11-7-2018

SAT-based Explicit LTLf Satisfiability Checking

Jianwen Li
Iowa State University, jianwen@iastate.edu

Kristin Y. Rozier
Iowa State University, kyrozier@iastate.edu

Geguang Pu
East China Normal University

Yueling Zhang
East China Normal University

Moshe Y. Vardi
Rice University

SAT-based Explicit LTLf Satisfiability Checking

Abstract

We present here a SAT-based framework for LTLf (Linear Temporal Logic on Finite Traces) satisfiability checking. We use propositional SAT-solving techniques to construct a transition system for the input LTLf formula; satisfiability checking is then reduced to a path-search problem over this transition system. Furthermore, we introduce CDLSC (Conflict-Driven LTLf Satisfiability Checking), a novel algorithm that leverages information produced by propositional SAT solvers from both satisfiability and unsatisfiability results. Experimental evaluations show that CDLSC outperforms all other existing approaches for LTLf satisfiability checking, by demonstrating an approximate four-fold speedup compared to the second-best solver.

Disciplines

Computer Sciences | Theory and Algorithms

Comments

This is a pre-print of the article Li, Jianwen, Kristin Y. Rozier, Geguang Pu, Yueling Zhang, and Moshe Y. Vardi. "SAT-based Explicit LTLf Satisfiability Checking." *arXiv preprint arXiv:1811.03176* (2018). Posted with permission.

SAT-based Explicit LTL_f Satisfiability Checking*

Jianwen Li, Kristin Y. Rozier
Iowa State University
Ames, IA, USA
{jianwen,kyrozier}@iastate.edu

Geguang Pu, Yueling Zhang
East China Normal University
Shanghai, China
{ggpu,ylzhang}@sei.ecnu.edu.cn

Moshe Y. Vardi
Rice University
Houston, TX, USA
vardi@cs.rice.edu

Abstract

We present here a SAT-based framework for LTL_f (Linear Temporal Logic on Finite Traces) satisfiability checking. We use propositional SAT-solving techniques to construct a transition system for the input LTL_f formula; satisfiability checking is then reduced to a path-search problem over this transition system. Furthermore, we introduce CDLSC (Conflict-Driven LTL_f Satisfiability Checking), a novel algorithm that leverages information produced by propositional SAT solvers from both satisfiability and unsatisfiability results. Experimental evaluations show that CDLSC outperforms all other existing approaches for LTL_f satisfiability checking, by demonstrating an approximate four-fold speedup compared to the second-best solver.

Introduction

Linear Temporal Logic over Finite Traces, or LTL_f, is a formal language gaining popularity in the AI community for formalizing and validating system behaviors. While standard Linear Temporal Logic (LTL) is interpreted on infinite traces (Pnueli 1977), LTL_f is interpreted over finite traces (De Giacomo and Vardi 2013). While LTL is typically used in formal-verification settings, where we are interested in nonterminating computations, cf. (Vardi 2007), LTL_f is more attractive in AI scenarios focusing on finite behaviors, such as planning (Bacchus and Kabanza 1998; De Giacomo and Vardi 1999; Calvanese *et al.* 2002; Patrizi *et al.* 2011; Camacho *et al.* 2017), plan constraints (Bacchus and Kabanza 2000; Gabaldon 2004), and user preferences (Bienvenu *et al.* ; Bienvenu *et al.* 2011; Sohrabi *et al.* 2011). Due to the wide spectrum of applications of LTL_f in the AI community (De Giacomo *et al.* 2014), it is worthwhile to study and develop an efficient framework for solving LTL_f-reasoning problems. Just as propositional satisfiability checking is one of the most fundamental propositional reasoning tasks, LTL_f satisfiability checking is a fundamental task for LTL_f reasoning.

Given an LTL_f formula, the satisfiability problem asks whether there is a finite trace that satisfies the formula. A "classical" solution to this problem is to reduce it to the LTL satisfiability problem (De Giacomo and Vardi 2013). The advantage of this approach is that the LTL satisfiability problem has been studied for at least a decade, and many mature tools are available, cf. (Rozier and Vardi 2007; Rozier and Vardi 2010). Thus, LTL_f satisfiability checking can benefit from progress in LTL satisfiability checking. There is, however, an inherent drawback that an extra cost has to be paid when checking LTL formulas, as the tool searches for a "lasso" (a lasso consists of a finite path plus a cycle, representing an infinite trace), whereas models of LTL_f formulas are just finite traces. Based on this motivation, (Li *et al.* 2014) presented a tableau-style algorithm for LTL_f satisfiability checking. They showed that the dedicated tool, *Aalta-finite*, which conducts an explicit-state search for a satisfying trace, outperforms extant tools for LTL_f satisfiability checking.

The conclusion of a dedicated solver being superior to LTL_f satisfiability checking from (Li *et al.* 2014), seems to be out of date by now because of the recent dramatic improvement in propositional SAT solving, cf. (Malik and Zhang 2009). On one hand, SAT-based techniques have led to a significant improvement on LTL satisfiability checking, outperforming the tableau-based techniques of *Aalta-finite* (Li *et al.* 2014). (Also, the SAT-based tool *ltl2sat* for LTL_f satisfiability checking outperforms *Aalta-finite* on particular benchmarks (Fionda and Greco 2016).) On the other hand, SAT-based techniques are now dominant in symbolic model checking (Cavada *et al.* 2014; Vizel *et al.* 2015). Our preliminary evaluation indicates that LTL_f satisfiability checking via SAT-based model checking (Bradley 2011; Een *et al.* 2011) or via SAT-based LTL satisfiability checking (Li *et al.* 2015) both outperform the tableau-based tool *Aalta-finite*. Thus, the question raised initially in (Rozier and Vardi 2007) needs to be re-opened with respect to LTL_f satisfiability checking: is it best to reduce to SAT-based model checking or develop a dedicated SAT-based tool?

Inspired by (Li *et al.* 2015), we present an explicit-state SAT-based framework for LTL_f satisfiability. We construct the LTL_f *transition system* by utilizing SAT solvers to compute the states explicitly. Furthermore, by making use of both satisfiability and unsatisfiability information from SAT solvers, we propose a *conflict-driven* algorithm, CDLSC, for efficient LTL_f satisfiability checking. We show that by specializing the transition-system approach of (Li *et al.* 2015) to LTL_f and its finite-trace semantics, we get a frame-

*Geguang Pu and Kristin Y. Rozier are corresponding authors.

work that is significantly simpler and yields a much more efficient algorithm CDLSC than the one in (Li *et al.* 2015).

We conduct a comprehensive comparison among different approaches. Our experimental results show that the performance of CDLSC dominates all other existing LTL_f-satisfiability-checking algorithms. On average, CDLSC achieves an approximate four-fold speed-up, compared to the second-best solution (IC3 (Bradley 2011)+K-LIVE (Claessen and Sörensson 2012)) tested in our experiments. Our results re-affirm the conclusion of (Li *et al.* 2014) that the best approach to LTL_f satisfiability solving is via a dedicated tool, based on explicit-state techniques.

LTL over Finite Traces

Given a set $\mathcal{P}$ of atomic propositions, an LTL_f formula ϕ has the form:

$$\phi ::= \mathsf{tt} \mid p \mid \neg\phi \mid \phi \wedge \phi \mid \mathcal{X}\phi \mid \phi\mathcal{U}\phi;$$

where tt is true, $\neg$ is the negation operator, $\wedge$ is the and operator, $\mathcal{X}$ is the strong Next operator and $\mathcal{U}$ is the Until operator. We also have the duals ff (false) for tt, $\vee$ for $\wedge$, $\mathcal{N}$ (weak Next) for $\mathcal{X}$ and $\mathcal{R}$ for $\mathcal{U}$. A *literal* is an atom $p \in \mathcal{P}$ or its negation ($\neg p$). Moreover, we use the notation $\mathcal{G}\phi$ (Globally) and $\mathcal{F}\phi$ (Eventually) to represent $\mathsf{ff}\mathcal{R}\phi$ and $\mathsf{tt}\mathcal{U}\phi$. Notably, $\mathcal{X}$ is the standard *next* operator, while $\mathcal{N}$ is *weak next*; $\mathcal{X}$ requires the existence of a successor state, while $\mathcal{N}$ does not. Thus $\mathcal{N}\phi$ is always true in the last state of a finite trace, since no successor exists there. This distinction is specific to LTL_f.

LTL_f formulas are interpreted over finite traces (De Giacomo and Vardi 2013). Given an atom set $\mathcal{P}$, we define $\Sigma = 2^{\mathcal{P}}$ be the family of sets of atoms. Let $\xi \in \Sigma^+$ be a finite nonempty trace, with $\xi = \sigma_0\sigma_1 \ldots \sigma_n$. we use $|\xi| = n + 1$ to denote the length of ξ. Moreover, for $0 \leq i \leq n$, we denote $\xi[i]$ as the i-th position of ξ, and ξ_i to represent $\sigma_i\sigma_{i+1} \ldots \sigma_n$, which is the suffix of ξ from position i. We define the satisfaction relation $\xi \models \phi$ as follows:

- $\xi \models \mathsf{tt}$; and $\xi \models p$, if $p \in \mathcal{P}$ and $p \in \xi[0]$;
- $\xi \models \neg\phi$, if $\xi \not\models \phi$;
- $\xi \models \phi_1 \wedge \phi_2$, if $\xi \models \phi_1$ and $\xi \models \phi_2$;
- $\xi \models \mathcal{X}\phi$ if $|\xi| > 1$ and $\xi_1 \models \psi$;
- $\xi \models (\phi_1\mathcal{U}\phi_2)$, if there exists $0 \leq i < |\xi|$ such that $\xi_i \models \phi_2$ and for every $0 \leq j < i$ it holds that $\xi_j \models \phi_1$;

Definition 1 (LTL_f Satisfiability Problem)**.** *Given an LTL_f formula ϕ over the alphabet Σ, we say ϕ is satisfiable iff there is a finite nonempty trace $\xi \in \Sigma^+$ such that $\xi \models \phi$.*

Notations. We use $cl(\phi)$ to denote the set of subformulas of ϕ. Let A be a set of LTL_f formulas, we denote $\bigwedge A$ to be the formula $\bigwedge_{\psi \in A}\psi$. The two LTL_f formulas ϕ_1, ϕ_2 are semantically equivalent, denoted as $\phi_1 \equiv \phi_2$, iff for every finite trace ξ, $\xi \models \phi_1$ iff $\xi \models \phi_2$. Obviously, we have $(\phi_1 \vee \phi_2) \equiv \neg(\neg\phi_1 \wedge \neg\phi_2)$, $\mathcal{N}\psi \equiv \neg\mathcal{X}\neg\psi$ and $(\phi_1\mathcal{R}\phi_2) \equiv \neg(\neg\phi_1\mathcal{U}\neg\phi_2)$.

We say an LTL_f formula ϕ is in *Tail Normal Form* (TNF) if ϕ is in *Negated Normal Form* (NNF) and $\mathcal{N}$-free. It is trivial to know that every LTL_f formula has an equivalent NNF. Assume ϕ is in NNF, $\mathsf{tnf}(\phi)$ is defined as $t(\phi) \wedge \mathcal{F}Tail$, where $Tail$ is a new atom to identify the last state of satisfying traces (Motivated from (De Giacomo and Vardi 2013)), and $t(\phi)$ is an LTL_f formula defined recursively as follows: (1) $t(\phi) = \phi$ if ϕ is tt, ff or a literal; (2) $t(\mathcal{X}\psi) = \neg Tail \wedge \mathcal{X}(t(\psi))$; (3) $t(\mathcal{N}\psi) = Tail \vee \mathcal{X}(t(\psi))$; (4) $t(\phi_1 \wedge \phi_2) = t(\phi_1) \wedge t(\phi_2)$; (5) $t(\phi_1 \vee \phi_2) = t(\phi_1) \vee t(\phi_2)$; (6) $t(\phi_1\mathcal{U}\phi_2) = (\neg Tail \wedge t(\phi_1))\mathcal{U}t(\phi_2)$; (7) $t(\phi_1\mathcal{R}\phi_2) = (Tail \vee t(\phi_1))\mathcal{R}t(\phi_2)$.

Theorem 1. *ϕ is satisfiable iff $\mathsf{tnf}(\phi)$ is satisfiable.*

In the rest of the paper, unless clearly specified, the input LTL_f formula is in TNF.

Approach Overview

There is a Non-deterministic Finite Automaton (NFA) $\mathcal{A}_\phi$ that accepts exactly the same language as an LTL_f formula ϕ (De Giacomo and Vardi 2013). Instead of constructing the NFA for ϕ, we generate the corresponding *transition system* (Definition 5), by leveraging SAT solvers. The transition system can be considered as an intermediate structure of the NFA, in which every state consists of a set of subformulas of ϕ.

The classic approach to generate the NFA from an LTL_f formula, i.e. Tableau Construction (Gerth *et al.* 1995), creates the set of all one-transition next states of the current state. However, the number of these states is extremely large. To mitigate the overload, we leverage SAT solvers to compute the next states of the current state iteratively. Although both approaches share the same worst case (computing all states in the state space), our new approach is better for on-the-fly checking, as it computes new states only if the satisfiability of the formula cannot be determined based on existing states.

We show the SAT-based approach via an example. Consider the formula $\phi = (\neg Tail \wedge a)\mathcal{U}b$. The initial state s_0 of the transition system is $\{\phi\}$. To compute the next states of s_0, we translate ϕ to its equivalent *neXt Normal Form* (XNF), e.g. $\mathsf{xnf}(\phi) = (b \vee ((\neg Tail \wedge a) \wedge \mathcal{X}\phi))$, see Definition 4. If we replace $\mathcal{X}\phi$ in $\mathsf{xnf}(\phi)$ with a new propositions p_1, the new formula, denoted $\mathsf{xnf}(\phi)^p$, is a pure Boolean formula. As a result, a SAT solver can compute an assignment for the formula $\mathsf{xnf}(\phi)^p$. Assume the assignment is $\{a, \neg b, \neg Tail, p_1\}$, then we can induce that $(a \wedge \neg b \wedge \neg Tail \wedge \mathcal{X}\phi) \Rightarrow \phi$ is true, which indicates $\{\phi\} = s_0$ is a one-transition next state of s_0, i.e. s_0 has a self-loop with the label $\{a, \neg b, \neg Tail\}$. To compute another next state of s_0, we add the constraint $\neg p_1$ to the input of the SAT solver. Repeat the above process and we can construct all states in the transition system.

Checking the satisfiability of ϕ is then reduced to finding a *final state* (Definition 6) in the corresponding transition system. Since ϕ is in TNF, a final state s meets the constraint that $Tail \wedge \mathsf{xnf}(\bigwedge s)^p$ (recall s is a set of subformulas of ϕ) is satisfiable. For the above example, the initial state s_0 is actually a final state, as $Tail \wedge \mathsf{xnf}(\phi)^p$ is satisfiable. Because all states computed by the SAT solver in the transition system are reachable from the initial state, we can prove that ϕ is satisfiable iff there is a final state in the system (Theorem 4).

We present a conflict-driven algorithm, i.e. CDLSC, to accelerate the satisfiability checking. CDLSC maintains a *conflict sequence* $\mathcal{C}$, in which each element, denoted as $\mathcal{C}[i]$

$(0 \leq i < |\mathcal{C}|)$, is a set of states in the transition system that cannot reach a final state in i steps. Starting from the initial state, CDLSC iteratively checks whether a final state can be reached, and makes use of the conflict sequence to accelerate the search. Consider the formula $\phi = (\neg Tail)\mathcal{U}a \wedge (\neg Tail)\mathcal{U}(\neg a) \wedge (\neg Tail)\mathcal{U}b \wedge (\neg Tail)\mathcal{U}(\neg b) \wedge (\neg Tail)\mathcal{U}c$. In the first iteration, CDLSC checks whether the initial state $s_0 = \{\phi\}$ is a final state, i.e. whether $Tail \wedge \mathsf{xnf}(\phi)^p$ is satisfiable. The answer is negative, so s_0 cannot reach a final state in 0 steps and can be added into $\mathcal{C}[0]$. However, we can do better by leveraging the Unsatisfiable Core (UC) returned from the SAT solver. Assume that we get the UC $u_1 = \{(\neg Tail)\mathcal{U}a, (\neg Tail)\mathcal{U}(\neg a)\}$. That indicates every state s containing u, i.e. $s \supseteq u$, is not a final state. As a result, we can add u instead of s_0 into $\mathcal{C}[0]$, making the algorithm much more efficient.

Now in the second iteration, CDLSC first tries to compute a one-transition next state of s_0 that is not included in $\mathcal{C}[0]$. (Otherwise the new state cannot reach a final state in 0 step.) This can be encoded as a Boolean formula $\mathsf{xnf}(\phi)^p \wedge \neg(p_1 \wedge p_2)$ where p_1, p_2 represent $\mathcal{X}((\neg Tail)\mathcal{U}a)$ and $\mathcal{X}((\neg Tail)\mathcal{U}(\neg a))$ respectively. Assume the new state $s_1 = \{(\neg Tail)\mathcal{U}a, (\neg Tail)\mathcal{U}b, (\neg Tail)\mathcal{U}(\neg b), (\neg Tail)\mathcal{U}c\}$ is generated from the assignment of the SAT solver. Then CDLSC checks whether s_1 can reach a final state in 0 step, i.e. $\mathsf{xnf}(\bigwedge s_1)^p \wedge Tail$ is satisfiable. The answer is negative and we can add the UC $u_2 = \{(\neg Tail)\mathcal{U}b, (\neg Tail)\mathcal{U}(\neg b)\}$ to $\mathcal{C}[0]$ as well. Now to compute a next state of s_0 that is not included in $\mathcal{C}[0]$, the encoded Boolean formula becomes $\mathsf{xnf}(\phi)^p \wedge \neg(p_1 \wedge p_2) \wedge \neg(p_3 \wedge p_4)$ where p_3, p_4 represent $\mathcal{X}((\neg Tail)\mathcal{U}b)$ and $\mathcal{X}((\neg Tail)\mathcal{U}(\neg b))$ respectively. Assume the new state $s_2 = \{(\neg Tail)\mathcal{U}a, (\neg Tail)\mathcal{U}b, (\neg Tail)\mathcal{U}c\}$ is generated from the assignment of the SAT solver. Since $\mathsf{xnf}(\bigwedge s_2)^p \wedge Tail$ is satisfiable, s_2 is a final state and we conclude that the formula ϕ is satisfiable. In principle, there are a total of $2^5 = 32$ states in the transition system of ϕ, but CDLSC succeeds to find the answer by computing only 3 of them (including the initial state).

CDLSC also leverages the conflict sequence to accelerate checking unsatisfiable formulas. Similar to Bounded Model Checking (BMC) (Biere *et al.* 1999), CDLSC searches the model iteratively. However, BMC invokes only 1 SAT call for each iteration, while CDLSC invokes multiple SAT calls. CDLSC is more like an IC3-style algorithm, but achieves a much simpler implementation by using UC instead of the *Minimal Inductive Core* (MIC) like IC3 (Bradley 2011).

SAT-based Explicit-State Checking

Given an LTL_f formula ϕ, we construct the LTL_f *transition system* (Li *et al.* 2014; Li *et al.* 2015) by SAT solvers and then check the satisfiability of the formula over its corresponding transition system.

LTL_f Transition System

First, we show how one can consider LTL_f formulas as propositional ones. This requires considering temporal subformulas as *propositional atoms*.

Definition 2 (Propositional Atoms)**.** *For an LTL_f formula ϕ, we define the set of propositional atoms of ϕ, i.e. $\mathsf{PA}(\phi)$, as follows: (1) $\mathsf{PA}(\phi) = \{\phi\}$ if ϕ is an atom, Next, Until or Release formula; (2) $\mathsf{PA}(\phi) = \mathsf{PA}(\psi)$ if $\phi = (\neg\psi)$; (3) $\mathsf{PA}(\phi) = \mathsf{PA}(\phi_1) \cup \mathsf{PA}(\phi_2)$ if $\phi = (\phi_1 \wedge \phi_2)$ or $(\phi_1 \vee \phi_2)$.*

Consider $\phi = (a \wedge ((\neg Tail \wedge a)\mathcal{U}b) \wedge \neg(\neg Tail \wedge \mathcal{X}(a \vee b)))$. We have $\mathsf{PA}(\phi) = \{a, Tail, ((\neg Tail \wedge a)\mathcal{U}b), (\mathcal{X}(a \vee b))\}$. Intuitively, the propositional atoms are obtained by treating all temporal subformulas of ϕ as atomic propositions. Thus, an LTL_f formula ϕ can be viewed as a propositional formula over $\mathsf{PA}(\phi)$.

Definition 3. *For an LTL_f formula ϕ, let ϕ^p be ϕ considered as a propositional formula over $\mathsf{PA}(\phi)$.* A propositional assignment *A of ϕ^p, is in $2^{\mathsf{PA}(\phi)}$ and satisfies $A \models \phi^p$.*

Consider the formula $\phi = (a \vee (\neg Tail \wedge a)\mathcal{U}b) \wedge (b \vee (Tail \vee c)\mathcal{R}d)$. From Definition 3, ϕ^p is $(a \vee p_1) \wedge (b \vee p_2)$ where p_1, p_2 are two Boolean variables representing the truth values of $(\neg Tail \wedge a)\mathcal{U}b$ and $(Tail \vee c)\mathcal{R}d$. Moreover, the set $\{\neg a, p_1((\neg Tail \wedge a)\mathcal{U}b), \neg b, p_2((Tail \vee c)\mathcal{R}d)\}$ is a propositional assignment of ϕ^p. In the rest of the paper, we do not introduce the intermediate variables and directly say $\{\neg a, (\neg Tail \wedge a)\mathcal{U}b, \neg b, (Tail \vee c)\mathcal{R}d\}$ is a propositional assignment of ϕ^p. The following theorem shows the relationship between the propositional assignment of ϕ^p and the satisfaction of ϕ.

Theorem 2. *For an LTL_f formula ϕ and a finite trace ξ, $\xi \models \phi$ implies there exists a propositional assignment A of ϕ^p such that $\xi \models \bigwedge A$; On the other hand, $\xi \models \bigwedge A$ where A is a propositional assignment of ϕ^p, also implies $\xi \models \phi$.*

We now introduce the *neXt Normal Form* (XNF) of LTL_f formulas, which is useful for the construction of the transition system.

Definition 4 (neXt Normal Form)**.** *An LTL_f formula ϕ is in* neXt Normal Form *(XNF) if there are no Until or Release subformulas of ϕ in $\mathsf{PA}(\phi)$.*

For example, $\phi = ((\neg Tail \wedge a)\mathcal{U}b)$ is not in XNF, while $(b \vee (\neg Tail \wedge a \wedge (\mathcal{X}((\neg Tail \wedge a)\mathcal{U}b))))$ is. Every LTL_f formula ϕ has a linear-time conversion to an equivalent formula in XNF, which we denoted as $\mathsf{xnf}(\phi)$.

Theorem 3. *For an LTL_f formula ϕ, there is a corresponding LTL_f formula $\mathsf{xnf}(\phi)$ in XNF such that $\phi \equiv \mathsf{xnf}(\phi)$. Furthermore, the cost of the conversion is linear.*

Observe that when ϕ is in XNF, there can be only Next (no Until or Release) temporal formulas in the propositional assignment of ϕ^p. For $\phi = b \vee (a \wedge \neg Tail \wedge \mathcal{X}(a\mathcal{U}b))$, the set $A = \{a, \neg b, \neg Tail, \mathcal{X}(a\mathcal{U}b)\}$ is a propositional assignment of ϕ^p. Based on LTL_f semantics, we can induce from A that if a finite trace ξ satisfying $\xi[0] \supseteq \{a, \neg b, \neg Tail\}$ and $\xi_1 \models a\mathcal{U}b$, $\xi \models \phi$ is true. This motivates us to construct the transition system for ϕ, in which $\{a\mathcal{U}b\}$ is a next state of $\{\phi\}$ and $\{a, \neg b, \neg Tail\}$ is the transition label between these two states.

Let ϕ be an LTL_f formula and A be a propositional assignment of ϕ^p, we denote $L(A) = \{l | l \in A \text{ is a literal}\}$ and $X(A) = \{\theta | \mathcal{X}\theta \in A\}$. Now we define the *transition system* for an LTL_f formula.

Definition 5. *Given an* LTL_f *formula* ϕ *and its literal set* $\mathcal{L}$, *let* $\Sigma = 2^{\mathcal{L}}$. *We define the* transition system $T_\phi = (S, s_0, T)$ *for* ϕ, *where* $S \subseteq 2^{cl(\phi)}$ *is the set of states,* $s_0 = \{\phi\} \in S$ *is the* initial state, *and*

- $T : S \times \Sigma \rightarrow 2^S$ *is the transition relation, where* $s_2 \in T(s_1, \sigma)$ $(\sigma \in \Sigma)$ *holds iff there is a propositional assignment* A *of* $\mathsf{xnf}(\bigwedge s_1)^p$ *such that* $\sigma \supseteq L(A)$ *and* $s_2 = X(A)$.

A run *of* T_ϕ *on a finite trace* $\xi(|\xi| = n > 0)$ *is a finite sequence* $s_0, s_1, \ldots, s_n$ *such that* s_0 *is the initial state and* $s_{i+1} \in T(s_i, \xi[i])$ *holds for all* $0 \leq i < n$.

We define the notation $|r|$ for a run r, to represent the length of r, i.e. number of states in r. We say state s_2 is reachable from state s_1 in $i(i \geq 0)$ steps (resp. in up to i steps), if there is a run r on some finite trace ξ leading from s_1 to s_2 and $|r| = i$ (resp. $|r| \leq i$). In particular, we say s_2 is a *one-transition next state* of s_1 if s_2 is reachable from s_1 in 1 steps. Since a state s is a subset of $cl(\phi)$, which essentially is a formula with the form of $\bigwedge_{\psi \in s} \psi$, we mix the usage of the state and formula in the rest of the paper. That is, a state can be a formula of $\bigwedge_{\psi \in s} \psi$, and a formula ϕ can be a set of states, i.e. $s \in \phi$ iff $s \Rightarrow \phi$.

Lemma 1. *Let* $T_\phi = (S, s_0, T)$ *be the transition system of* ϕ. *Every state* $s \in S$ *is reachable from the initial state* s_0.

Definition 6 (Final State)**.** *Let* s *be a state of a transition system* T_ϕ. *Then* s *is a* final state *of* T_ϕ *iff the Boolean formula* $Tail \wedge (\mathsf{xnf}(s))^p$ *is satisfiable.*

By introducing the concept of *final state*, we are able to check the satisfiability of the LTL_f formula ϕ over T_ϕ.

Theorem 4. *Let* ϕ *be an* LTL_f *formula. Then* ϕ *is satisfiable iff there is a final state in* T_ϕ.

An intuitive solution from Theorem 4 to check the satisfiability of ϕ is to construct states of T_ϕ until (1) either a final state is found by Definition 6, meaning ϕ is satisfiable; or (2) all states in T_ϕ are generated but no final state can be found, meaning ϕ is unsatisfiable. This approach is simple and easy to implement, however, it does not perform well according to our preliminary experiments.

Conflict-Driven LTL_f Satisfiability Checking

In this section, we present a conflict-driven algorithm for LTL_f satisfiability checking. The new algorithm is inspired by (Li *et al.* 2015), where information of both satisfiability and unsatisfiability results of SAT solvers are used. The motivation is as follows: In Definition 6, if the Boolean formula $Tail \wedge \mathsf{xnf}(s)^p$ is unsatisfiable, the SAT solver is able to provide a UC (Unsatisfiable Core) c such that $c \subseteq s$ and $Tail \wedge \mathsf{xnf}(c)^p$ is still unsatisfiable. It means that c represents a set of states that are not final states. By adding a new constraint $\neg(\bigwedge_{\psi \in c} \mathcal{X}\psi)$, the SAT solver can provide a model (if exists) that avoids re-generation of those states in c, which accelerates the search of final states. More generally, we define the *conflict sequence*, which is used to maintain all information of UCs acquired during the checking process.

Definition 7 (Conflict Sequence)**.** *Given an* LTL_f *formula* ϕ, *a conflict sequence* $\mathcal{C}$ *for the transition system* T_ϕ *is a finite sequence of set of states such that:*

1. *The initial state* $s_0 = \{\phi\}$ *is in* $\mathcal{C}[i]$ *for* $0 \leq i < |\mathcal{C}|$;
2. *Every state in* $\mathcal{C}[0]$ *is not a final state;*
3. *For every state* $s \in \mathcal{C}[i+1]$ $(0 \leq i < |\mathcal{C}| - 1)$, *all the one-transition next states of* s *are included in* $\mathcal{C}[i]$.

We call each $\mathcal{C}[i]$ *is a* frame, *and* i *is the* frame level.

In the definition, $|\mathcal{C}|$ represents the length of $\mathcal{C}$ and $\mathcal{C}[i]$ denotes the i-th element of $\mathcal{C}$. Consider the transition system shown in Figure 1, in which s_0 is the initial state and s_4 is the final state. Based on Definition 7, the sequence $\mathcal{C} = \{s_0, s_1, s_2, s_3\}, \{s_0, s_1\}, \{s_0\}$ is a conflict sequence. Notably, the conflict sequence for a transition system may not be unique. For the above example, the sequence $\{s_0, s_1\}, \{s_0\}$ is also a conflict sequence for the system. This suggests that the construction of a conflict sequence is algorithm-specific. Moreover, it is not hard to induce that every non-empty prefix of a conflict sequence is also a conflict sequence. For example, a prefix of $\mathcal{C}$ above, i.e. $\{s_0, s_1, s_2, s_3\}, \{s_0, s_1\}$, is a conflict sequence. As a result, a conflict sequence can be constructed iteratively, i.e. the elements can be generated (and updated) in order. Our new algorithm is motivated by these two observations.

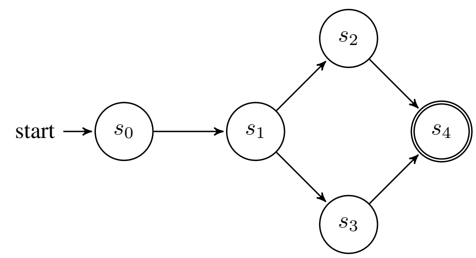

Figure 1: An example transition system for the conflict sequence.

An inherent property of conflict sequences is described in the following lemma.

Lemma 2. *Let* ϕ *be an* LTL_f *formula with a conflict sequence* $\mathcal{C}$ *for the transition system* T_ϕ, *then* $\bigcap_{0 \leq j \leq i} \mathcal{C}[j] (0 \leq i < |\mathcal{C}|)$ *represents a set of states that cannot reach a final state in up to* i *steps.*

Proof. We first prove $\mathcal{C}[i](i \geq 0)$ is a set of states that cannot reach a final state in i step. Basically from Definition 7, $\mathcal{C}[0]$ is a set of states that are not final states. Inductively, assume $\mathcal{C}[i](i \geq 0)$ is a set of states that cannot reach a final state in i steps. From Item 3 of Definition 7, every state $s \in \mathcal{C}[i+1]$ satisfies all its one-transition next states are in $\mathcal{C}[i]$, thus every state $s \in \mathcal{C}[i+1]$ cannot reach a final state in $i+1$ steps. Now since $\mathcal{C}[i](i \geq 0)$ is a set of states that cannot reach a final state in i steps, $\bigcap_{0 \leq j \leq i} \mathcal{C}[j]$ is a set of states that cannot reach a final state in up to i steps. □

We are able to utilize the conflict sequence to accelerate the satisfiability checking of LTL_f formulas, using the theoretical foundations provided by Theorem 5 and 6 below.

Theorem 5. *The* LTL_f *formula* ϕ *is satisfiable iff there is a run* $r = s_0, s_1, \ldots, s_n (n \geq 0)$ *of* T_ϕ *such that (1)* s_n *is a final state; and (2)* s_i $(0 \leq i \leq n)$ *is not in* $\mathcal{C}[n-i]$ *for every conflict sequence* $\mathcal{C}$ *of* T_ϕ *with* $|\mathcal{C}| > n - i$.

Proof. ($\Leftarrow$) Since s_n is a final state, ϕ is satisfiable according to Theorem 4. ($\Rightarrow$) Since ϕ is satisfiable, there is a finite trace ξ such that the corresponding run r of T_ϕ on ξ ends with a final state (according to Theorem 4). Let r be $s_0 \to s_1 \to \ldots s_n$ where s_n is the final state. It holds that s_i ($0 \leq i \leq n$) is a state that can reach a final state in $n-i$ steps. Moreover for every $\mathcal{C}$ of T_ϕ with $|\mathcal{C}| > n-i, \mathcal{C}[n-i]$ ($\mathcal{C}[n-i]$ is meaningless when $|\mathcal{C}| \leq n-i$) represents a set of states that cannot reach a final state in $n-i$ steps (From the proof of Lemma 2). As a result, it is true that s_i is not in $\mathcal{C}[n-i]$ if $|\mathcal{C}| > n-i$. □

Theorem 5 suggests that to check whether a state s can reach a final state in i steps ($i \geq 1$), finding a one-transition next state s' of s that is not in $\mathcal{C}[i-1]$ is necessary; as $s' \in \mathcal{C}[i-1]$ imples s' cannot reach a final state in $i-1$ steps (From the proof of Lemma 2). If all one-transition next states of s are in $\mathcal{C}[i-1]$, s cannot reach a final state in i steps.

Theorem 6. *The* LTL_f *formula* ϕ *is unsatisfiable iff there is a conflict sequence* $\mathcal{C}$ *and* $i \geq 0$ *such that* $\bigcap_{0\leq j\leq i} \mathcal{C}[j] \subseteq \mathcal{C}[i+1]$.

Proof. ($\Leftarrow$) $\bigcap_{0\leq j\leq i} \mathcal{C}[j] \subseteq \mathcal{C}[i+1]$ is true implies that $\bigcap_{0\leq j\leq i} \mathcal{C}[j] = \bigcap_{0\leq j\leq i+1} \mathcal{C}[j]$ is true. Also from Lemma 2 we know $\bigcap_{0\leq j\leq i} \mathcal{C}[j]$ is a set of states that cannot reach a final state in up to i steps. Since $\phi \in \mathcal{C}[i]$ is true for each $i \geq 0$, ϕ is in $\bigcap_{0\leq j\leq i} \mathcal{C}[j]$. Moreover, $\bigcap_{0\leq j\leq i} \mathcal{C}[j] = \bigcap_{0\leq j\leq i+1} \mathcal{C}[j]$ is true implies all reachable states from ϕ are included in $\bigcap_{0\leq j\leq i} \mathcal{C}[j]$. We have known all states in $\bigcap_{0\leq j\leq i} \mathcal{C}[j]$ are not final states, so ϕ is unsatisfiable.

($\Rightarrow$) If ϕ is unsatisfiable, every state in T_ϕ is not a final state. Let S be the set of states of T_ϕ. According to Lemma 2, $\bigcap_{0\leq j\leq i} \mathcal{C}[j] (i \geq 0)$ contains the set of states that are not final in up to i steps. Now we let $\mathcal{C}$ satisfy that $\bigcap_{0\leq j\leq i} \mathcal{C}[j] (i \geq 0)$ contains all states that are not final in up to i steps, so $\bigcap_{0\leq j\leq i} \mathcal{C}[j]$ includes all reachable states from ϕ, as ϕ is unsatisfiable. However, because $\bigcap_{0\leq j\leq i} \mathcal{C}[j] \supseteq \bigcap_{0\leq j\leq i+1} \mathcal{C}[j] \supseteq S (i \geq 0)$, there must be an $i \geq 0$ such that $\bigcap_{0\leq j\leq i} \mathcal{C}[j] = \bigcap_{0\leq j\leq i+1} \mathcal{C}[j]$, which indicates that $\bigcap_{0\leq j\leq i} \mathcal{C}[j] \subseteq \mathcal{C}[i+1]$ is true. □

Algorithm Design. The algorithm, named CDLSC (Conflict-Driven LTL_f Satisfiability Checking), constructs the transition system on-the-fly. The initial state s_0 is fixed to be $\{\phi\}$ where ϕ is the input formula. From Definition 6, whether a state s is final is reducible to the satisfiability checking of the Boolean formula $Tail \wedge \mathsf{xnf}(s)^p$. If s_0 is a final state, there is no need to maintain the conflict sequence in CDLSC, and the algorithm can return SAT immediately; Otherwise, the conflict sequence is maintained as follows.

- In CDLSC, every element of $\mathcal{C}$ is a set of set of subformulas of the input formula ϕ. Formally, each $\mathcal{C}[i]$ ($i \geq 0$) can be represented by the LTL_f formula $\bigvee_{c\in\mathcal{C}[i]} \bigwedge_{\psi\in c} \psi$ where c is a set of subformulas of ϕ. We mix-use the notation $\mathcal{C}[i]$ for the corresponding LTL_f formula as well. Every state s satisfying $s \Rightarrow \mathcal{C}[i]$ is included in $\mathcal{C}[i]$.
- $\mathcal{C}$ is created iteratively. In each iteration $i \geq 0$, $\mathcal{C}[i]$ is initialized as the empty set.
- To compute elements in $\mathcal{C}[0]$, we consider an existing state s (e.g. s_0). If the Boolean formula $Tail \wedge \mathsf{xnf}(s)^p$ is unsatisfiable, s is not a final state and can be added into $\mathcal{C}[0]$ from Item 2 of Definition 7. Moreover, CDLSC leverages the Unsatisfiable Core (UC) technique from the SAT community to add a set of states, all of which are not final and include s, to $\mathcal{C}[0]$. This set of states, denoted as c, is also represented by a set of LTL_f formulas and satisfies $c \subseteq s$. The detail to obtain c is discussed below.
- To compute elements in $\mathcal{C}[i+1]$ ($i \geq 0$), we consider the Boolean formula $(\mathsf{xnf}(s) \wedge \neg\mathcal{X}(\mathcal{C}[i]))^p$, where $\mathcal{X}(\mathcal{C}[i])$ represents the LTL_f formula $\bigvee_{c\in\mathcal{C}[i]} \bigwedge_{\psi\in c} \mathcal{X}(\psi)$. The above Boolean formula is used to check whether there is a one-transition next state of s that is not in $\mathcal{C}[i]$. If the formula is unsatisfiable, all the one-transition next states of s are in $\mathcal{C}[i]$, thus s can be added into $\mathcal{C}[i+1]$ according to Item 3 of Definition 7. Similarly, we also utilize the UC technique to obtain a subset c of s, such that c represents a set of states that can be added into $\mathcal{C}[i+1]$.

As shown above, every Boolean formula sent to a SAT solver has the form of $(\mathsf{xnf}(s) \wedge \theta)^p$ where s is a state and θ is either $Tail$ or $\neg\mathcal{X}(\mathcal{C}[i])$. Since every state s consists of a set of LTL_f formulas, the Boolean formula can be rewritten as $\alpha_1 = (\bigwedge_{\psi\in s} \mathsf{xnf}(\psi) \wedge \theta)^p$. Moreover, we introduce a new Boolean variable p_ψ for each $\psi \in s$, and re-encode the formula to be $\alpha_2 = \bigwedge_{\psi\in s} p_\psi \wedge (\bigwedge_{\psi\in s}(\mathsf{xnf}(\psi) \vee \neg p_\psi) \wedge \theta)^p$. α_2 is satisfiable iff α_1 is satisfiable, and A is an assignment of α_2 iff $A \backslash \{p_\psi | \psi \in s\}$ is an assignment of α_1. Sending α_2 instead of α_1 to the SAT solver that supports assumptions (e.g. Minisat (Eén and Sörensson 2003)) enables the SAT solver to return the UC, which is a set of s, when α_2 is unsatisfiable. For example, assume $s = \{\psi_1, \psi_2, \psi_3\}$ and α_2 is sent to the SAT solver with $\{p_{\psi_i} | i \in \{1, 2, 3\}\}$ being the assumptions. If the SAT solver returns unsatisfiable and the UC $\{p_{\psi_1}\}$, the set $c = \{\psi_1\}$, which represents every state including ψ_1, is the one to be added into the corresponding $\mathcal{C}[i]$. We use the notation $get\_uc()$ for the above procedure.

The pseudo-code of CDLSC is shown in Algorithm 1. Line 1-2 considers the situation when the input formula ϕ is a final state itself. Otherwise, the first frame $\mathcal{C}[0]$ is initialized to $\{\phi\}$ (Line 3), and the current frame level is set to 0 (Line 4). After that, the loop body (Line 5-11) keeps updating the elements of $\mathcal{C}$ iteratively, until either the procedure $try\_satsify$ returns true, which means to find a model of ϕ, or the procedure inv_found returns true, which is the implementation of Theorem 6. The loop continues to create a new frame in $\mathcal{C}$ if neither of the procedures succeeds to return true. To describe conveniently, we say every run of the while loop body in Algorithm 1 is an *iteration*.

The procedure $try\_satisfy$ is responsible for updating $\mathcal{C}$. Taking an formula ϕ and the frame level $frame\_level$ currently working on, $try\_satisfy$ returns true iff a model of ϕ can be found, with the length of $frame\_level + 1$. As shown in Algorithm 2, $try\_satisfy$ is implemented in a recursive way. Each time it checks whether a next state of the input ϕ, which belongs to a lower level (than the input $frame\_level$) frame can be found (Line 2). If the result is positive and such a new state ϕ' is constructed, $try\_satisfy$

Algorithm 1 Implementation of CDLSC

Input: An LTL$_f$ formula ϕ.
Output: SAT or UNSAT.

1: **if** $Tail \wedge \mathsf{xnf}(\phi)^p$ is satisfiable **then**
2:     **return** SAT;
3: Set $\mathcal{C}[0] := \{\phi\}$;
4: Set $frame\_level := 0$;
5: **while** true **do**
6:     **if** $try\_satisfy(\phi, frame\_level)$ returns true **then**
7:         **return** SAT;
8:     **if** $inv\_found(frame\_level)$ returns true **then**
9:         **return** UNSAT;
10:     $frame\_level := frame\_level + 1$;
11:     Set $\mathcal{C}[frame\_level] = \emptyset$;

first checks whether ϕ' is a final state when $frame\_level$ is 0 (in which case returns true). If ϕ' is not a final state, a UC is extracted from the SAT solver and added to $\mathcal{C}[0]$ (Line 5-11). If $frame\_level$ is not 0, $try\_satisfy$ recursively checks whether a model of ϕ' can be found with the length of $frame\_level$ (Line 12-13). If the result is negative and such a state cannot be constructed, a UC is extracted from the SAT solver and added into $\mathcal{C}[frame\_level+1]$ (Line 14-15).

Algorithm 2 Implementation of $try\_satisfy$

Input: ϕ: The formula is working on;
$frame\_level$: The frame level is working on.
Output: true or false.

1: Let $\psi = \neg\mathcal{X}(\mathcal{C}[frame\_level])$;
2: **while** $(\psi \wedge \mathsf{xnf}(\phi))^p$ is satisfiable **do**
3:     Let A be the model of $(\psi \wedge \mathsf{xnf}(\phi))^p$;
4:     Let $\phi' = X(A)$, i.e. be the next state of ϕ extracted from A;
5:     **if** $frame\_level == 0$ **then**
6:         **if** $Tail \wedge \mathsf{xnf}(\phi')^p$ is satisfiable **then**
7:             **return** true;
8:         **else**
9:             Let $c = get\_uc()$;
10:             Add c into $\mathcal{C}[frame\_level]$;
11:             Continue;
12:     **if** $try\_satisfy(\phi', frame\_level - 1)$ is true **then**
13:         **return** true;
14: Let $c = get\_uc()$;
15: Add c into $\mathcal{C}[frame\_level+1]$;
16: **return** false;

Notably, Item 1 of Definition 7, i.e. $\{\phi\} \in \mathcal{C}[i]$, is guaranteed for each $i \geq 0$, as the original input formula of $try\_satisfy$ is always ϕ (Line 6 in Algorithm 1) and there is some c (Line 15 in Algorithm 2) including $\{\phi\}$ that is added into $\mathcal{C}[i]$, if no model can be found in the current iteration.

The procedure inv_found in Algorithm 1 implements Theorem 6 in a straightforward way: It reduces the checking of whether $\bigcap_{0\leq j\leq i}\mathcal{C}[j] \subseteq \mathcal{C}[i+1]$ being true on some frame level i, to the satisfiability checking of the Boolean formula $\bigwedge_{1\leq j\leq i}\mathcal{C}[j] \Rightarrow \mathcal{C}[i+1]$. Finally, we state Theorem 7 below to provide the theoretical guarantee that CDLSC always terminates correctly.

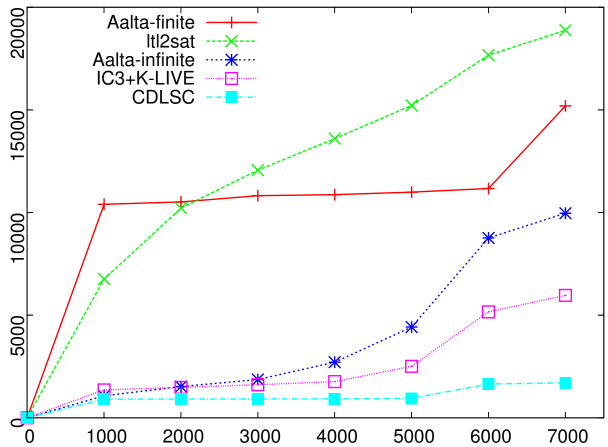

Figure 2: Result for LTL$_f$ Satisfiability Checking on LTL-as-LTL_f Benchmarks. The X axis represents the number of benchmarks, and the Y axis is the accumulated checking time (s).

Lemma 3. *After each iteration of CDLSC with no model found, the sequence $\mathcal{C}$ is a conflict sequence of T_ϕ for the transition system T_ϕ.*

Theorem 7. *The CDLSC algorithm terminates with a correct result.*

CDLSC is shown how to accelerate the checking of satisfiable formulas in the previous section. For unsatisfiable instances, consider $\phi = (\neg Tail)\mathcal{U}a \wedge (Tail)\mathcal{R}\neg a \wedge (\neg Tail)\mathcal{U}b$. CDLSC first checks that $Tail \wedge \mathsf{xnf}(\phi)^p$ is unsatisfiable, where the SAT solver returns $c = \{(\neg Tail)\mathcal{U}a, Tail\mathcal{R}\neg a\}$ as the UC. So c is added into $\mathcal{C}[0]$. Then CDLSC checks that $(\mathsf{xnf}(\phi) \wedge \neg\mathcal{X}(\mathcal{C}[0]))^p$ is still unsatisfiable, in which $c = \{(\neg Tail)\mathcal{U}a, Tail\mathcal{R}\neg a\}$ is still the UC. So c is added into $\mathcal{C}[1]$ as well. Since $\mathcal{C}[0] \subseteq \mathcal{C}[1]$ and according to Theorem 6, CDLSC terminates with the unsatisfiable result. In this case, CDLSC only visits one state for the whole checking process. For a more general instance like $\phi \wedge \psi$, where ψ is a large LTL$_f$ formula, checking by CDLSC enables to achieve a significantly improvement compared to the checking by traditional tableau approach.

Summarily, CDLSC is a conflict-driven on-the-fly satisfiability checking algorithm, which successfully leads to either an earlier finding of a satisfying model, or the faster termination with the unsatisfiable result.

Experimental Evaluation

Benchmarks We first consider the *LTL-as-*LTL$_f$ benchmark, which is evaluated by previous works on LTL$_f$ satisfiability checking (Li *et al.* 2014; Fionda and Greco 2016). This benchmark consists of 7442 instances that are originally LTL formulas but are treated as LTL$_f$ formulas, as both logics share the same syntax. Previous works (Li *et al.* 2014; Fionda and Greco 2016) have shown that the benchmark is useful to test the scalability of LTL$_f$ solvers.

Secondly, we consider the 7 LTL$_f$*-specific* patterns that are introduced in recent researches on LTL$_f$, e.g. (De Giacomo *et al.* 2014; Di Ciccio *et al.* 2016), and we create 100 instances for each pattern. As shown in Table 1, it is trivial to check the satisfiability of these LTL$_f$ patterns by most

Table 1 Results for LTL_f Satisfiability Checking on LTL_f-specific Benchmarks.

Type	Number	Result	IC3+K-LIVE	Aalta-finite	Aalta-infinite	ltl2sat	CDLSC
Alternate Response	100	sat	134	1	48	123	3
Alternate Precedence	100	sat	154	3	70	380	4
Chain Precedence	100	sat	127	2	45	83	2
Chain Response	100	sat	79	1	41	49	2
Precedence	100	sat	132	2	14	124	1
Responded Existence	100	sat	130	1	14	327	1
Response	100	sat	155	1	41	53	2
Practical Conjunction	1000	varies	1669	19564	4443	20477	115

tested solvers, as either they have small sizes or dedicated heuristics for LTL_f, which are encoded in both Aalta-finite and CDLSC, enable to solve them quickly. Inspired from the observation in (Li *et al.* 2013) that an LTL specification in practice is often the conjunction of a set of small and frequently-used patterns, we randomly choose a subset of the instances of the 7 patterns to imitate a real LTL_f specification in practice. We generate 1000 such instances as the *practical conjunction* pattern shown in the last row of Table 1. Unlike the random benchmarks in SAT community, which are often considered not interesting, we argue that the new practical conjunction pattern is a representative for real LTL_f specifications in industry.

Experimental Setup We implement CDLSC in C++, and use Minisat 2.2.0 (Eén and Sörensson 2003) as the SAT engine[1]. We compare it with two extant LTL_f satisfiability solvers: Aalta-finite (Li *et al.* 2014) and ltl2sat (Fionda and Greco 2016). We also compared with the state-of-art LTL solver Aalta-infinite (Li *et al.* 2015), using the LTL_f-to-LTL satisfiability-preserving reduction described in (De Giacomo and Vardi 2013). As LTL satisfiability checking is reducible to model checking, as described in (Rozier and Vardi 2007), we also compared with this reduction, using nuXmv with the IC3+K-LIVE back-end (Cavada *et al.* 2014), as an LTL_f satisfiability checker.

We ran the experiments on a RedHat 6.0 cluster with 2304 processor cores in 192 nodes (12 processor cores per node), running at 2.83 GHz with 48GB of RAM per node. Each tool executed on a dedicated node with a timeout of 60 seconds, measuring execution time with Unix `time`. Excluding timeouts, all solvers found correct verdicts for all formulas. All artifacts are available in the supplemental material.

Results Figure 2 shows the results for LTL_f satisfiability checking on LTL-as-LTL_f benchmarks. CDLSC outperforms all other approaches. On average, CDLSC performs about 4 times faster than the second-best approach IC3+K-LIVE (1705 seconds vs. 6075 seconds). CDLSC checks the LTL_f formula directly, while IC3+K-LIVE must take the input of the LTL formula translated from the LTL_f formula. As a result, IC3-KLIVE may take extra cost, e.g. finding a satisfying lasso for the model, to the satisfiability checking. Meanwhile, CDLSC can benefit from the heuristics dedicated for LTL_f that are proposed in (Li *et al.* 2014). Finally, the performance of ltl2sat is highly tied to its performance of unsatisfiability checking as most of the timeout cases for ltl2sat are unsatifiable. For Aalta-finite, its performance is restricted by the heavy cost of Tableau Construction.

Table 1 shows the results for LTL_f-specific experiments. Columns 1-3 show the types of LTL_f formulas under test, the number of test instances for each formula type, and the results by formula type. Columns 4-8 show the checking times by formula types in seconds. The dedicated LTL_f solvers perform extremely fast on the seven scalable pattern formulas (Column 5 and 8), because their heuristics work well on these patterns. For the difficult conjunctive benchmarks, CDLSC still outperforms all other solvers.

[1]https://github.com/lijwen2748/aaltaf

Discussion and Concluding Remarks

Bounded Model Checking (BMC) (Biere *et al.* 1999) is also a popular SAT-based technique, which is however, not necessary to compare. There are two ways to apply BMC to LTL_f satisfiability checking. The first one is to check the satisfiability of the LTL formula from the input LTL_f formula. (Li *et al.* 2015) has shown that this approach cannot perform better than IC3+K-LIVE, and the fact of CDLSC outperforming IC3+K-LIVE induces CDLSC also outperforms BMC. The second approach is to check the satisfiability of the LTL_f formula ϕ directly, by unrolling ϕ iteratively. In the worst case, BMC can terminate (with UNSAT) once the iteration reaches the upper bound. This is exactly what is implemented in ltl2sat (Fionda and Greco 2016).

In this paper, we introduce a new SAT-based framework, based on which we present a conflict-driven algorithm CDLSC, for LTL_f satisfiability checking. Our experiments demonstrate that CDLSC outperforms Aalta-infinite and IC3+K-LIVE, which are designed for LTL satisfiability checking, showing the advantage of a dedicated algorithm for LTL_f. Notably, CDLSC maintains a conflict sequence, which is similar to the state-of-art model checking technique IC3 (Bradley 2011). CDLSC does not require the conflict sequence to be monotone, and simply use the UC from SAT solvers to update the sequence. Meanwhile, IC3 requires the sequence to be strictly monotone, and has to compute its dedicated MIC (Minimal Inductive Core) to update the sequence. We conclude that CDLSC outperforms other existing approaches for LTL_f satisfiability checking.

Acknowledgement. We thank anonymous reviewers for the helpful comments. This work is supported in part by NASA ECF NNX16AR57G, NSF CAREER Award CNS-1552934, NSF grants CCF-1319459 and IIS-1527668, NSF Expeditions in Computing project "ExCAPE: Expeditions in Computer Augmented Program Engineering", NSFC projects No. 6157297, No. 61632005 and No. 61532019, and China HGJ project No. 2017ZX01038102-002.

References

[Bacchus and Kabanza 1998] F. Bacchus and F. Kabanza. Planning for temporally extended goals. *Ann. of Mathematics and Artificial Intelligence*, 22:5–27, 1998.

[Bacchus and Kabanza 2000] F. Bacchus and F. Kabanza. Using temporal logic to express search control knowledge for planning. *Artificial Intelligence*, 116(1–2):123–191, 2000.

[Bienvenu *et al.*] M. Bienvenu, C. Fritz, and S. McIlraith. In *KR*, Lake District, UK.

[Bienvenu *et al.* 2011] M. Bienvenu, C. Fritz, and S. A. McIlraith. Specifying and computing preferred plans. *Artificial Intelligence*, 175(7Â¨C8):1308 – 1345, 2011.

[Biere *et al.* 1999] A. Biere, A. Cimatti, E.M. Clarke, and Y. Zhu. Symbolic model checking without BDDs. In *Proc. 5th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems*, volume 1579 of *Lecture Notes in Computer Science*. Springer, 1999.

[Bradley 2011] A. Bradley. SAT-based model checking without unrolling. In Ranjit Jhala and David Schmidt, editors, *Verification, Model Checking, and Abstract Interpretation*, volume 6538 of *LNCS*, pages 70–87. Springer, 2011.

[Calvanese *et al.* 2002] D. Calvanese, G. De Giacomo, and M.Y. Vardi. Reasoning about actions and planning in LTL action theories. In *Principles of Knowledge Representation and Reasoning*, pages 593–602. Morgan Kaufmann, 2002.

[Camacho *et al.* 2017] A. Camacho, J.A. Baier, C. Muise, and A.S. McIlraith. Bridging the gap between LTL synthesis and automated planning. Technical report, U. Toronto, 2017.

[Cavada *et al.* 2014] R. Cavada, A. Cimatti, M. Dorigatti, A. Griggio, A. Mariotti, A. Micheli, S. Mover, M. Roveri, and S. Tonetta. The NuXMV symbolic model checker. In *CAV*, pages 334–342, 2014.

[Claessen and Sörensson 2012] K. Claessen and N. Sörensson. A liveness checking algorithm that counts. In *FMCAD*, pages 52–59. IEEE, 2012.

[De Giacomo and Vardi 1999] G. De Giacomo and M.Y. Vardi. Automata-theoretic approach to planning for temporally extended goals. In *Proc. European Conf. on Planning*, Lecture Notes in AI 1809, pages 226–238. Springer, 1999.

[De Giacomo and Vardi 2013] G. De Giacomo and M. Vardi. Linear temporal logic and linear dynamic logic on finite traces. In *IJCAI*, pages 2000–2007. AAAI Press, 2013.

[De Giacomo *et al.* 2014] G. De Giacomo, R. De Masellis, and M. Montali. Reasoning on LTL on finite traces: Insensitivity to infiniteness. In *AAAI*, pages 1027–1033, 2014.

[Di Ciccio *et al.* 2016] C. Di Ciccio, F. Maggi, and J. Mendling. Efficient discovery of target-branched declare constraints. *Inf. Syst.*, 56(C):258–283, March 2016.

[Eén and Sörensson 2003] N. Eén and N. Sörensson. An extensible SAT-solver. In *SAT*, pages 502–518, 2003.

[Een *et al.* 2011] N. Een, A. Mishchenko, and R. Brayton. Efficient implementation of property directed reachability. In *FMCAD*, pages 125–134, 2011.

[Fionda and Greco 2016] V. Fionda and G. Greco. The complexity of LTL on finite traces: Hard and easy fragments. In *AAAI*, pages 971–977. AAAI Press, 2016.

[Gabaldon 2004] A. Gabaldon. Precondition control and the progression algorithm. In *KR*, pages 634–643. AAAI Press, 2004.

[Gerth *et al.* 1995] R. Gerth, D. Peled, M.Y. Vardi, and P. Wolper. Simple on-the-fly automatic verification of linear temporal logic. In P. Dembiski and M. Sredniawa, editors, *Protocol Specification, Testing, and Verification*, pages 3–18. Chapman & Hall, 1995.

[Li *et al.* 2013] J. Li, L. Zhang, G. Pu, M. Vardi, and J. He. LTL satisfibility checking revisited. In *TIME*, pages 91–98, 2013.

[Li *et al.* 2014] J. Li, L. Zhang, G. Pu, M. Y. Vardi, and J. He. LTL_f satisfibility checking. In *ECAI*, pages 91–98, 2014.

[Li *et al.* 2015] J. Li, S. Zhu, G. Pu, and M.Y. Vardi. SAT-based explicit LTL reasoning. In *HVC*, pages 209–224. Springer, 2015.

[Malik and Zhang 2009] S. Malik and L. Zhang. Boolean satisfiability from theoretical hardness to practical success. *Commun. ACM*, 52(8):76–82, 2009.

[Patrizi *et al.* 2011] F. Patrizi, N. Lipoveztky, G. De Giacomo, and H. Geffner. Computing infinite plans for LTL goals using a classical planner. In *IJCAI*, pages 2003–2008. AAAI Press, 2011.

[Pnueli 1977] A. Pnueli. The temporal logic of programs. In *IEEE FOCS*, pages 46–57, 1977.

[Rozier and Vardi 2007] K.Y. Rozier and M.Y. Vardi. LTL satisfiability checking. In *SPIN*, volume 4595 of *LNCS*, pages 149–167. Springer, 2007.

[Rozier and Vardi 2010] K.Y. Rozier and M.Y. Vardi. LTL satisfiability checking. *STTT*, 12(2):123–137, 2010.

[Sohrabi *et al.* 2011] S. Sohrabi, J. A. Baier, and S. A. McIlraith. Preferred explanations: Theory and generation via planning. In *AAAI*, pages 261–267, August 2011.

[Vardi 2007] M.Y. Vardi. Automata-theoretic model checking revisited. In *VMCAI*, LNCS 4349, pages 137–150. Springer, 2007.

[Vizel *et al.* 2015] Y. Vizel, G. Weissenbacher, and S. Malik. Boolean satisfiability solvers and their applications in model checking. *Proceedings of the IEEE*, 103(11):2021–2035, 2015.

Missing Proofs

Proof of Theorem 1

We first introduce the following lemmas that are useful for the proof.

Lemma 4. *If* $\mathsf{tnf}(\phi)$ *is satisfiable, there is a non-empty finite trace* ξ *such that* $\neg Tail \in \xi[i]$ *for* $0 \le i < |\xi| - 1$, $Tail \in \xi[|\xi| - 1]$ *and* $\xi \models \mathsf{tnf}(\phi)$.

Proof. Since $\mathsf{tnf}(\phi)$ is satisfiable, there is a non-empty finite trace ξ' such that $\xi' \models \mathsf{tnf}(\phi)$. Recall that $\mathsf{tnf}(\phi)$ has the form of $t(\phi) \wedge FTail$, so $\xi' \models \mathsf{tnf}(\phi)$ implies $\xi' \models t(\phi)$ and there is $\le k < |\xi'|$ such that $Tail \in \xi'[k]$ and $Tail \notin \xi'[j]$ for every $j < k$. We define $tp(\xi') = \xi'[0]\xi'[1]\ldots\xi'[k]$, and first prove that $\xi' \models t(\phi)$ implies $tp(\xi') \models t(\phi)$. Let $\xi = tp(\xi')$, and we prove by induction over the type of ϕ that $\xi \models t(\phi)$.

1. If $\phi = \mathsf{tt}$, then $t(\phi) = \mathsf{tt}$ and of course $\xi \models t(\phi)$;
2. If $\phi = l$ is a literal, then $t(\phi) = l$ and $\xi' \models t(\phi)$ implies $l \in \xi'[0] = \xi[0]$. Therefore, $\xi \models t(\phi)$;
3. If $\phi = \phi_1 \wedge \phi_2$, then $t(\phi) = t(\phi_1) \wedge t(\phi_2)$, and $\xi' \models t(\phi)$ implies $\xi' \models t(\phi_1)$ and $\xi' \models t(\phi_2)$. By hypothesis assumption, $\xi' \models t(\phi_1)$ implies $\xi \models t(\phi_1)$ and $\xi' \models t(\phi_2)$ implies $\xi \models t(\phi_2)$. So $\xi \models t(\phi)$ is true. If $\phi = \phi_1 \vee \phi_2$, the proof is similar;
4. If $\phi = X\psi$, then $t(\phi) = \neg Tail \wedge X(t(\psi))$, and $\xi' \models t(\phi)$ implies that $Tail \notin \xi'[0]$ and $\xi'_1 \models t(\psi)$. Let $\xi_1 = tp(\xi'_1)$. By hypothesis assumption, $\xi'_1 \models t(\psi)$ implies $\xi_1 \models t(\psi)$ is true. Moreover, because $Tail \notin \xi'[0]$, $\xi = tp(\xi') = \xi'[0] \cdot tp(\xi'_1) = \xi'[0] \cdot \xi_1$ from its definition. As a result, $\xi \models t(\phi)$ is true;
5. If $\phi = N\psi$, then $t(\phi) = Tail \vee X(t(\psi)) = Tail \vee (\neg Tail \wedge X(t(\psi)))$, and $\xi' \models t(\phi)$ implies that $Tail \in \xi'[0]$ or $\xi' \models \neg Tail \wedge X(t(\psi))$. In the first case, $\xi = \xi'[0]$ and obviously $\xi \models t(\phi)$. For the second case, the proof is the same as that if $\phi = X\psi$;
6. If $\phi = \phi_1 U \phi_2$, then $t(\phi) = (\neg Tail \wedge t(\phi_1)) U t(\phi_2)$, and $\xi' \models t(\phi)$ implies that there is $0 \le i < |\xi'|$ such that $\xi'_i \models t(\phi_2)$ and for every $0 \le j < i$ it holds $\xi'_j \models \neg Tail \wedge t(\phi_1)$. As a result, we have that $\xi = tp(\xi') = \xi'[0] \ldots \xi'[i-1] \cdot tp(\xi'_i)$, and thus $\xi_i = tp(\xi'_i)$ and $\xi_j = \xi'[j] \ldots \xi'[i-1] \cdot tp(\xi'_i) = tp(\xi'_j)$. By hypothesis assumption, $\xi'_i \models t(\phi_2)$ implies $\xi_i \models t(\phi_2)$ and $\xi'_j \models \neg Tail \wedge t(\phi_1)$ implies $\xi_j \models \neg Tail \wedge t(\phi_1)$. As a result, $\xi \models t(\phi)$ is true;
7. If $\phi = \phi_1 R \phi_2$, then $t(\phi) = (Tail \vee t(\phi_1)) R t(\phi_2)$, and $\xi' \models t(\phi)$ implies that for all $0 \le i < |\xi'|$ it holds that, $\xi'_i \models t(\phi_2)$ or there is $0 \le j < i$ such that $\xi'_j \models Tail \vee t(\phi_1)$. Since $\xi = tp(\xi')$, so $\xi_i = tp(\xi'_i)$ for $0 \le i < |\xi|$. By hypothesis assumption, $\xi'_i \models t(\phi_2)$ implies $\xi_i \models t(\phi_2)$ for every $0 \le i < |\xi| - 1$. Moreover, it is true that $Tail \in \xi[|\xi| - 1]$, which implies $\xi[|\xi| - 1] \models Tail \vee t(\phi_1)$. Therefore, we have that $\xi \models t(\phi)$.

Because $tp(\xi') \models t(\phi)$ is true, and $tp(\xi') \models FTail$ is obviously true, we prove finally that $\xi = tp(\xi') \models \mathsf{tnf}(\phi)$. □

Lemma 5. *Let* ξ, ξ' *are two non-empty finite traces satisfying* $|\xi| = |\xi'|$ *and* $\xi'[i] = \xi[i]$ *for* $0 \le i < |\xi| - 1$ *as well as* $\xi'[|\xi| - 1] = \xi[|\xi| - 1] \cup \{Tail\}$. *Then* $\xi \models \phi$ *iff* $\xi' \models \mathsf{tnf}(\phi)$.

Proof. We prove by induction over the type of ϕ.

1. If ϕ is tt, ff or a literal l, obviously $\xi \models \phi$ holds iff $\xi' \models \mathsf{tnf}(\phi)$ holds;
2. If $\phi = \neg\psi$, then $\xi \models \phi$ holds iff $\xi \not\models \psi$ holds. By hypothesis assumption, $\xi \not\models \psi$ holds iff $\xi' \not\models \mathsf{tnf}(\psi)$ holds, which means $\xi \models \phi$ holds iff $\xi' \models \mathsf{tnf}(\phi)$ holds;
3. If $\phi = X\psi$, then $\xi \models \phi$ holds iff $|\xi| > 1$ and $\xi_1 \models \psi$ holds. By hypothesis assumption, $\xi_1 \models \psi$ holds iff $\xi'_1 \models \mathsf{tnf}(\psi)$ holds, and $\xi'_1 \models \mathsf{tnf}(\psi)$ holds iff $\xi' \models \neg Tail \wedge X(\mathsf{tnf}(\psi))$ holds (because $\neg Tail \in \xi'[0]$). As a result, we have the following equations:

$$\begin{aligned} &\xi \models X\psi \\ \Leftrightarrow\ &\xi' \models \neg Tail \wedge X(\mathsf{tnf}(\psi)) \\ \Leftrightarrow\ &\xi' \models \neg Tail \wedge X(t(\psi) \wedge FTail) \\ \Leftrightarrow\ &\xi' \models \neg Tail \wedge X(t(\psi)) \wedge FTail \end{aligned}$$

   Since $\mathsf{tnf}(\phi) = \neg Tail \wedge X(t(\psi)) \wedge FTail$, so $\xi \models \phi$ iff $\xi' \models \mathsf{tnf}(\phi)$ is true;
4. If $\phi = \phi_1 \wedge \phi_2$, then $\xi \models \phi$ holds iff both $\xi \models \phi_1$ and $\xi \models \phi_2$ hold. By hypothesis assumption, we have $\xi \models \phi_1$ holds iff $\xi' \models \mathsf{tnf}(\phi_1)$ holds, and $\xi \models \phi_2$ holds iff $\xi' \models \mathsf{tnf}(\phi_2)$ holds. As a result, $\xi \models \phi$ holds iff $\xi' \models \mathsf{tnf}(\phi_1) \wedge \mathsf{tnf}(\phi_2) = t(\phi_1) \wedge t(\phi_2) \wedge FTail = t(\phi_1 \wedge \phi_2) \wedge FTail = \mathsf{tnf}(\phi_1 \wedge \phi_2)$ holds;
5. If $\phi = \phi_1 U \phi_2$, then $\xi \models \phi$ holds iff there exists $0 \le i < |\xi|$ such that $\xi_i \models \phi_2$, and for every $0 \le j < i$ it holds that $\xi_j \models \phi_1$. By hypothesis assumption, $\xi_i \models \phi_2$ holds iff $\xi'_i \models \mathsf{tnf}(\phi_2)$ holds, and moreover, $\xi_j \models \phi_1$ holds iff $\xi'_j \models \mathsf{tnf}(\phi_1)$ holds. Because of $0 \le j < i$ and $0 \le i < |\xi|$, j does not equal to $|\xi| - 1$, which means $\neg Tail \in \xi'[j]$. As a result, $\xi'[j] \models \neg Tail \wedge \mathsf{tnf}(\phi_1)$. Therefore, $\xi'_i \models \phi_2$ holds and for every $0 \le j < i$, $\xi'_j \models \neg Tail \wedge \mathsf{tnf}(\phi_1)$ is true, which means $\xi' \models (\neg Tail \wedge \mathsf{tnf}(\phi_1)) U \mathsf{tnf}(\phi_2)$ is true. Finally, we have

$$\begin{aligned} &\xi \models \phi_1 U \phi_2 \\ \Leftrightarrow\ &\xi' \models (\neg Tail \wedge \mathsf{tnf}(\phi_1)) U \mathsf{tnf}(\phi_2) \\ \Leftrightarrow\ &\xi' \models (\neg Tail \wedge t(\phi_1) \wedge FTail) U (t(\phi_2) \wedge FTail) \\ \Leftrightarrow\ &\xi' \models (\neg Tail \wedge t(\phi_1)) U t(\phi_2) \wedge FTail \\ \Leftrightarrow\ &\xi' \models \mathsf{tnf}(\phi) \end{aligned}$$

   The proof is done.

□

We are ready now to prove Theorem 1.

Proof. ($\Rightarrow$) If ϕ is satisfiable, there is a non-empty finite trace ξ such that $\xi \models \phi$. From Lemma 5, we know that there is a corresponding finite trace ξ' satisfying $|\xi| = |\xi'|$ and $\xi'[i] = \xi[i]$ for $0 \le i < |\xi| - 1$ as well as $\xi'[|\xi| - 1] = \xi[|\xi| - 1] \cup \{Tail\}$ such that $\xi' \models \mathsf{tnf}(\phi)$. So $\mathsf{tnf}(\phi)$ is satisfiable.

($\Leftarrow$) If $\mathsf{tnf}(\phi)$ is satisfiable, there is a finite trace ξ' satisfying $Tail \notin \xi'[i]$ for $0 \le i < |\xi| - 1$ and $Tail \in \xi'[|\xi| - 1]$ such that $\xi' \models \mathsf{tnf}(\phi)$, from Lemma 4. Moreover, according to Lemma 5, there is a corresponding finite trace satisfying $|\xi| = |\xi'|$ and $\xi[i] = \xi'[i]$ for $0 \le i < |\xi| - 1$ as well as $Tail \notin \xi[|\xi| - 1]$ such that $\xi \models \phi$. So ϕ is satisfiable. □

Proof of Theorem 2

Proof. ($\Rightarrow$) Base case: when ϕ is a literal, Next, Unitl or Release formula, it is true since there is only one propositional assignment of ϕ^p, i.e. $A = \{\phi\}$. Inductive step: if $\phi = \phi_1 \wedge \phi_2$, $\xi \models \phi$ implies $\xi \models \phi_1$ and $\xi \models \phi_2$. By assumption hypothesis, there is A_i of ϕ_i^p ($i = 1, 2$) such that $\xi \models \bigwedge A_i$. Let $A = A_1 \cup A_2$, and a consistent A, in which either ψ or $\neg\psi$ cannot be, must exists (A may not be unique because A_1 and A_2 may not be unique). Otherwise, there is $\psi \in A_1$ and $\neg\psi \in A_2$ such that ξ cannot model $\bigwedge A_1$ and $\bigwedge A_2$ at the same time, which is a contradiction. So A is a propositional assignment of ϕ^p and $\xi \models \bigwedge A$. The proof for $\phi = \phi_1 \vee \phi_2$ is similar.

($\Leftarrow$) A is a propositional assignment of ϕ^p, so $A \models \phi^p$ implies $(\bigwedge A) \Rightarrow \phi$. Therefore, $\xi \models \bigwedge A$ implies that $\xi \models \phi$. □

Proof of Theorem 3

Proof. First, $\mathsf{xnf}(\phi)$ can be constructed recursively as follows: (1) $\mathsf{xnf}(\phi) = \phi$, when ϕ is tt, ff, a literal or $\mathcal{X}\psi$ (Note ϕ is $\mathcal{N}$-free); (2) $\mathsf{xnf}(\phi_1 \ o \ \phi_2) = \mathsf{xnf}(\phi_1) \ o \ \mathsf{xnf}(\phi_2)$, where o is $\wedge$ or $\vee$; (3) $\mathsf{xnf}(\phi_1 \mathcal{U} \phi_2) = \mathsf{xnf}(\phi_2) \vee (\mathsf{xnf}(\phi_1) \wedge \mathcal{X}(\phi_1 \mathcal{U} \phi_2))$; and (4) $\mathsf{xnf}(\phi_1 \mathcal{R} \phi_2) = \mathsf{xnf}(\phi_2) \wedge (\mathsf{xnf}(\phi_1) \vee \mathcal{X}(\phi_1 \mathcal{R} \phi_2))$; Since the construction is built on two expansion rules of Unitl and Release, and the expansion stops once the Until and Release are in the scope of Next, it preserves the equivalence $\phi \equiv \mathsf{xnf}(\phi)$, and the cost is at most linear. □

Proof of Lemma 1

Proof. Basically, for $s \in T(s_0, \sigma)$ $(\sigma \in \Sigma)$, since there is a propositional assignment A of $\mathsf{xnf}(\bigwedge s_0)^p$ such that $\sigma \supseteq L(A)$ and $s = X(A)$, s is reachable from s_0 in one step. Inductively, assume s is reachable from s_0 in k $(k \geq 1)$ steps. For $s' \in T(s, \sigma)$ $(\sigma \in \Sigma)$, similarly we have s' is reachable from s in one step. As a result, s' is reachable from s_0 in $k+1$ steps. □

Proof of Theorem 4

We first introduce the following lemma that is used for the proof.

Lemma 6. *s is a final state of T_ϕ, iff there is a finite trace ξ with $|\xi| = 1$ such that $\xi \models s$.*

Proof. From Definition 6, s is a final state iff there is a propositional assignment A of the Boolean formula $Tail \wedge (\mathsf{xnf}(s))^p$ and $Tail \in A$. Recall that every Next subformula in s is associated with $\neg Tail$, so $Tail \in A$ holds iff no Next subformula is in A, and thus iff $L(A) \models \mathsf{xnf}(s)^p$ holds. Let $\xi = \sigma$ $(\sigma \in \Sigma)$ such that $\sigma \supseteq L(A)$, and obviously $\xi \models s$. □

Now we start to prove Theorem 4.

Proof. $(\Rightarrow)$ Since ϕ is satisfiable, there is a finite trace $\xi \models \phi$. Assume $|\xi| = n (n > 0)$. Based on Theorem 2, there is a propositional assignment A_0 of $\mathsf{xnf}(\phi)^p$ such that $\xi \models \bigwedge A_0$. And according to Definition 5, there is a transition $s_1 \in T(s_0, \sigma_0)$ in T_ϕ where $s_0 = \phi$, $\sigma_0 \supseteq L(A_0)$ and $s_1 = X(A_0)$. Moreover, we have that $\xi_1 \models s_1$. Recursively,, we can prove that for $n > i \geq 0$, there is a transition $s_{i+1} \in T(s_i, \sigma_i)$ in T_ϕ such that $\sigma_i \supseteq L(A_i)$, $s_{i+1} = X(A_i)$ for some propositional assignment A_i of $\mathsf{xnf}(s_i)^p$, and $\xi_{i+1} \models s_{i+1}$ holds. For $i = n-1$, since $|\xi_i| = 1$ and $\xi_i \models s_i$, s_i is a final state according to Lemma 6, and it is reachable from s_0 based on Lemma 1.

$(\Leftarrow)$ Let s be a final state in T_ϕ, and it is reachable from the initial state s_0 from Lemma 1. Assume a run $r = s_0, \ldots, s_{n-1}, s (n >= 0)$ (when $n = 0$, $s = s_0$ is the initial state) of T_ϕ on $\xi' = \sigma_0, \sigma_1, \ldots, \sigma_{n-1}$ leads from ϕ to s. Moreover according to Lemma 6, there is a finite trace ξ'' with $|\xi''| = 1$ such that $\xi'' \models s$. Let $\xi = \xi' \cdot \xi'' = \sigma_0 \sigma_1, \ldots \sigma_n (n \geq 0)$ where $\xi'' = \sigma_n$, and now we prove that $\xi \models \phi$. The proof can be achieved by induction from n to 0. Basically, $(\xi_n = \sigma_n) \models s$ is obviously true. Inductively assume $\xi_i \models s_i$ for $n \geq i \geq 1$, so $\xi_{i-1} = \xi[i-1] \cdot \xi_i$ satisfies $\xi[i-1] \supseteq L$ and $\xi_i \models s_i$ for some $s_i \in T(s_{i-1}, L)$ from the definition of T_ϕ, which means $\xi_{i-1} \models s_{i-1}$. When $i = 0$, we prove that $(\xi = \xi_0) \models (s_0 = \phi)$. □

Proof of Lemma 3

Proof. First, CDLSC sets $\mathcal{C}[0] = \{\phi\}$ after checking $Tail \wedge \mathsf{xnf}(\phi)^p$ is unsatisfiable, which meets Item 2 of Definition 7. Secondly after each iteration $i \geq 0$, $try\_satisfy$ guarantees that $\{\phi\}$ is added into each $\mathcal{C}[i]$ if no model is found, which meets Item 1 of Definition 7. By enumerating Line 10 and 15 in $try\_satisfy$, we have that $\mathsf{xnf}(s) \wedge \neg\mathcal{X}(\mathcal{C}[i])$ is unsatisfiable for $s \in \mathcal{C}[i+1] (0 \leq i \leq |\mathcal{C}| - 1)$, which meets Item 3 of Definition 7. So $\mathcal{C}$ is a conflict sequence after each iteration with no model found. □

Proof of Theorem 7

Proof. CDLSC runs iteratively, so CDLSC terminates iff either the procedure $try\_satisfy$ or inv_found returns true for some iteration. From Lemma 3, $\mathcal{C}$ is a conflict sequence after each iteration if no model found . After each iteration, $try\_satisfy$ returns true iff a final state is found (Line 6-7) based on Theorem 5. Meanwhile, inv_found returns true iff ϕ is unsatisfiable because of Theorem 6. As a result, there is always such an iteration, after which CDLSC can terminate and terminate correctly. □