination of Item 6 and 7.

In a nutshell, the SMV model for $\theta$ has the analogous structure in Table 1.

**Encoding heuristics for MLTL$_0$ formulas.** We also encode the rules shown in Lemma 2 to prune the state space for checking the satisfiability of MLTL$_0$ formulas. These rules are encoded using the INVAR constraint in the SMV model. Taking the $\mathcal{U}$ formula as an example, we encode $T\_(\psi_1 \mathcal{U}_{[0,a]}\psi_2) \wedge T\_(\psi_1 \mathcal{U}_{[0,a-1]}\psi_2) \leftrightarrow T\_(\psi_1 \mathcal{U}_{[0,a-1]}\psi_2)$ ($a > 0$) for each $\psi_1 \mathcal{U}_{[0,a]}\psi_2$ in $cl^*(\varphi)$. Similar encodings also apply to the $\mathcal{R}$ formulas in $cl^*(\varphi)$. Theorem 4 below guarantees the correctness of the translation, and it can be proved by induction over the type of $\varphi$ and the construction of the SMV model.

**Theorem 4.** *The MLTL formula $\varphi$ is satisfiable iff the corresponding SMV model $M_\varphi$ violates the LTL property $\Box\neg Tail$.*

There are different techniques that can be used for LTL model checking. Based on the latest evaluation of LTL satisfiability checking [33], the KLIVE [24] back-end implemented in the SMV model checker nuXmv [34] produces the best performance. We thus choose KLIVE as our model-checking technique for MLTL-SAT.

**Bounded MLTL-SAT.** Although MLTL-SAT is reducible to the satisfiability problem of other well-explored logics, with established off-the-shelf satisfiability solvers, a dedicated solution based on inherent properties of MLTL may be superior. One

---

[3] A temporary variable is introduced in the DEFINE statement rather than the VAR statement of the SMV model, as it will be automatically replaced with those in VAR statements.

**Table 1**

The SMV encoding for MLTL formula $\varphi$.

| |
|---|
| **VAR** |
| $\quad$ $a$: Boolean; //for each atom $a$ in $\varphi$ |
| $\quad$ ... |
| $\quad$ $Tail$: Boolean; //for $Tail$; |
| $\quad$ $\mathcal{X}\_\psi$: Boolean; //for each $\mathcal{U}$ and $\mathcal{R}$ formula in $cl^*(\varphi)$; |
| $\quad$ ... |
| $\quad$ $\mathcal{N}\_\psi$: Boolean; //for each $\mathcal{U}$ and $\mathcal{R}$ formula in $cl^*(\varphi)$; |
| $\quad$ ... |
| **INIT** |
| $\quad$ $e(\varphi)$; |
| **DEFINE** |
| $\quad$ $T\_\psi := \mathcal{X}\_(\psi_1 U_{[a-1,b-1]}\psi_2)$; // for $\psi_1\mathcal{U}_{[a,b]}\psi_2$ and $b \geq a > 0$ |
| $\quad$ ... |
| **INVAR** // for MLTL$_0$ encoding only |
| $\quad$ $T\_(\psi_1\mathcal{U}_{[0,a]}\psi_2) \wedge T\_(\psi_1\mathcal{U}_{[0,a-1]}\psi_2) \leftrightarrow T\_(\psi_1\mathcal{U}_{[0,a-1]}\psi_2)$ && |
| $\quad$ ... |
| **TRANS** |
| $\quad$ $(\mathcal{X}\_\psi \leftrightarrow (\neg Tail \wedge next(e(\psi))))$ && |
| $\quad$ ... && |
| $\quad$ $(\mathcal{N}\_\psi \leftrightarrow (Tail \vee next(e(\psi))))$ && |
| $\quad$ ... && TRUE; |
| **LTLSPEC** |
| $\quad$ $\square\neg Tail$; |
| **FAIRNESS** TRUE |

intuition is, since all intervals in MLTL formulas are bounded, the satisfiability of the formula can be reduced to Bounded Model Checking (BMC) [36].

**Theorem 5.** *Given an MLTL formula $\varphi$ with $K$ as the largest natural in the intervals of $\varphi$, $\varphi$ is satisfiable iff there is a finite trace $\pi$ with $|\pi| \leq K \cdot |cl(\varphi)|$ such that $\pi \models \varphi$.*

**Proof.** From Lemma 1, there is an LTL$_f$ formula $\theta$ of $\varphi$, of size of $O(K \cdot |cl(\varphi)|)$, that is equivalent to $\varphi$. Moreover, $\theta$ contains only $\mathcal{X}$ and $\mathcal{N}$ temporal operators, the number of which is less than $K \cdot |cl(\varphi)|$. Let $T(\theta)$ be the set of temporal operators in $\theta$, $|T(\theta)|$ denote the size of $T(\theta)$, and nnf$(\theta)$ be the NNF (Negation Normal Form) of $\theta$. An LTL$_f$ formula is in NNF if every negation operator $\neg$ appears only in front of atoms of the formula. For the LTL$_f$ formula $\theta$, there is a NNF MLTL formula nnf$(\theta)$ such that $\theta \equiv$ nnf$(\theta)$, where nnf$(\theta)$ can be obtained by making use of the dual operators. Consider $\theta = \neg(a \wedge (b\mathcal{U}c))$, nnf$(\theta)$ is $(\neg a) \vee ((\neg b)\mathcal{R}(\neg c))$. Moreover, the conversion cost is linear to the size of $\theta$.

By construction, nnf$(\theta) \equiv \theta$ and $|T(\theta)| = |T(\text{nnf}(\theta))|$ are true. We now prove that, for a finite trace $\xi \models$ nnf$(\theta)$, there is a prefix $\xi'$ of $\xi$ such that $\xi' \models$ nnf$(\theta)$ and $|\xi'| \leq |T(\text{nnf}(\theta))| + 1$. If $|\xi| \leq |T(\text{nnf}(\theta))| + 1$, then $\xi'$ is $\xi$ itself. So we only need to consider the situation when $|\xi| > |T(\text{nnf}(\theta))| + 1$.

- If nnf$(\theta)$ is a literal, $\xi \models$ nnf$(\theta)$ implies $\xi[0] \models$ nnf$(\varphi)$. Let $\xi' = \xi[0]$ and it is true that $|\xi'| \leq |T(\text{nnf}(\theta))| + 1 = 1$;
- If nnf$(\theta) = \psi_1 \wedge \psi_2$, $\xi \models$ nnf$(\theta)$ implies $\xi \models \psi_1$ and $\xi \models \psi_2$. By induction, there are $\eta$ and $\eta'$, which are prefixes of $\xi$ such that $\eta \models \psi_1$, $|\eta| \leq |T(\psi_1)| + 1$ and $\eta' \models \psi_2$, $|\eta'| \leq |T(\psi_2)| + 1$. Assume wlog that $|\eta| \geq |\eta'|$, and let $\xi' = \eta$. We know that $\xi' \models$ nnf$(\theta)$ and $|\xi'| = |T(\psi_1)| + 1 \leq |T(\text{nnf}(\theta))| + 1$ is true. The proof is analogous if nnf$(\theta) = \psi_1 \vee \psi_2$;
- If nnf$(\theta) = \mathcal{X}\psi$, $\xi \models$ nnf$(\theta)$ implies that $\xi_1 \models \psi$. By there is a prefix $\xi_1'$ of $\xi_1$ such that $\xi_1' \models \psi$ and $|\xi_1'| \leq |T(\psi)| + 1$. Let $\xi' = \xi[0] \cdot \xi_1'$, and we know that $\xi' \models$ nnf$(\theta)$ is true, and $|\xi'| = |T(\psi)| + 1 + 1 \leq |T(\text{nnf}(\theta))| + 1$;
- If nnf$(\theta) = \mathcal{N}\psi$, and since we only consider the case when $|\xi| > |T(\text{nnf}(\theta))| + 1$, we have that $\xi \models$ nnf$(\theta)$ implies that $\xi_1 \models \psi$. As a result, the proof for the case of $\mathcal{N}$ formula is the same as that of $\mathcal{X}$ formula.

Since we proved that $\xi \models$ nnf$(\theta)$ implies there is a prefix $\xi'$ of $\xi$ such that $\xi' \models$ nnf$(\theta)$ and $|\xi'| \leq |T(\text{nnf}(\theta)|) + 1$; it is also true that $\xi \models \theta$ implies there is a prefix $\xi'$ of $\xi$ such that $\xi' \models \theta$ and $|\xi'| \leq |T(\theta)| + 1 \leq K \cdot |cl(\varphi)|$; and thus we prove that $\xi \models \varphi$ implies there is a prefix $\xi'$ of $\xi$ such that $\xi' \models \varphi$ and $|\xi'| \leq K \cdot |cl(\varphi)|$. That means, whenever $\varphi$ is satisfiable, there is a trace $\xi' \models \varphi$ with the size bounded by $K \cdot |cl(\varphi)|$. $\square$

Theorem 5 states that the satisfiability of a given MLTL formula can be reduced to checking for the existence of a satisfying trace. To apply the BMC technique in nuXmv, we compute and set the maximal depth of BMC to be the value of $K \cdot |cl(\varphi)|$ for a given MLTL formula $\varphi$. The input SMV model for BMC is still $M_\varphi$, as described in Section 4.2. **However**

**to ensure correct BMC checking in nuXmv, the constraint "FAIRNESS TRUE" has to be added into the SMV model.**[4] The LTLSPEC remains $\Box\neg Tail$. According to Theorem 5, $\varphi$ is satisfiable iff the model checker returns a counterexample by using the BMC technique within the maximal depth of $K \cdot |cl(\varphi)|$.

### 4.3. MLTL-SAT via SMT solving

Another approach to solve MLTL-SAT is via SMT solving, considering that using SMT solvers to handle intervals in MLTL formulas is straightforward. Since the input logic of SMT solvers is First-Order Logic, we must first translate the MLTL formula to its equi-satisfiable formula in First-Order Logic over the natural domain $N$. We assume that readers are familiar with First-Order Logic and only focus on the translation. Given an MLTL formula $\varphi$ and the alphabet $\Sigma$, we construct the corresponding formula in First-Order Logic over $N$ in the following way.

1. For each $p \in \Sigma$, define a corresponding function $f_p : Int \rightarrow Bool$ such that $f_p(k)$ is true ($k \in N$) iff there is a satisfying (finite) trace $\pi$ of $\varphi$ and $p$ is in $\pi[k]$.
2. The First-Order Logic formula $\text{fol}(\varphi, k, len)$ for $\varphi$, where $k, len \in N$, is constructed recursively as below:
   - $\text{fol}(true, k, len) = (len > k)$ and $\text{fol}(false, k, len) = false$;
   - $\text{fol}(p, k, len) = (len > k) \wedge f_p(k)$ for $p \in \Sigma$;
   - $\text{fol}(\neg\xi, k, len) = (len > k) \wedge \neg\text{fol}(\xi, k, len)$;
   - $\text{fol}(\xi \wedge \psi, k, len) = \text{fol}(\xi, k, len) \wedge \text{fol}(\psi, k, len)$;
   - $\text{fol}(\xi \, \mathcal{U}_{[a,b]} \, \psi, k, len) = \exists i.( \, (a + k \leq i \leq b + k) \wedge \text{fol}(\psi, i, len) \wedge \forall j.((a + k \leq j < i) \rightarrow \text{fol}(\xi, j, len)))$.

In the formula $\text{fol}(\varphi, k, len)$, $k$ represents the index of the (finite) trace from which $\varphi$ is evaluated, and $len$ indicates the length of the trace satisfying $\varphi$ at the index $k$. Since the formula is constructed recursively, we need to introduce $k$ to record the index. Meanwhile, $len$ is necessary because the MLTL semantics, which is interpreted over finite traces, constrains the lengths of the satisfying traces of the Until formulas. The following theorem guarantees that MLTL-SAT is reducible to the satisfiability of First-Order Logic.

**Theorem 6.** *For an MLTL formula $\varphi$, $\varphi$ is satisfiable iff the corresponding First-Order Logic formula $\exists len.\text{fol}(\varphi, 0, len)$ is satisfiable.*

**Proof.** Let the alphabet of $\varphi$ be $\Sigma$, and $\pi \in (2^{\Sigma})^*$ be a finite trace. For each $p \in \Sigma$, we define the function $f_p : Int \rightarrow Bool$ as follows: $f_p(k) = true$ iff $p \in \pi[k]$ if $0 \leq k < |\pi|$. We now prove by induction over the type of $\varphi$ and the construction of $\text{fol}(\varphi, k, len)$ with respect to $\varphi$ that $\pi_k \models \varphi$ holds iff $\{f_p | p \in \Sigma\}$ is a model of $\text{fol}(\varphi, k, |\pi|)$: here $|\pi|$ is the length of $\pi$. The cases when $\varphi$ is true or false are trivial.

- If $\varphi = p$ is an atom, $\pi_k \models \varphi$ holds iff $p \in \pi[k]$ (i.e., $\pi_k[0]$) is true, which means $f_p(k) = true$. As a result, $\{f_p\}$ is a model of $\text{fol}(\varphi, k, |\pi|)$, which implies that $\pi_k \models \varphi$ holds iff $\{f_p | p \in \Sigma\}$ is a model of $\text{fol}(\varphi, k, |\pi|)$.
- If $\varphi = \neg\xi$, $\pi_k \models \varphi$ holds iff $\pi_k \not\models \xi$ holds. By hypothesis assumption, $\pi_k \models \xi$ holds iff $\{f_p | p \in \Sigma\}$ is a model of $\text{fol}(\xi, k, |\pi|)$, which is equivalent to saying $\pi_k \not\models \xi$ holds iff $\{f_p | p \in \Sigma\}$ is not a model of $\text{fol}(\xi, k, |\pi|)$. As a result, $\pi_k \models \neg\xi$ holds iff $\{f_p | p \in \Sigma\}$ is a model of $\neg\text{fol}(\xi, k, |\pi|)$.
- If $\varphi = \xi \wedge \psi$, $\pi_k \models \varphi$ holds iff $\pi_k \models \xi$ and $\pi_k \models \psi$. By hypothesis assumption, $\pi_k \models \xi$ (resp. $\pi_k \models \psi$) holds iff $\{f_p | p \in \Sigma\}$ is a model of $\text{fol}(\xi, k, |\pi|)$ (resp. $\text{fol}(\psi, k, |\pi|)$). According to the construction of the fol function, $\{f_p | p \in \Sigma\}$ is a model of $\text{fol}(\xi \wedge \psi, k, |\pi|)$. As a result, $\pi_k \models \xi \wedge \psi$ holds iff $\{f_p | p \in \Sigma\}$ is a model of $\text{fol}(\xi \wedge \psi, k, |\pi|)$.
- If $\varphi = \xi \, \mathcal{U}_{[a,b]} \, \psi$, $\pi_k \models \varphi$ holds iff there is $a + k \leq i \leq b + k$ such that $\pi_i \models \psi$ and $\pi_j \models \xi$ holds for every $a + k \leq j < i$. By hypothesis assumption, $\pi_i \models \psi$ holds iff $\{f_p | p \in \Sigma\}$ is a model of $\text{fol}(\psi, i, len)$, and $\pi, j \models \xi$ holds iff $\{f_p | p \in \Sigma\}$ is a model of $\text{fol}(\xi, j, |\pi|)$. Moreover, $|\pi| > a + k$ must be true according to the MLTL semantics. As a result, $\{f_p | p \in \Sigma\}$ is a model of $\text{fol}(\varphi, k, |\pi|)$, which implies that $\pi_k \models \xi \, \mathcal{U}_{[a,b]}\psi$ holds iff $\{f_p | p \in \Sigma\}$ is a model of $\text{fol}(\xi \, \mathcal{U}_{[a,b]} \, \psi, k, |\pi|)$.

This proof holds for all values of $k$, including the special case where $k = 0$. □

We then encode $\exists len.\text{fol}(\varphi, 0, len)$ into the SMT-LIB v2 format [37], which is the input of most modern SMT solvers; we call the full SMT-LIB v2 encoding $\text{SMT}(\varphi)$. We first use the "declare-fun" command to declare a function $f_a : Int \rightarrow Bool$ for each $p \in \Sigma$. We also define the function $f_\varphi : Int \times Int \rightarrow Bool$ for the First-Order Logic formula $\text{fol}(\varphi, k, len)$. The corresponding SMT-LIB v2 command is "define-fun $f_\varphi$ ((k Int) (len Int)) Bool $S(\text{fol}(\varphi, k, len))$", where $S(\text{fol}(\varphi, k, len))$ is the SMT-LIB v2 implementation of $\text{fol}(\varphi, k, len)$. In detail, $S(\text{fol}(\varphi, k, len))$ is acquired recursively as follows.

- $S(\text{fol}(p, k, len)) \longrightarrow (\text{and } (> len \ k) \ (f_p \ k))$

---

[4] Based on comments in emails from the nuXmv developers which are quoted as below: "... there is a known issue with BMC and LTL properties, namely that the BMC command currently implemented in nuXmv doesn't ensure that the returned counterexamples of properties of the form 'G formula' are infinite traces... just adding a 'FAIRNESS TRUE' to the model should make BMC work as expected.

**Table 2**

The SMT-LIB v2 template for SMT($\varphi$).

| |
|---|
| (declare-fun $f_a$ (Int) Bool) //declare corresponding function for $a \in \Sigma$ |
| $\dots$ |
| //define function for fol($\varphi, k, len$) |
| (define-fun $f_\varphi$ ((k Int) (len Int)) Bool $S$(fol($\varphi, k, len$))) |
| (assert (exists ((len Int)) ($f_\varphi$ 0))) |
| (check-sat) |

- $S(\neg \text{fol}(\varphi, k, len)) \longrightarrow$ (and ($>$ len k) (not $S$(fol($\varphi, k$))))
- $S(\text{fol}(\varphi_1 \land \psi, k, len)) \longrightarrow$ (and $S$(fol($\varphi_1, k, len$)) $S$(fol($\psi, k, len$)))
- $S(\text{fol}(\varphi_1 \ \mathcal{U}_{[a,b]} \ \psi, k, len)) \longrightarrow$ (exists ($i$ $Int$) (and ($\leq$ (+ $a$ $k$) $i$) ($\geq$ $i$ (+ $b$ $k$)) $S$(fol($\psi, i, len$)) (forall ($j$ $Int$) ($\Rightarrow$ (and ($\leq$ (+ $a$ $k$) $j$) ($<$ $j$ $i$)) $S$(fol($\varphi_1, j, len$))))))

In the above, we introduce the "forall" and "exists" quantifiers to encode the Until operator in a natural way. Considering that all $U$ operators are bounded in MLTL, the quantifiers can be erased from the encoding in a similar way as what we define the formula $f(\varphi)$ in Lemma 1. However, from our preliminary tests, such quantifier-free encoding has the following two drawbacks: (1) the constructed FOL formula can become too large to read as the bounds of the temporal operators increase. And (2) the performance for the quantifier-free formula seems not as competitive as that for the formula with quantifiers above, as the bounds become larger. The conjecture is that there are optimizations inside the SMT solver which speedup the checking process. Therefore, we choose the encoding way with quantifiers.

Finally, we use the "assert" command "(assert (exists (($len$ Int)) ($f_\varphi$ 0 len)))" together with the "(check-sat)" command to request SMT solvers for the satisfiability of $\exists len.\text{fol}(\varphi, 0, len)$. In a nutshell, the general framework of the SMT-LIB v2 format for SMT($\varphi$) (i.e., $\exists len.\text{fol}(\varphi, 0, len)$) is shown in Table 2, and the correctness is guaranteed by Theorem 7 below.

**Example 1.** The SMT encoding for $(F[0, 10001]\neg p) \land (G[0, 10000]p)$ is shown as below:

```
(declare-fun f_p (Int) Bool)
(declare-fun f (Int) Bool)
(assert (= (f 0)
          (and
            (forall ((x Int))
              (implies
                (and
                 (<= 0 x)
                 (<= x 10000)
                )
                (f_p x)
              )
            )
            (exists ((x Int))
              (and
                (and
                 (<= 0 x)
                 (<= x 10001)
                )
                ( not (f_p x) )
              )
            )
          )
      )
 )
(assert (f 0))
(check-sat)
```

**Theorem 7.** *The First-Order Logic formula* $\exists len. f ol(\varphi, 0, len)$ *is satisfiable iff the SMT solver returns SAT with the input SMT($\varphi$).*

An inductive proof for the theorem can be conducted according to the construction of SMT($\varphi$). Notably, there is no difference between the SMT encoding for MLTL formulas and that for MLTL$_0$ formulas, as the SMT-based encoding does not require unrolling the temporal operators in the formula.

*An alternative SMT encoding.* Since SMT is essentially a combination of different theories, the performance of SMT solving may hence vary on the theories used by the solver. Consider the encoding $S(\text{fol}(\varphi, 0, len))$ above, an alternative to the

**Table 3**

List of solvers and their runtime flags.

| Encoding | MLTLconverter flag | Solver | Solver flag |
|---|---|---|---|
| LTL | -ltl | aalta | default |
| LTL$_f$ | -ltlf | aaltaf | default |
| SMV | -smv | nuXmv | -source bmc.cmd (BMC) <br> -source klive.cmd (KLIVE) |
| SMT-LIB v2 | -smtlib | Z3 | -smt2 |

*read_model*

*flatten_hierarchy*                          *read_model*

*encode_variables*                           *flatten_hierarchy*

*build_boolean_model*                        *encode_variables*

*bmc_setup*                                   *build_boolean_model*

*go_bmc*                                      *check_ltlspec_klive -d*

*check_ltlspec_bmc -k MAX*                    *quit*

*quit*

**Fig. 1.** nuXmv commands for BMC (left) and KLIVE (right).

function for each variable in $\Sigma$ is to use an array instead, in which the Array (or ArrayEx) theory can apply. Therefore, each $f_a$ in the previous encoding corresponds to an array $A_a$ in the alternative one, with $(f_a\ k)$ being replaced by "(select $A_a$ k)". It is interesting to explore how different SMT theories affect the satisfiability-checking performance, and we will answer this question in the next section.

## 5. Experimental evaluations

**Tools and platform.** We implemented the translator MLTLconverter in C++, including encodings for an MLTL formula as equi-satisfiable LTL and LTL$_f$ formulas, and corresponding SMV and SMT-LIB v2 models. We leverage the extant LTL solver aalta [33], LTL$_f$ solver aaltaf [30], SMV model checker nuXmv [34], and the SMT solver Z3 [38] to check the satisfiability of the input MLTL formula in their respective encodings from MLTLconverter. The solvers, including the runtime flags we used, are summarized in Table 3. We evaluated both BMC and KLIVE [24] model-checking back-ends in nuXmv, and the corresponding commands are shown in Fig. 1. Notably in the figure, the maximal length "*MAX*" to run BMC is computed dynamically for each MLTL formula, based on Theorem 5.

All experiments were executed on Rice University's NOTS cluster,[5] running RedHat 5, with 226 dual socket compute blades housed within HPE s6500, HPE Apollo 2000, and Dell PowerEdge C6400 chassis. All the nodes are interconnected with 10 GigE network. Each satisfiability check over one MLTL formula and one solver was executed with exclusive access to one CPU and 8 GB RAM with a timeout of one hour, as measured by the Linux `time` command. We assigned a time penalty of one hour to benchmarks that segmentation fault or timeout.

**Experimental goals.** We evaluate performance along three metrics. (1) Each satisfiability check has two parts: the encoding time (consumed by MLTLconverter) and the solving time (consumed by solvers). We evaluate how each encoding affects the performance of both stages of MLTL-SAT. (2) We comparatively analyze the performance and scalability of end-to-end MLTL-SAT via LTL-SAT, LTL$_f$-SAT, LTL model checking, and our new SMT-based approach. (3) We evaluate the performance and scalability for MLTL$_0$ satisfiability checking using MLTL$_0$-SAT encoding heuristics (Lemma 2).

**Benchmarks.** There are few MLTL (or even MTL-over-naturals) benchmarks available for evaluation. Previous works on MTL-over-naturals [29,2,26] mainly focus on the theoretic exploration of the logic. To enable rigorous experimental evaluation, we develop three types of benchmarks, motivated by the generation of LTL benchmarks [15].[6]

1. *Random MLTL formulas (R)*: We generated 10,000 R formulas, varying the formula length $L$ (20, 40, 60, 80, 100), the number of variables $N$ (1, 2, 3, 4, 5), and the probability of the appearance of the $\mathcal{U}$ operator $P$ (0.33, 0,5, 0.7, 0.95); for each $(L, N, P)$ we generated 100 formulas. For every $\mathcal{U}$ operator, we randomly chose an interval $[i, j]$ where $i \geq 0$ and $j \leq 100$.
2. *NASA-Boeing MLTL formulas (NB)*: We use challenging benchmarks [39] created from projects at NASA [40,41] and Boeing [42]. We extract 63 real-life LTL requirements from the SMV models of the benchmarks, and then randomly generate an

---

**Table 4**

**LTL$_f$-Specific Benchmarks**: formulas specifically designed for $LTL_f$ from previous works, adapted to be benchmarks for our experiments. To create benchmarks from Declare Templates, we substituted variables for branches, then created formula-generating scripts. Notably, $\mathcal{W}$ is noted as *Weak Until*, whose semantics can be represented as $a\mathcal{W}b \equiv (a\mathcal{U}b|\Box a)$.

| Name | LTL$_f$ formalization | Description | Answer |
|---|---|---|---|
| Declare Patterns | | From [43] | sat* |
| Existence | $\Diamond a$ | $a$ must be executed at least once | sat* |
| Absence 2 | $\neg\Diamond(a \wedge \Diamond a)$ | $a$ can be executed at most once | sat* |
| Choice | $\Diamond a \vee \Diamond b$ | $a$ or $b$ must be executed | sat* |
| Exclusive Choice | $(\Diamond a \vee \Diamond b) \wedge \neg(\Diamond a \wedge \Diamond b)$ | Either $a$ or $b$ must be executed, but not both | sat* |
| Resp. existence | $\Diamond a \rightarrow \Diamond b$ | If $a$ is executed, then $b$ must be executed as well | sat* |
| Coexistence | $(\Diamond a \rightarrow \Diamond b) \wedge (\Diamond b \rightarrow \Diamond a)$ | Either $a$ and $b$ are both executed, or none of them is executed | sat* |
| Response | $\Box(a \rightarrow \Diamond b)$ | Every time $a$ is executed, $b$ must be executed afterwards | sat* |
| Precedence | $\neg b \mathcal{W} a$ | $b$ can be executed only if $a$ has been executed before | sat* |
| Succession | $\Box(a \rightarrow \Diamond b) \wedge (\neg b \mathcal{W} a)$ | $b$ must be executed after $a$, and $a$ must precede $b$ | sat* |
| Alt. Response | $\Box(a \rightarrow \mathcal{X}(\neg a U b))$ | Every $a$ must be followed by $b$, without any other $b$ in between | sat* |
| Alt. Precedence | $(\neg b \mathcal{W} a) \wedge \Box(b \rightarrow \mathcal{X}(\neg b \mathcal{W} a))$ | Every $b$ must be preceded by $a$, without any other $b$ in between | sat* |
| Alt. Succession | $\Box(a \rightarrow \mathcal{X}(\neg a U b)) \wedge (\neg b \mathcal{W} a) \wedge \Box(b \rightarrow \mathcal{X}(\neg \mathcal{W} a))$ | Combination of alternate response and alternate precedence | sat* |
| Chain Response | $\Box(a \rightarrow \mathcal{X} b)$ | If $a$ is executed then $b$ must be executed next | sat* |
| Chain Precedence | $\Box(\mathcal{X} b \rightarrow a)$ | Task $b$ can be executed only immediately after $a$ | sat* |
| Chain Succession | $\Box(a \leftrightarrow \mathcal{X} b)$ | Tasks $a$ and $b$ must be executed next to each other | sat* |
| Not Coexistence | $\neg(\Diamond a \wedge \Diamond b)$ | Only one among tasks $a$ and $b$ can be executed, but not both | sat* |
| Neg. Succession | $\Box(a \rightarrow \neg\Diamond b)$ | Task $a$ cannot be followed by $b$, and $b$ cannot be preceded by $a$ | sat* |
| Neg. Chain Succession | $\Box(a \leftrightarrow \mathcal{X}\neg b)$ | Tasks $a$ and $b$ cannot be executed next to each other | sat* |
| End | $\Diamond(a \wedge \neg\mathcal{X}(a \vee \neg a))$ | $a$ occurs *last*; translated to $LTL_f$ from [44] | sat* |
| Declare Templates | | formula-generating code inspired by constraints from [45] | sat* |
| RespondedExistence (n) | $\Diamond x \rightarrow \Diamond\left(\bigvee_{i=1}^{n} y_i\right)$ | | sat* |
| Response (n) | $\Box\left(x \rightarrow \Diamond\left(\bigvee_{i=1}^{n} y_i\right)\right)$ | | sat* |
| AlternateResponse(n) | $\Box\left(x \rightarrow \mathcal{X}\left(\neg x\mathcal{U}\bigvee_{i=1}^{n} y_i\right)\right)$ | | sat* |
| ChainResponse(n) | $\Box\left(x \rightarrow \mathcal{X}\left(\bigvee_{i=1}^{n} y_i\right)\right)$ | | sat* |
| Precedence(n) | $(\neg x)\mathcal{W}\left(\bigvee_{i=1}^{n} y_i\right)$ | | sat* |
| AlternatePrecedence(n) | $Precedence(n) \wedge \Box(x \rightarrow \mathcal{X}Precedence(n))$ | | sat* |
| ChainPrecedence(n) | $\Box\left((\mathcal{X} x) \rightarrow \left(\bigvee_{i=1}^{n} y_i\right)\right)$ | | sat* |

   interval for each temporal operator. (We replace each $\mathcal{X}$ with $\Box_{[1,1]}$.) We create 3 groups of such formulas (63 in each) to test the scalability of different approaches, by restricting the maximal number of the intervals to be 1,000, 10,000, and 100,000 respectively.

3. *Random MLTL$_0$ formulas (R0)*: We generated 500 R0 formulas in the same way as the R formulas, except that every generated interval was restricted to start from 0; we generated sets of five for each $(L, N, P)$. This small set of R benchmarks serve to compare the performance on MLTL$_0$ formulas whose SMV encodings were created with/without heuristics.

4. Unsatisfiable Random Conjunctive formulas (RC): Our preliminary evaluations show that 98% of the formulas collected in the above three benchmarks are satisfiable, which makes it hard to evaluate different approaches on unsatisfiability checking. Inspired from the fact shown in [30] that large conjunctive formulas tends to be unsatisfiable, we construct *random conjunctive formulas* for MLTL as follows.

   • A random conjunctive MLTL formula with the length $n$ has the form of $\bigwedge_{1 \leq i \leq n} C_i$ where $C_i$ is a small MLTL patterns that are widely used in practice.

   • Despite little work has investigated the common-used MLTL formulas, there are 26 off-the-shelf LTL$_f$ patterns collected in Table 4. As a result, we construct the MLTL patterns in demand from the corresponding LTL$_f$ ones. Informally for each LTL$_f$ pattern in Table 4, we replace the $\mathcal{X}$ operator with $\Diamond_{[1,1]}$ and the $\Diamond$, $\Box$ and $\mathcal{U}$ operators with $\Diamond_{[l,h]}$, $\Box_{[l,h]}$ and $\mathcal{U}_{[l,h]}$ respectively by randomly choosing $l$ and $h$ such that $l \leq h$. Notably, the $\mathcal{W}$ operator shown in Table 4 can be replaced by $\Box$ and $\mathcal{U}$, i.e. $\xi\mathcal{W}\psi \equiv \Box\xi \vee \xi\mathcal{U}\psi$.

   • We originally generated three groups of 1,000 RC formulas each, varying the number of conjuncts $C$ (5, 10, 15, 20, 25), the number of variables $N$ (1, 2, 3, 4, 5), and the interval ranges $R$ ([0, 50], [0, 100], [0,500]); for each $(C, N, R)$ we generated 100 formulas. Recall that we aim to generate unsatisfiable MLTL formulas, so we first run a preliminary evaluation on these formulas and then select 800 unsatisfiable instances as our RC benchmark.

**Correctness checking.** We compared the verdicts from all solvers for every test instance and found no inconsistencies, excluding segmentation faults. This exercise aided with verification of our implementations of the translators, including diagnosing the need for including FAIRNESS TRUE in BMC models.

**Experimental results.** Fig. 2 compares encoding times for the R benchmark formulas. We find that (1) Encoding MLTL as either LTL and LTL$_f$ is not scalable even when the intervals in the formula are small; (2) The cost of MLTL-to-SMV encoding is comparable to that from MLTL to SMT-LIB v2. Although the cost of encoding MLTL as LTL/LTL$_f$ and SMV are
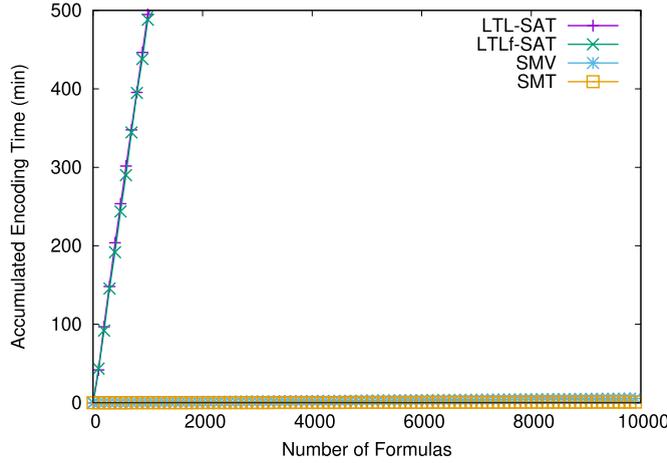
**Fig. 2.** Cactus plot for different MLTL encodings on R formulas: LTL-SAT and LTL$_f$-SAT lines overlap; SMV and SMT lines overlap.
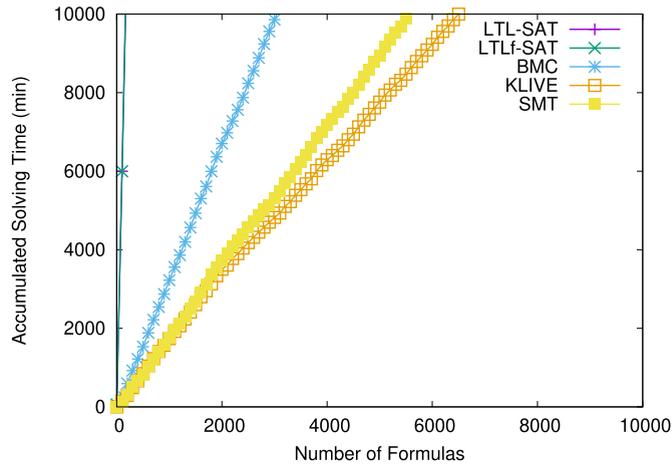


**Fig. 3.** Cactus plot for different MLTL solving approaches on R formulas: LTL-SAT and LTL$_f$-SAT lines overlap.

in $O(K \cdot |cl(\varphi)|)$, where $K$ is the maximal interval length in $\varphi$, the practical gap between the LTL/LTL$_f$ encodings and SMV encoding affirms our conjecture that the SMV model is more compact in general than the corresponding LTL/LTL$_f$ formulas. Also because $K$ is kept small in the R formulas, the encoding cost between SMV and SMT-LIB v2 becomes comparable.

Fig. 3 shows total satisfiability checking times for R benchmarks. Recall that the inputs of both BMC and KLIVE approaches are SMV models. The MLTL-SAT via KLIVE is the fastest solving strategy for MLTL formulas with interval ranges of less than 100. The portion of satisfiable/unsatisfiable formulas of this benchmark is approximate 4/1. Although BMC is known to be good at detecting counterexamples with short lengths, it does not perform as well as the KLIVE and SMT approaches on checking satisfiable formulas since only longer counterexamples (with length greater than 1000) exist for most of these formulas. While nuXmv successfully checked all such models, Fig. 4 shows that increasing the interval range constraint results in segmentation faults; more than half of our benchmarks produced this outcome for formulas with allowed interval ranges of up to 600. Meanwhile, the solving solutions via LTL-SAT/LTL$_f$-SAT are definitely not competitive for any interval range.

The SMT-based approach dominates the model-checking-approaches when considering scalable NB benchmarks, as shown in Fig. 5. Here, e.g., "BMC-1000" means using BMC to check the group of benchmarks with a maximal interval range of 1,000, and the analogous meaning applies to "KLIVE-1000". Since we consider two different SMT encodings in this paper, we use "Z3-F" to represent the encoding with the uninterpreted function theory and "Z3-A" for the one with the array theory. Due to segmentation faults, "BMC-1000" and "KLIVE-1000" have almost the same performance because the SMV models generated from our translator MLTLconverter are too large for nuXmv to handle. The performance of the model-checking approaches is constrained by the scalability of the model checker (nuXmv). However, the SMT encoding does not face such a bottleneck; see "Z3-F-1000," "Z3-F-10000," and "Z3-F-100000" in Fig. 5. We conclude that the SMT approach is the best available strategy for MLTL satisfiability checking.
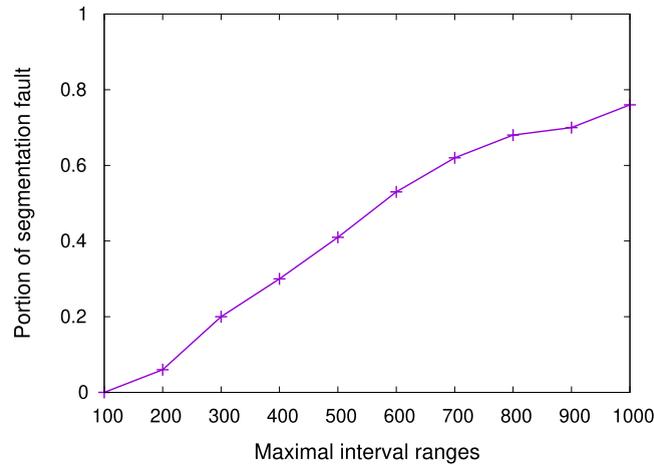
**Fig. 4.** Proportion of segmentation faults for sets of 200 R formulas with maximal interval ranges varying from 100 to 1000.
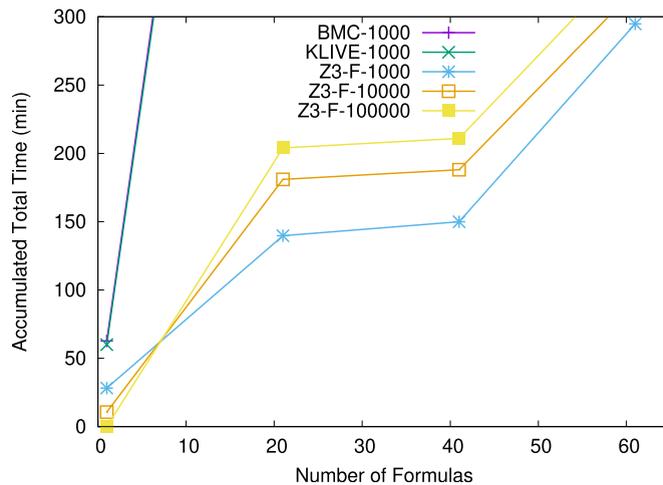


**Fig. 5.** Cactus plot for BMC, KLIVE and SMT-solving approaches on the NB benchmarks; BMC and KLIVE overlap.

We then evaluated the performance between the two different SMT encodings, as shown in Fig. 6. It turns out that there is no big performance gap between these two encodings: the encoding with the array theory performs only slightly less well than the one with the uninterpreted function theory. The conclusion that changing different encoding ways may affect the satisfiability-checking performance significantly, as shown in [17] for the SMV encodings, seems not directly applicable to the SMT encodings.

We also evaluated the performance of model-checking-based approaches on the R0 formulas, observing that there is an exponential complexity gap between MLTL-SAT and $MLTL_0$-SAT. Fig. 7 compares the performance of satisfiability solving via the BMC and KLIVE approaches. There is no significant improvement when the SMV encoding heuristics for $MLTL_0$ are applied. For the BMC solving approach, performance is largely unaffected by encoding heuristics. For the KLIVE solving approach, encoding heuristics decrease solving performance. The results support the well-known phenomenon that the theoretic analysis and the practical evaluations do not always match.

Finally, we compared different approaches on checking unsatisfiable random conjunction formulas, as shown in Fig. 8. The results indicate that the SMT approach performs best for checking unsatisfiability. The reason why there is a big performance gap between the other two approaches and the SMT ones is because the benchmarks contain those formulas whose interval ranges are greater than 100. Both the BMC and KLIVE solving techniques cannot perform well on the formulas whose interval ranges are greater than 100. However, the conclusion that the model-checking approach performs best still preserves on unsatisfiable formulas whose interval ranges are smaller than 100.

We summarize with the following five conclusions. (1) For satisfiability checking of MLTL formulas, the new SMT-based approach is best. (2) For satisfiability checking of MLTL formulas with interval ranges less than 100, the MLTL-SAT via KLIVE approach is fastest. (3) The above two observations on both satisfiable and unsatisfiable formulas. (4) The SMT encodings with different theories do not perform very differently in evaluation. (5) The dedicated encoding heuristics for $MLTL_0$ do not
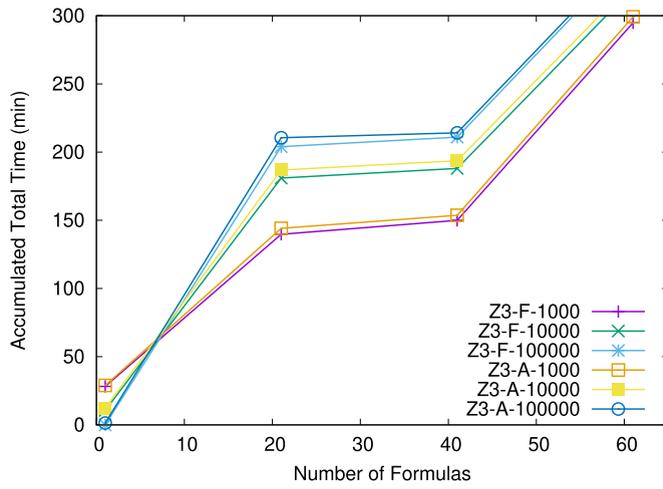
**Fig. 6.** Cactus plot for the two different SMT-solving approaches on the NB benchmarks.
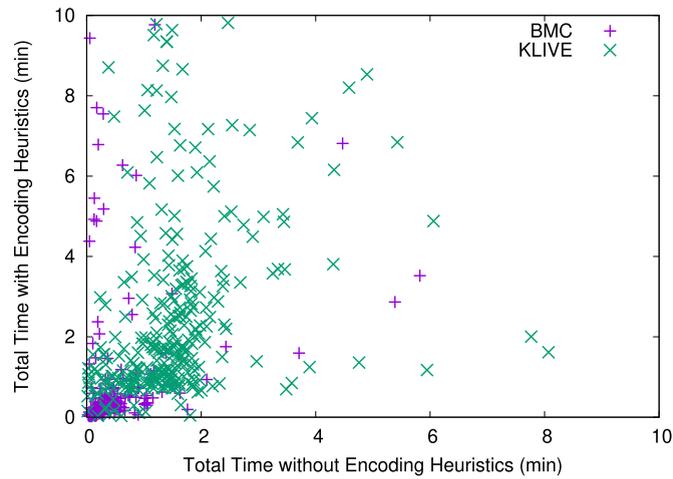


**Fig. 7.** Scatter plot for both the BMC and KLIVE approaches to check $MLTL_0$ formulas with/without encoding heuristics.
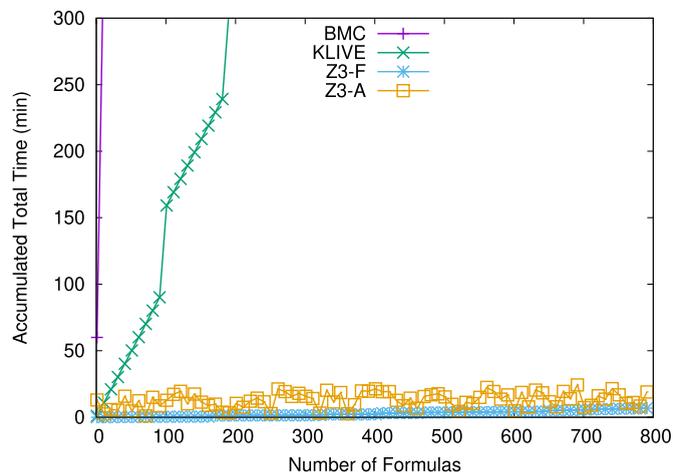


**Fig. 8.** Cactus plot for different approaches on the unsatisfiable random conjunction (RC) benchmarks.

significantly improve the satisfiability checking time of MLTL$_0$-SAT over MLTL-SAT. They do not solve the nuXmv scalability problem;

## 6. Discussion and conclusion

Metric Temporal Logic (MTL) was first introduced in [2], for describing continuous behaviors interpreted over infinite real-time traces. The later variants Metric Interval Temporal Logic (MITL) [46], and Bounded Metric Temporal Logic (BMTL) [47] are also interpreted over infinite traces. Intuitively, MLTL is a combination of MITL and BMTL that allows only bounded, discrete (over natural domain) intervals that are interpreted over finite traces. There are several previous works on the satisfiability of MITL, though their tools only support the infinite semantics. Bounded satisfiability checking for MITL formulas is proposed in [48], and the reduction from MITL to LTL is presented in [49]. Since previous works focus on MITL over infinite traces and there is no trivial way to reduce MLTL over finite traces to MITL over infinite traces, the previous methodologies are not comparable to those presented in this paper. This includes the SMT-based solution of reducing MITL formulas to equi-satisfiable Constraint LTL formulas [23]. Compared to that, our new SMT-based approach more directly encodes MLTL formulas into the SMT language without translation through an intermediate language.

The contribution of a complete, correct, and open-source MLTL satisfiability checking algorithm and tool opens up avenues for a myriad of future directions, as we have now made possible specification debugging MLTL formulas in design-time verification and benchmark generation for runtime verification. We plan to explore alternative encodings for improving the performance of MLTL satisfiability checking and work toward developing an optimized multi-encoding approach, following the style of the previous study for LTL [17]; the current SMT model generated from the MLTL formula uses a relatively simple theory (uninterpreted functions). We also plan to explore lazy encodings from MLTL formulas to SMT models. For example, instead of encoding the whole MLTL formula into a monolithic SMT model, we may be able to decrease overall satisfiability-solving time by encoding the MLTL formula in parts with dynamic ordering similar to [39]. To make the output of SMT-based MLTL satisfiability checking more usable, we plan to investigate translations from the functions returned from Z3 for satisfiable instances into more easily parsable satisfying assignments.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgment

## References

[1] T. Reinbacher, K.Y. Rozier, J. Schumann, Temporal-logic based runtime observer pairs for system health management of real-time systems, in: TACAS, in: LNCS, vol. 8413, Springer-Verlag, 2014, pp. 357–372.
[2] R. Alur, T.A. Henzinger, Real-time logics: complexity and expressiveness, in: LICS, IEEE, 1990, pp. 390–401.
[3] O. Maler, D. Nickovic, Monitoring temporal properties of continuous signals, in: FORMATS, Springer, 2004, pp. 152–166.
[4] G. De Giacomo, M. Vardi, Linear temporal logic and linear dynamic logic on finite traces, in: IJCAI, AAAI Press, 2013, pp. 2000–2007.
[5] A. Pnueli, The temporal logic of programs, in: IEEE FOCS, 1977, pp. 46–57.
[6] J. Geist, K.Y. Rozier, J. Schumann, Runtime observer pairs and bayesian network reasoners on-board fpgas: flight-certifiable system health management for embedded systems, in: RV, vol. 8734, Springer-Verlag, 2014, pp. 215–230.
[7] J. Schumann, K.Y. Rozier, T. Reinbacher, O.J. Mengshoel, T. Mbaya, C. Ippolito, Towards real-time, on-board, hardware-supported sensor and software health management for unmanned aerial systems, Int. J. Progn. Heal. Manag. 6 (1) (2015) 1–27.
[8] K.Y. Rozier, J. Schumann, C. Ippolito, Intelligent hardware-enabled sensor and software safety and health management for autonomous UAS, Technical Memorandum NASA/TM-2015-218817, NASA Ames Research Center, Moffett Field, CA 94035, May 2015.
[9] J. Schumann, P. Moosbrugger, K.Y. Rozier, R2U2: monitoring and diagnosis of security threats for unmanned aerial systems, in: RV, Springer-Verlag, 2015.
[10] J. Schumann, P. Moosbrugger, K.Y. Rozier, Runtime analysis with R2U2: a tool exhibition report, in: RV, Springer-Verlag, 2016.
[11] P. Moosbrugger, K.Y. Rozier, J. Schumann, R2U2: monitoring and diagnosis of security threats for unmanned aerial systems, in: FMSD, 2017, pp. 1–31.
[12] Runtime Verification Benchmark Competition, https://www.rv-competition.org/2018-2/, 2018.
[13] F.B. Kessler, nuXmv 1.1.0 (2016-05-10) Release Notes, https://es-static.fbk.eu/tools/nuxmv/downloads/NEWS.txt, 2016.
[14] K.Y. Rozier, Specification: the biggest bottleneck in formal methods and autonomy, in: VSTTE, in: LNCS, vol. 9971, Springer-Verlag, 2016, pp. 1–19.
[15] K. Rozier, M. Vardi, LTL satisfiability checking, in: SPIN, in: LNCS, vol. 4595, Springer, 2007, pp. 149–167.
[16] K. Rozier, M. Vardi, LTL satisfiability checking, Int. J. Softw. Tools Technol. Transf. 12 (2) (2010) 123–137.
[17] K. Rozier, M. Vardi, A multi-encoding approach for LTL symbolic satisfiability checking, in: 17th International Symposium on Formal Methods, FM2011, in: LNCS, vol. 6664, Springer-Verlag, 2011, pp. 417–431.

[18] K. Rozier, M. Vardi, Deterministic compilation of temporal safety properties in explicit state model checking, in: 8th Haifa Verification Conference, HVC2012, in: Lecture Notes in Computer Science (LNCS), vol. 7857, Springer-Verlag, 2012, pp. 243–259.

[19] R. Bloem, H. Chockler, M. Ebrahimi, O. Strichman, Synthesizing non-vacuous systems, in: A. Bouajjani, D. Monniaux (Eds.), Verification, Model Checking, and Abstract Interpretation, Springer International Publishing, Cham, 2017, pp. 55–72.

[20] R. Armoni, L. Fix, A. Flaisher, O. Grumberg, N. Piterman, M. Vardi, Enhanced vacuity detection for linear temporal logic, in: Computer Aided Verification, Proc. 15th International Conference, Springer, 2003, pp. 368–380.

[21] K.Y. Rozier, On the evaluation and comparison of runtime verification tools for hardware and cyber-physical systems, in: RV-CUBES, vol. 3, Kalpa Publications, 2017, pp. 123–137.

[22] J. Li, K.Y. Rozier, MLTL benchmark generation via formula progression, in: Runtime Verification, Springer International Publishing, 2018, pp. 426–433.

[23] M. Bersani, M. Rossi, P.S. Pietro, An SMT-based approach to satisfiability checking of MITL, Inf. Comput. 245 (C) (2015) 72–97, https://doi.org/10.1016/j.ic.2015.06.007.

[24] K. Claessen, N. Sörensson, A liveness checking algorithm that counts, in: FMCAD, IEEE, 2012, pp. 52–59.

[25] J. Li, M. Vardi, K. Rozier, Satisfiability checking for mission-time ltl, in: I. Dillig, S. Tasiran (Eds.), Computer Aided Verification, Springer International Publishing, Cham, 2019, pp. 3–22.

[26] R. Alur, T. Henzinger, A really temporal logic, J. ACM 41 (1) (1994) 181–204.

[27] C. Furia, P. Spoletini, Tomorrow and All Our Yesterdays: MTL Satisfiability over the Integers, Springer, 2008, pp. 126–140.

[28] A. Sistla, E. Clarke, The complexity of propositional linear temporal logic, J. ACM 32 (1985) 733–749.

[29] R. Alur, T. Feder, T. Henzinger, The benefits of relaxing punctuality, J. ACM 43 (1) (1996) 116–146.

[30] J. Li, L. Zhang, G. Pu, M.Y. Vardi, J. He, LTL$_f$ satisfibility checking, in: ECAI, 2014, pp. 91–98.

[31] G.D. Giacomo, M. Vardi, Synthesis for LTL and LDL on finite traces, in: IJCAI, 2015, pp. 1558–1564.

[32] P. Pandya, S. Shah, The unary fragments of metric interval temporal logic: bounded versus lower bound constraints, in: Int'l Symp. on Automated Technology for Verification and Analysis, Springer, 2012, pp. 77–91.

[33] J. Li, S. Zhu, G. Pu, M. Vardi, SAT-based explicit LTL reasoning, in: HVC, Springer, 2015, pp. 209–224.

[34] R. Cavada, A. Cimatti, M. Dorigatti, A. Griggio, A. Mariotti, A. Micheli, S. Mover, M. Roveri, S. Tonetta, The NuXMV symbolic model checker, in: CAV, 2014, pp. 334–342.

[35] K. McMillan, Symbolic model checking: an approach to the state explosion problem, Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA, USA, 1992, uMI. Order No. GAX92-24209.

[36] A. Biere, A. Cimatti, E. Clarke, Y. Zhu, Symbolic model checking without BDDs, in: TACAS, in: LNCS, vol. 1579, Springer, 1999.

[37] C. Barrett, A. Stump, C. Tinelli, The SMT-LIB standard: version 2.0, in: Workshop on Satisfiability Modulo Theories, 2010.

[38] L.D. Moura, N. Bjørner, Z3: an efficient SMT solver, in: TACAS, Springer-Verlag, 2008, pp. 337–340.

[39] R. Dureja, K.Y. Rozier, More scalable ltl model checking via discovering design-space dependencies (D3), in: D. Beyer, M. Huisman (Eds.), Tools and Algorithms for the Construction and Analysis of Systems, Springer International Publishing, Cham, 2018, pp. 309–327.

[40] M. Gario, A. Cimatti, C. Mattarei, S. Tonetta, K.Y. Rozier, Model checking at scale: automated air traffic control design space exploration, in: Proceedings of 28th International Conference on Computer Aided Verification, CAV 2016, in: LNCS, vol. 9780, Springer, Toronto, ON, Canada, 2016, pp. 3–22.

[41] C. Mattarei, A. Cimatti, M. Gario, S. Tonetta, K.Y. Rozier, Comparing different functional allocations in automated air traffic control design, in: Proceedings of Formal Methods in Computer-Aided Design, FMCAD 2015, IEEE/ACM, Austin, Texas, USA, 2015.

[42] M. Bozzano, A. Cimatti, A. Fernandes Pires, D. Jones, G. Kimberly, T. Petri, R. Robinson, S. Tonetta, Formal design and safety analysis of AIR6110 wheel brake system, in: CAV, 2015.

[43] G. De Giacomo, R. De Masellis, M. Montali, Reasoning on LTL on finite traces: insensitivity to infiniteness, in: AAAI, 2014, pp. 1027–1033.

[44] C.D. Ciccio, M. Mecella, On the discovery of declarative control flows for artful processes, ACM Trans. Manag. Inf. Syst. 5 (4) (2015) 24:1–24:37, https://doi.org/10.1145/2629447.

[45] C. Di Ciccio, F. Maggi, J. Mendling, Efficient discovery of target-branched declare constraints, Inf. Syst. 56 (C) (2016) 258–283, https://doi.org/10.1016/j.is.2015.06.009.

[46] R. Alur, T. Henzinger, Reactive modules, in: Proc. 11th IEEE Symp. on Logic in Computer Science, 1996, pp. 207–218.

[47] J. Ouaknine, J. Worrell, Some recent results in metric temporal logic, in: F. Cassez, C. Jard (Eds.), Formal Modeling and Analysis of Timed Systems, Springer, Berlin, Heidelberg, 2008, pp. 1–13.

[48] M. Pradella, A. Morzenti, P. Pietro, Bounded satisfiability checking of metric temporal logic specifications, ACM Trans. Softw. Eng. Methodol. 22 (3) (2013) 20:1–20:54.

[49] U. Hustadt, A. Ozaki, C. Dixon, Theorem proving for metric temporal logic over the naturals, in: L. de Moura (Ed.), Automated Deduction – CADE 26, Springer International Publishing, Cham, 2017, pp. 326–343.