

A Unified Automata-Theoretic Approach to LTL_f Modulo Theories

Marco Faella^{a,*},¹ and Gennaro Parlato^{b,**},¹

^aUniveristy of Naples Federico II, Italy

^bUniversity of Molise, Italy

ORCID (Marco Faella): <https://orcid.org/0000-0001-7617-5489>, ORCID (Gennaro Parlato): <https://orcid.org/0000-0002-8697-2980>

Abstract. We present a novel automata-based approach to address linear temporal logic modulo theory (LTL_f^{MT}) as a specification language for data words. LTL_f^{MT} extends LTL_f by replacing atomic propositions with quantifier-free multi-sorted first-order formulas interpreted over arbitrary theories. While standard LTL_f is reduced to finite automata, we reduce LTL_f^{MT} to symbolic data-word automata (SDWAs), whose transitions are guarded by constraints from underlying theories. Both the satisfiability of LTL_f^{MT} and the emptiness of SDWAs are undecidable, but the latter can be reduced to a system of constrained Horn clauses, which are supported by efficient solvers and ongoing research efforts. We discuss multiple applications of our approach beyond satisfiability, including model checking and runtime monitoring. Finally, a set of empirical experiments shows that our approach to satisfiability works at least as well as a previous custom solution.

1 Introduction

Linear Temporal Logic (LTL), pioneered by Pnueli in 1977 [47], provides a powerful framework for analyzing temporal properties of dynamic systems, including computer programs, hardware, and protocols. Its applications include model checking and synthesis of digital circuitry, program verification, cybersecurity, and robotics. While standard LTL assumes an infinite sequence of observations (which is not always appropriate), LTL_f , introduced by De Giacomo and Vardi in 2013 [12], focuses on finite traces instead. Since then, LTL_f has gained prominence in AI and business process modeling [13].

Both LTL and LTL_f are tailored for reasoning about system behaviors defined using propositional logic formulas with Boolean variables—representing states or events. They cannot natively handle systems with data, whose states may involve variables from complex domains like integers or real numbers. To address this limitation, more expressive logics, such as first-order theories, may be required to specify properties of real-world systems operating on data.

We consider LTL_f^{MT} , an extension of LTL_f where atomic propositions are replaced by quantifier-free multi-sorted first-order formulas interpreted over arbitrary theories. This extension is a fragment of the logic introduced by Geatti et al. [27] and aligns with a recent effort on extending Monadic Second Order (MSO) logic to data trees [19].

LTL_f^{MT} can express temporal properties of infinite-state systems described by numerical variables. Notably, it allows the specification of relationships between the current and next values of a state variable, a powerful feature that renders the logic undecidable. For instance, the formula $G(\odot x = x + 1)$ signifies the constant increase of the variable x at each time step until the end of the finite trace.

We present a *unifying and general framework* to tackle various problems related to LTL_f^{MT} using symbolic data-word automata (SDWAs) and on constrained Horn clauses (CHCs). SDWAs extend finite automata (FAs) by equipping states and alphabet with values taken from possibly infinite domains, and their transitions are guarded by formulas taken from quantifier-free first-order theories. Notably, the emptiness of SDWAs can be reduced to the satisfiability of CHCs, a type of logic that has proven successful in program verification [5]. Efficient algorithmic solutions and tools for CHCs continue to improve, as witnessed by the competition CHC-COMP [22]. Our reduction to SDWAs, akin to the classical automata-theoretic approaches to linear temporal logic, enables us to address diverse decision problems for LTL_f^{MT} .

First, we handle the *satisfiability* problem for LTL_f^{MT} , an undecidable problem that can be reduced to the emptiness problem for SDWAs, and subsequently to the satisfiability problem of CHCs through the previously mentioned constructions. This enables the utilization of efficient procedures offered by various solvers, including the SMT solver Z3 [14]. We implemented our approach in a prototype tool and used it to assess the satisfiability of a benchmark set, which includes cases used in [27] and our running example. Geatti et al. [27] exclusively concentrate on satisfiability, presenting a sound and incomplete procedure for LTL_f^{MT} . They employ a tableau originally introduced for LTL by Reynolds [48], iteratively unfolding for an increasing number of steps. Its status is checked at each iteration using an SMT solver. While their approach adeptly handles the logic's time dimension, it lacks the capability to exploit the fixpoint features of modern CHC solvers. Consequently, contradictory temporal requests cannot be detected without a fixed horizon, as illustrated by the formula G_{ANDF} in Section 7. Overall, our experiments demonstrate that our approach handles all benchmarks belonging to our fragment, showcasing similar or superior performance compared to the tool implementing the approach of [27].

As a second application of our framework, we delve into the *model-checking* problem for infinite state-transition systems against LTL_f^{MT} specifications. Remarkably, to the best of our knowledge, this

* Corresponding Author. Email: m.faella@unina.it.

** Corresponding Author. Email: gennaro.parlato@unimol.it.

¹ Equal contribution.

problem has not been explored in the existing literature. Leveraging the closure of SDWAs under intersection, we can model the state-transition system itself as an SDWA, allowing us to intersect it with the SDWA derived from the negation of the specification, and subsequently check its emptiness. Given that the emptiness problem is addressed through CHC solving, we can again employ this efficient technology to resolve the model-checking problem, marking a significant advancement in this unexplored domain.

In the final application of our automata-based framework, we focus on the implementation of *runtime monitors* for LTL_f^{MT} specifications. These monitors serve to evaluate the correctness of a specific run based on log files or on-the-fly assessments. Our solution provides both preliminary and final verdicts, where a final verdict indicates that no change in the evaluation of the specification can occur. Such early termination (a.k.a. *anticipatory monitoring* [38]) is clearly desirable for performance reasons. Our approach involves tracking the state of the SDWA corresponding to the given specification and solving two CHC instances derived from the current state of the monitor. We also suggest more scalable solutions by employing over-approximations, albeit with slightly reduced precision.

Our technical contributions can be summarized as follows:

- We enhance a previously established automata construction designed for LTL_f^{MT} monitoring and linear arithmetic (refer to [23]) by incorporating an off-the-shelf LTL_f procedure and tool, specifically SPOT [18].
- We introduce SDWAs as a state-based representation of the syntactically convenient LTL_f^{MT}, and demonstrate its versatility across multiple applications.
- We leverage the connection between SDWAs and CHCs, illustrating that modern CHC solvers can effectively address fundamental problems related to LTL_f^{MT}.
- We present a series of satisfiability experiments showing that the CHC-based approach is simpler and more general than a previous custom approach (see [27]) but also exhibits comparable efficiency.

Supplementary material is in the technical report [21].

2 Data Words

In this section, we introduce data words that we use to represent executions involving variables with unbounded values.

Given two integers i, j , with $i \leq j$, we denote by $[i, j]$ the set of integers k satisfying $i \leq k \leq j$, and by $[j]$ the set $[1, j]$.

Data signatures and data alphabets. A *data signature* \mathcal{S} is a set of pairs $\{id_i : type_i\}_{i \in [n]}$, where id_i is a field name, and $type_i$ is the type of id_i . Common types include integers, floating point rationals and real numbers, the Boolean type \mathbb{B} and the bit vectors of fixed length. An *evaluation* ν of a data signature \mathcal{S} is a map that associates each field name id in \mathcal{S} with a value of the corresponding type, denoted by $\nu.id$. We denote the set of all evaluations of \mathcal{S} by $\mathcal{E}(\mathcal{S})$. Henceforth, we will use the term *symbol* to refer to an element in $\mathcal{E}(\mathcal{S})$, while we refer to the set $\mathcal{E}(\mathcal{S})$ as the *data alphabet* of \mathcal{S} .

Data languages. A *data word* over a data signature \mathcal{S} is a finite sequence $w = w_1 w_2 \dots w_n$ where each element w_i is a symbol of $\mathcal{E}(\mathcal{S})$. The *length* of a data word w , denoted by $|w|$, is the number of symbols in the sequence. The *empty data word*, denoted by ϵ , is the data word with no data symbols, i.e., $|\epsilon| = 0$. We denote the set of all data words over $\mathcal{E}(\mathcal{S})$ by $\mathcal{E}(\mathcal{S})^*$. A *data language* over $\mathcal{E}(\mathcal{S})$ is any subset of $\mathcal{E}(\mathcal{S})^*$. For any data word w of length n and $i \in [n]$, we denote its *prefix* $w_1 \dots w_i$ by $w_{\leq i}$, and its *suffix* $w_i \dots w_n$ by $w_{\geq i}$.

Data logic. This paper extends LTL_f by integrating data constraints expressed in first-order logic with equality, following standard syntax and semantics [44]. To handle multiple data types in the fields of symbols forming data words, we utilize many-sorted signatures. Specifically, we employ a many-sorted first-order logic \mathcal{D} with sorts $data_1, \dots, data_n$. Each $data_i$ has a corresponding logic \mathcal{D}_{data_i} , including function symbols of type $data_i^h \rightarrow data_i$ and relation symbols of type $data_i^h \rightarrow \mathbb{B}$, with arity h . These logics encompass features like integer or real arithmetic, arrays, etc. Henceforth, we assume that \mathcal{D} comprises constant symbols \mathcal{D}^{con} , relation symbols \mathcal{D}^{rel} , and function symbols \mathcal{D}^{fun} .

3 LTL_f Modulo Theories

We present LTL_f^{MT}, an extension of LTL_f [12], tailored for expressing data-word properties via embedded \mathcal{D} -formulas. LTL_f^{MT} is a fragment of Geatti et al.'s logic [27], distinguished by the absence of quantified variables, and it extends the logic of Felli et al. [23] by allowing more general data constraints. Henceforth, we assume that \mathcal{S} is the data signature of the data words under consideration.

Syntax. Terms of LTL_f^{MT} are defined as follows:

$$t := \bigcirc^n id \mid \bigcirc^n id \mid c \mid f(t_1, \dots, t_k),$$

where $n \geq 0$, id is a field in \mathcal{S} , $c \in \mathcal{D}^{\text{con}}$, and f is a k -arity function symbol in \mathcal{D}^{fun} . The grammar includes the *next term constructor* \bigcirc^n and the *weak next term constructor* \bigcirc^n . The superscript n denotes “ n steps in the future” (called *lookahead*). For instance, $\bigcirc^2 x$ represents the value of field x two steps ahead, requiring those steps to exist. $\bigcirc^2 x$ denotes the same value but does not require those steps to exist. Instead of $\bigcirc^0 id$ and $\bigcirc^0 id$, we can simply write id .

Formulas of LTL_f^{MT} are generated by the following grammar:

$$\varphi := p(t_1, \dots, t_k) \mid \neg \varphi \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 U \varphi_2 \mid X \varphi,$$

where $p \in \mathcal{D}^{\text{rel}}$ is a relation (or *predicate*) symbol of arity k applied to terms t_1, \dots, t_k . We use \neg for negation, \wedge for conjunction, and U for the *until operator*. The temporal operator X signifies *tomorrow* (distinct from *next* for clarity). In the grammars, $f(t_1, \dots, t_k)$ and $p(t_1, \dots, t_k)$ are assumed to be well-typed, ensuring that the function f and relation p can operate on the data types of t_1, \dots, t_k .

Semantics. The LTL_f^{MT} formulas are interpreted over data words, in accordance with an interpretation \mathbb{D} of the symbols of the data logic \mathcal{D} . For a given data word w , a term t within the formula is interpreted as a value $\llbracket t \rrbracket_w$ derived from the underlying data logic. However, if the term attempts to read beyond the end of the word, it will be assigned one of two special values: ∇ (a strong error) or $\tilde{\nabla}$ (a weak error). In particular, the semantics for function terms gives higher priority to strong errors compared to weak errors arising from the arguments. Formally:

$$\begin{aligned} \llbracket \bigcirc^n id \rrbracket_w &= \begin{cases} w_{n+1}.id & \text{if } |w| \geq n + 1 \\ \nabla & \text{otherwise.} \end{cases} \\ \llbracket \bigcirc^n id \rrbracket_w &= \begin{cases} w_{n+1}.id & \text{if } |w| \geq n + 1 \\ \tilde{\nabla} & \text{otherwise.} \end{cases} \\ \llbracket c \rrbracket_w &= \mathbb{D}(c) \\ \llbracket f(t_1, \dots, t_k) \rrbracket_w &= \begin{cases} \mathbb{D}(f)(\llbracket t_1 \rrbracket_w, \dots, \llbracket t_k \rrbracket_w) & \text{if } \llbracket t_i \rrbracket_w \notin \{\nabla, \tilde{\nabla}\} \text{ for all } i \\ \nabla & \text{if } \llbracket t_i \rrbracket_w = \nabla \text{ for at least one } i \\ \tilde{\nabla} & \text{otherwise.} \end{cases} \end{aligned}$$

Formulas of LTL_f^{MT} are interpreted over data words. For a data word w and a LTL_f^{MT} formula φ , we define the *satisfaction relation* $w \models \varphi$ (i.e., w is a model of φ) as follows:

$$\begin{aligned}
w \models p(t_1, \dots, t_k) & \text{ iff } \left(\llbracket t_i \rrbracket_w \notin \{\nabla, \tilde{\nabla}\} \text{ for all } i, \text{ and} \right. \\
& \quad \mathbb{D}(p)(\llbracket t_1 \rrbracket_w, \dots, \llbracket t_k \rrbracket_w) \text{ holds} \left. \right) \text{ or} \\
& \quad \left(\llbracket t_i \rrbracket_w \neq \nabla \text{ for all } i, \text{ and} \right. \\
& \quad \quad \left. \llbracket t_i \rrbracket_w = \tilde{\nabla} \text{ for at least one } i \right) \\
w \models \neg\varphi & \text{ iff } w \not\models \varphi \\
w \models \varphi_1 \wedge \varphi_2 & \text{ iff } w \models \varphi_1 \text{ and } w \models \varphi_2 \\
w \models \varphi_1 U \varphi_2 & \text{ iff there is } i \in \llbracket w \rrbracket \text{ s.t. } w_{\geq i} \models \varphi_2, \text{ and} \\
& \quad \text{for all } j \in [i-1] \text{ it holds } w_{\geq j} \models \varphi_1 \\
w \models X\varphi & \text{ iff } |w| > 1 \text{ and } w_{\geq 2} \models \varphi.
\end{aligned}$$

The semantics for predicates gives higher priority to strong errors over weak errors: if any of the terms within the predicate yields a strong error, the predicate is false; otherwise, if weak errors arise from the terms, the predicate is true. The predicate is evaluated according to its meaning in the data theory only when no errors occur. We define the *language* of an LTL_f^{MT} formula φ , denoted by $\mathcal{L}(\varphi)$, as the set of all data words $w \in \mathcal{E}(\mathcal{S})^*$ such that $w \models \varphi$.

Derived operators. We also introduce supplementary logical and temporal operators derived from the fundamental operators outlined in the grammar for defining formulas. The logical operators include \vee for disjunction, \rightarrow for implication, \leftrightarrow for bi-implication, and the constants *true* and *false*. The additional temporal operators are $F\varphi$ (eventually) and $G\varphi$ (globally), representing *true* $U\varphi$ and $\neg F\neg\varphi$, respectively. Using the standard Tomorrow operator X , we define the *weak Tomorrow operator* \tilde{X} as $\tilde{X}\varphi \stackrel{\text{def}}{=} \text{last} \vee X\varphi$, where *last* is the formula $\neg X \text{true}$. While X requires the existence of the next symbol in the word, \tilde{X} does not. Therefore, $\tilde{X}\varphi$ is always true on the last symbol of a data word. We write X^n and \tilde{X}^n to signify n repetitions of the corresponding operator.

Example 1. Consider a simple **temperature control system** with four variables: a Boolean *heat* indicating heating status, a real-valued *temp* that represents the temperature in the building, and two integer-valued variables *e* (energy consumption) and *t* (time). The system follows these rules:

- Time is measured in discrete units of hours.
- Heating increases temperature by 1.5 degrees per hour.
- Without heating, the temperature decreases by 1 degree hourly.
- Minimum heating duration is 4 hours.
- Minimum heating-off duration is 2 hours.
- Each heating hour consumes 1 unit of energy.

These rules can be expressed with the LTL_f^{MT} formula φ_{rules} :

$$\begin{aligned}
G \left(\right. & \quad \odot t = t + 1 \\
& \quad \wedge \text{heat} \rightarrow (\odot e = e + 1 \wedge \odot \text{temp} = \text{temp} + 1.5) \\
& \quad \wedge \neg \text{heat} \rightarrow (\odot e = e \wedge \odot \text{temp} = \text{temp} - 1) \\
& \quad \wedge (\neg \text{heat} \wedge X \text{heat}) \rightarrow (\tilde{X}^2 \text{heat} \wedge \tilde{X}^3 \text{heat} \wedge \tilde{X}^4 \text{heat}) \\
& \quad \wedge (\text{heat} \wedge X \neg \text{heat}) \rightarrow \tilde{X}^2 \neg \text{heat} \\
& \quad \left. \wedge e = 0 \wedge t = 0. \right)
\end{aligned}$$

The **satisfiability problem** (SAT) for a given LTL_f^{MT} formula φ with an associated data signature \mathcal{S} asks whether $\mathcal{L}(\varphi)$ is empty. The following theorem can be proved by reducing the halting problem for 2-counter machines [10] to it.

Theorem 1. *The SAT problem for LTL_f^{MT} is undecidable.*

Normal Form. We show that removing the weak next term constructor \odot from the syntax does not affect expressiveness. This result is used to streamline the automaton construction outlined later.

Lemma 1. *For all LTL_f^{MT} formulas φ , there is an equivalent LTL_f^{MT} formula φ' without the weak next term constructor \odot .*

Proof. Consider a subformula θ of φ of the form $p(t_1, \dots, t_k)$, where at least one term includes the weak next term constructor \odot . Let a (resp., b) denote the maximum lookahead of \odot (resp., \bigcirc) in t_1, \dots, t_k . For each $i \in [k]$, let t'_i be the same as t_i except that \odot^n is replaced by \bigcirc^n . Replace θ with $(p(t'_1, \dots, t'_k) \vee \neg X^a \text{true}) \wedge X^b \text{true}$. The formula φ' is obtained by applying the above substitution to all subformulas of the type $p(t_1, \dots, t_k)$ occurring in φ . It is straightforward to prove the equivalence of φ and φ' . \square

Then, we demonstrate that the lookahead can be limited to the values $\{0, 1\}$ while maintaining (un)satisfiability, at the cost of enlarging the signature, similarly to Lemma 7 in [23]. We do not use this result in the automaton construction in Section 5, because, while the enlarged signature is inconsequential for satisfiability checking, it makes model-checking and monitoring applications more cumbersome, and we want the automaton to work uniformly on all three applications.

Lemma 2. *For all LTL_f^{MT} formulas φ without \odot , there exists an equi-satisfiable LTL_f^{MT} formula φ' in which every term of the form \bigcirc^n satisfies $n \in \{0, 1\}$.*

Proof. Let \mathcal{S} be the signature of φ , and let n^* be the maximum lookahead of \bigcirc in φ . To decrease the lookahead to a maximum of 1, we enrich the signature with n^* extra copies of the data fields, i.e., we set:

$$\mathcal{S}' \stackrel{\text{def}}{=} \bigcup_{i=0}^{n^*} \{(id_i : \text{type}) \mid (id : \text{type}) \in \mathcal{S}\}.$$

Then, let ψ be the formula on \mathcal{S}' obtained from φ by replacing each term of the type $\bigcirc^n id$ with id_n . The formula φ' required by the statement is the following:

$$\varphi' \stackrel{\text{def}}{=} \psi \wedge G \bigwedge_{i=1}^{n^*} (\bigcirc^1 id_i = id_{i-1}).$$

Proving that φ and φ' are equisatisfiable is straightforward. \square

4 Symbolic Data-Word Automata

In this section, we introduce *Symbolic Data-Word Automata* (SDWAs), a new automaton type extending traditional finite state automata (FA) to work with data words. Unlike FAs, SDWAs utilize two distinct data signatures to define alphabets and state sets. The transition function includes constraints from a given data logic \mathcal{D} . SDWAs can be seen as a linear (as opposed to branching) variant of *symbolic data-tree automata* from [19]. In the next section, we use SDWAs as the pivotal model allowing us to address various decision problems related to LTL_f^{MT} , by reducing them to the emptiness problem for SDWAs.

Definition 1. A symbolic data-word automaton \mathcal{A} with data logic \mathcal{D} , is a 5-tuple $(\mathcal{S}^\Sigma, \mathcal{S}^Q, \psi^0, \psi^\Delta, \psi^F)$ where:

\mathcal{S}^Σ is the alphabet data signature defining the word alphabet $\Sigma = \mathcal{E}(\mathcal{S}^\Sigma)$. Each field in \mathcal{S}^Σ is typed with a sort of \mathcal{D} ;

\mathcal{S}^Q is the state data signature defining the set of states $Q = \mathcal{E}(\mathcal{S}^Q)$. Each field in \mathcal{S}^Q is typed with a sort of \mathcal{D} ;

$\psi^0(q)$ is a \mathcal{D} -formula on the free variable q of type \mathcal{S}^Q ,² defining the set of initial states, i.e., the set of all the elements $q \in Q$ such that $\psi^0(q)$ evaluates to true;

$\psi^\Delta(q, a, q')$ is the transition constraint, defined as a \mathcal{D} formula, where q and q' are variables of type \mathcal{S}^Q , and a is a variable of type \mathcal{S}^Σ ;

$\psi^F(q)$ is a \mathcal{D} formula on the free variable q , defining the set of final states $F \subseteq Q$, i.e., the set of all the elements $q \in Q$ such that $\psi^F(q)$ evaluates to true.

\mathcal{A} is designed to accept \mathcal{S}^Σ -words. A data word w is considered accepted by \mathcal{A} if there exists a function $\pi : [0, l] \rightarrow Q$, where $l = |w|$, such that the following conditions hold: (i) $\psi^0(\pi(0))$ is true, (ii) for all $i \in [0, l - 1]$, $\psi^\Delta(\pi(i), w_{i+1}, \pi(i + 1))$ is true, and (iii) $\psi^F(\pi(l))$ holds.

The set of all \mathcal{S}^Σ -words accepted by \mathcal{A} forms its language, denoted by $\mathcal{L}(\mathcal{A})$. \mathcal{A} is **deterministic** if (i) for all $q \in Q$ and $a \in \Sigma$ there is exactly one state $q' \in Q$ such that $\psi^\Delta(q, a, q')$ holds, and (ii) there is exactly one state q such that $\psi^0(q)$ holds. The following result is proved in the technical report [21] and used to address the model-checking problem in Section 6.2.

Theorem 2 (CLOSURE UNDER INTERSECTION). *Given two SDWAs \mathcal{A}_1 and \mathcal{A}_2 , we can effectively construct an SDWA \mathcal{A}_\cap such that $\mathcal{L}(\mathcal{A}_\cap) = \mathcal{L}(\mathcal{A}_1) \cap \mathcal{L}(\mathcal{A}_2)$.*

4.1 Solving the Emptiness Problem with CHCs

The *emptiness problem* for symbolic data-word automata consists in determining if a given SDWA \mathcal{A} recognizes any word, i.e., whether $\mathcal{L}(\mathcal{A})$ is empty.

The undecidability of the problem follows from Theorem 1 and the SDWA construction in Section 5. However, we show that the emptiness problem can be effectively reduced to the satisfiability of a CHC system, which are often solvable with efficient off-the-shelf solvers.

Constrained Horn Clauses. We fix a set R of uninterpreted fixed-arity relation symbols representing the unknowns in the system. A *Constrained Horn Clause*, or CHC, is a formula of the form $H \leftarrow C \wedge B_1 \wedge \dots \wedge B_n$ where:

- C is a constraint over the background data logic \mathcal{D} and does not include any application of symbols in R ;
- B_i is an application $p(v_1, \dots, v_k)$ of a relation symbol $p \in R$ to first-order variables v_1, \dots, v_k ;
- H is the clause *head* and, similarly to B_i , is an application $p(v_1, \dots, v_k)$ of a relation symbol $p \in R$ to the first-order variables, or it is *false*;
- all first-order variables in the signature of predicates and constraints are implicitly universally quantified.

A finite set \mathcal{H} of CHCs is a *system*, corresponding to the first-order formula obtained by taking the conjunction of all its CHCs.

² Since q is a variable of type $\mathcal{S}^Q = \{id_i : type_i\}_{i \in [n]}$, $\psi^0(q)$ becomes a \mathcal{D} -formula on the variables $q.id_1, \dots, q.id_n$.

The semantics of constraints is assumed to be given as a structure. A system \mathcal{H} with relation symbols R is *satisfiable* if there exists an interpretation for each predicate in R such that all clauses in \mathcal{H} are valid under that interpretation.

In constraint logic programming, every CHC system \mathcal{H} has a unique minimal model, computed as the fixpoint of an operator derived from its clauses [49, 36]. This semantics justifies the correctness of the reduction defined below.

Reduction. We propose a linear-time reduction from the emptiness problem for SDWAs to the satisfiability of CHC systems. The translation of SDWAs, akin to programs with scalar variables, into CHCs exhibits similarities with the use of proof rules in software verification [41, 30].

Given an SDWA $\mathcal{A} = (\mathcal{S}^\Sigma, \mathcal{S}^Q, \psi^0, \psi^\Delta, \psi^F)$, with structured variables q and q' of type \mathcal{S}^Q , a of type \mathcal{S}^Σ , and $h(\cdot)$ as an uninterpreted predicate, we map \mathcal{A} to the CHC system $\mathcal{H}_\mathcal{A}$, consisting of the following CHCs:

$$\mathcal{H}_\mathcal{A} \stackrel{\text{def}}{=} \begin{cases} h(q) & \leftarrow \psi^0(q) \\ h(q') & \leftarrow h(q) \wedge \psi^\Delta(q, a, q') \\ \text{false} & \leftarrow h(q) \wedge \psi^F(q). \end{cases}$$

Theorem 3 (EMPTINESS). *The language $\mathcal{L}(\mathcal{A})$ of an SDWA \mathcal{A} is empty iff the CHC system $\mathcal{H}_\mathcal{A}$ is satisfiable.*

5 The Automata-Theoretic Construction

This section constructs a deterministic SDWA equivalent to the provided LTL_f^{MT} formula φ in three steps:

1. Convert φ into an LTL_f formula φ' by abstracting away the data constraints.
2. Employ classical results to obtain a deterministic finite automaton $\mathcal{A}_{\varphi'}$ equivalent to φ' .
3. Convert $\mathcal{A}_{\varphi'}$ into an SDWA by checking the data constraints on the appropriate data.

The process is similar to a result for a fragment of the logic MSO-D [19]. In contrast, our paper presents a deterministic automaton and manages an arbitrary lookahead of the next term constructor.

5.1 Transforming to LTL_f

Here, we define the LTL_f formula $abs(\varphi)$, crafted to abstract data constraints from φ . Thanks to Lemma 1, we can assume that φ does not contain the \ominus constructor.

Consider all subformulas of the form $p(t_1, \dots, t_k)$ within φ , denoted as $\theta_1, \dots, \theta_m$. For each $i \in [m]$, introduce a new atomic proposition b_i , and let n_i be the maximum lookahead of the \bigcirc operator within θ_i . Construct the LTL_f formula $abs(\varphi)$ by replacing each occurrence of θ_i in φ with the formula $X^{n_i}b_i$. The rationale behind this transformation is that to establish the truth of θ_i , it may be necessary to read from the data word up to n_i positions ahead. As our ultimate goal is to construct a deterministic FA for φ , we defer the verification of θ_i by n_i steps, allowing the automaton to read and store all relevant data needed to establish its truth. Let $AP = \{b_1, \dots, b_m\}$ be the atomic propositions in $abs(\varphi)$.

To establish the semantic relation between φ and $abs(\varphi)$, we apply a similar abstraction to data words. Intuitively, we need to define the truth of each atomic proposition b_i based on the data present in the data word. To preserve the correspondence between $X^{n_i}b_i$ and

θ_i , the truth of b_i in a given position j in the abstract word must be the same as the value of θ_i in the data word n_i steps earlier than j . It is still not clear what should be the value of b_i in the first n_i positions of the abstract word, because for those positions there is not enough past data to evaluate θ_i . To solve this technical issue, we assume that each data type includes a *default value*, and we use those default values to evaluate θ_i when it reads beyond the start of the data word. This solution ends up attributing essentially arbitrary values to b_i in the early positions of the abstract word, but this does not thwart our construction. Indeed, since every occurrence of b_i in $\text{abs}(\varphi)$ appears in the context $X^{n_i}b_i$, $\text{abs}(\varphi)$ is insensitive to the value of b_i in the first n_i positions of the word.

Let def_k denote the data word of length k where all fields hold the default value of their type. For a data word w on the data signature \mathcal{S} of φ , let $\text{abs}(\varphi, w)$ be the word w' over the finite alphabet 2^{AP} , defined as follows:

- the length of w' is the same as the length of w ;
- for all $i = 1, \dots, m$ and $j = 1, \dots, |w|$,

$$b_i \in w'_j \iff \begin{cases} w_{\geq j-n_i} \models \theta_i & \text{if } j > n_i \\ \text{def}_{n_i-j+1} \cdot w \models \theta_i & \text{otherwise.} \end{cases}$$

Theorem 4. For all LTL_f^{MT} formulas φ and data words w on the same data signature, $w \models \varphi$ iff $\text{abs}(\varphi, w) \models \text{abs}(\varphi)$.

5.2 Building the SDWA

Recall the following well-known result about LTL_f:

Theorem 5 ([11]). For all LTL_f formulas φ , we can effectively construct an equivalent deterministic FA.

Let $\mathcal{A}_{\varphi'} = (2^{\text{AP}}, Q, q_0, \delta, F)$ be a deterministic FA equivalent to $\varphi' = \text{abs}(\varphi)$, where $\delta : Q \times 2^{\text{AP}} \rightarrow Q$. We define a deterministic SDWA $\mathcal{A}_{\varphi} = (\mathcal{S}, \mathcal{S}^Q, \psi^0, \psi^{\Delta}, \psi^F)$ with:

$\mathcal{S}^Q = \{\text{state} : Q\} \cup \{\text{data}^i : \mathcal{S}\}_{i \in [N]}$, where N is the maximum of the n_i 's, that is, the maximum lookahead of the \bigcirc term constructor in φ . In words, the state data signature of the automaton holds a state of the finite automaton $\mathcal{A}_{\varphi'}$ and N copies of the data signature of the original formula. The first component (i.e., *state*) is used to simulate a run of $\mathcal{A}_{\varphi'}$, and the second (i.e., *data*) stores the previous N symbols read from the data word. Specifically, the field data^i contains the symbol that was read i steps ago, as specified in the definition of ψ^{Δ} below.

$\psi^0(s)$ holds iff $s.\text{state} = q_0$, and all other fields contain the default value of their type.

$\psi^{\Delta}(s, a, s')$ holds iff (i) $s'.\text{data}^1 = a$, (ii) $s'.\text{data}^{i+1} = s.\text{data}^i$ for all $i \in [N-1]$, and (iii) it holds $\delta(s.\text{state}, \sigma) = s'.\text{state}$, where, for all $i \in [m]$:

$$b_i \in \sigma \text{ iff } s.\text{data}^{n_i} s.\text{data}^{n_i-1} \dots s.\text{data}^1 a \models \theta_i. \quad (1)$$

$\psi^F(s)$ holds iff $s.\text{state} \in F$.

Note that when the symbolic automaton begins reading a data word, its *buffer* of symbols kept in the *data* fields is initialized with default values, which are gradually replaced by those contained in the last N symbols read. Hence, the data constraints θ_i are initially evaluated over default data values, resulting in essentially arbitrary truth values. This is not a problem, because the formula φ' , which guides the *state* component of the symbolic automaton and prescribes which data constraints must be true at each step, is insensitive to the value of b_i in the first n_i positions of the word, as noted earlier.

We now state the main result of this section.

Theorem 6. For all LTL_f^{MT} formulas φ , we can effectively construct a deterministic SDWA \mathcal{A}_{φ} s.t. $\mathcal{L}(\mathcal{A}_{\varphi}) = \mathcal{L}(\varphi)$.

6 Applications

In this section, we demonstrate the versatility of our automata-theoretic approach in addressing diverse challenges associated with LTL_f^{MT}, namely satisfiability, model checking, and monitoring.

6.1 Satisfiability

The findings in the preceding sections offer a direct route to tackle the satisfiability problem for LTL_f^{MT}. Theorem 6 establishes that any LTL_f^{MT} formula can be converted into an equivalent SDWA. By leveraging Theorem 3, we transform the emptiness of an SDWA into the satisfiability of a system of CHCs. This reduction facilitates the application of advanced solution techniques, as illustrated in Section 7. There, we assess the satisfiability of diverse benchmark formulas using Z3 as the underlying CHC solver.

Theorem 7. The LTL_f^{MT} satisfiability problem can be effectively reduced to the emptiness problem for SDWAs.

6.2 Model Checking

The LTL_f^{MT} model-checking problem for an SDWA \mathcal{M} and an LTL_f^{MT} formula φ is to decide whether there exists a data word w in the language of \mathcal{M} that is a model of φ (i.e., $w \models \varphi$).

The undecidability of this problem stems from the undecidability of the emptiness problem for SDWAs. Nonetheless, we can readily devise a sound procedure for solving the model-checking problem by ultimately employing CHC solvers.

Our procedure is as follows. Firstly, we create an SDWA \mathcal{A}_{φ} such that $\mathcal{L}(\varphi) = \mathcal{L}(\mathcal{A}_{\varphi})$ (refer to Theorem 6). Subsequently, utilizing the construction associated with the closure of SDWAs under intersection (as detailed in the proof of Theorem 2), we generate a new SDWA \mathcal{A} such that $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{M}) \cap \mathcal{L}(\mathcal{A}_{\varphi})$. The determination of the answer to the model-checking problem relies on establishing whether $\mathcal{L}(\mathcal{A})$ is empty: a positive answer to the model-checking problem for \mathcal{M} and φ occurs iff the language of \mathcal{A} is non-empty.

Theorem 8. The LTL_f^{MT} model-checking problem can be effectively reduced to the emptiness problem for SDWAs.

When the CHC engine proves that the model-checking problem admits a positive answer, a data word satisfying the formula can be derived from the counterexample to the CHC system.

6.3 Runtime Monitoring

At times, the model checking of complex systems becomes intractable owing to the inherent complexity of the system. This leads to the adoption of less intricate verification approaches, such as runtime monitoring. In runtime monitoring, we continuously observe the system's execution and promptly notify any violation or fulfillment of the specified property. In this context, we delve into runtime monitoring concerning a given LTL_f^{MT} formula φ .

In defining the monitoring problem, we trace the satisfaction status of φ along a trace of \mathcal{M} , considering the trace fragment w that has been observed thus far. Following established conventions [3, 43, 23], we define $RV = \{\text{PS}, \text{CS}, \text{CV}, \text{PV}\}$ as the set of four *monitoring states*: current satisfaction (CS), permanent satisfaction (PS), current violation (CV), and permanent violation (PV).

Definition 2. For a given data word (trace) w and an LTL_f^{MT} formula φ , the **monitoring problem** is the task of determining the monitoring state $st \in RV$ that satisfies one of the following conditions:

- $st = CS$, $w \models \varphi$, and $ww' \not\models \varphi$ for some trace w' ;
- $st = PS$, $w \models \varphi$, and $ww' \models \varphi$ for every trace w' ;
- $st = CV$, $w \not\models \varphi$, and $ww' \models \varphi$ for some trace w' ;
- $st = PV$, $w \not\models \varphi$, and $ww' \not\models \varphi$ for every trace w' . \square

The following undecidability result should come as no surprise, and can be proved along similar lines as the previous ones.

Theorem 9. The monitoring problem for LTL_f^{MT} is undecidable (already for linear constraints).

We provide a sound procedure to solve the monitoring problem for LTL_f^{MT} . To achieve this, consider the deterministic SDWA $\mathcal{A}_\varphi = (\mathcal{S}, \mathcal{S}^Q, \psi^0, \psi^\Delta, \psi^F)$, as detailed in Theorem 6. Now, define two new SDWAs as follows:

$\mathcal{A}_\varphi(q, F)$ is derived from \mathcal{A}_φ by replacing the initial state with the state $q \in \mathcal{S}^Q$, and maintaining the same final states as \mathcal{A}_φ , i.e., $\psi^F(s) = (s.state \in F)$.

$\mathcal{A}_\varphi(q, \bar{F})$ is obtained from \mathcal{A}_φ by replacing the initial state with $q \in \mathcal{S}^Q$, and updating the final states with the formula $\psi^F(s) = (s.state \notin F)$.

Monitoring Procedure:

1. Utilizing a structured variable, denoted as q , we monitor the state of \mathcal{A}_φ as the analyzed system evolves. Initialization of q involves assigning it the unique state s that satisfies $\psi^0(s)$.
2. After each system transition, we update q with the unique state in which \mathcal{A}_φ moves to. We determine the monitoring state as follows:

CS: if $\psi^F(q) = true$, and $\mathcal{L}(\mathcal{A}_\varphi(q, \bar{F})) \neq \emptyset$;

PS: if $\psi^F(q) = true$, and $\mathcal{L}(\mathcal{A}_\varphi(q, \bar{F})) = \emptyset$;

CV: if $\psi^F(q) = false$, and $\mathcal{L}(\mathcal{A}_\varphi(q, F)) \neq \emptyset$;

PV: if $\psi^F(q) = false$, and $\mathcal{L}(\mathcal{A}_\varphi(q, F)) = \emptyset$.

The above procedure justifies the following result.

Theorem 10. The LTL_f^{MT} monitoring problem can be effectively reduced to the emptiness problem for SDWAs.

As per Theorem 9, the procedure might time out. One approach is to execute \mathcal{A}_φ along with the system and provide a conclusive answer only upon the system's termination. However, approximate solutions that anticipate an answer remain viable. By performing two emptiness checks on FAs corresponding respectively to $abs(\varphi)$ and to its complement (rather than on the SDWAs $\mathcal{A}_\varphi(q, \cdot)$), a correct answer is given when the result is PS or PV. If the result is CS (resp., CV), the actual monitoring state may be $\{CS, PS\}$ (resp., one of $\{CV, PV\}$). This imprecision does not introduce monitoring errors but extends monitoring until a definitive answer is attainable.

7 Satisfiability Experiments

In this section, we present experiments to evaluate the satisfiability of various LTL_f^{MT} formulas. The experiments include the temperature control system outlined in Example 1 and formulas from [27].




Figure 1. Architecture of the prototype implementation. Dashed transformations are performed manually.

Figure 1 illustrates the experimental toolchain. We manually translated each LTL_f^{MT} formula into LTL_f , as described in Section 5.1, and then into LTL, in order to feed it to the tool SPOT [18], generating a corresponding FA. Next, our script converts the FA into a CHC system encoding the emptiness of the corresponding symbolic SDWA, as detailed in Sections 5.2 and 4.1, representing each FA state as a distinct uninterpreted predicate. Finally, we use the SMT-solver Z3 to check the satisfiability of the CHC system. Theorem 3 establishes an inverse link between the satisfiability of the original LTL_f^{MT} formula and the CHC system. In this section, references to satisfiability refer to the original LTL_f^{MT} formula.

Experiments were performed on an AMD Ryzen 9 5900X (3.70Ghz) with 32GB of RAM and running Windows 10. The CHC solver was Z3, version 4.6.0 (64 bit). All formulas were also tested with the BLACK tool by Geatti et al. [27].

The experimental results are presented in Table 1.

Temperature Control. Reconsider Example 1, where the formula φ_{rules} describes a temperature control system. Suppose we want to verify the system's capability to bring the temperature back to at least 20°C after 24 hours, starting from 20°C and never dropping below 18°C, with an energy budget of N units, where N is a parameter. We express this requirement with the following formula:

$$\text{TEMPCTRL} \stackrel{\text{def}}{=} \varphi_{rules} \wedge (temp = 20) \wedge G(temp \geq 18) \\ \wedge X^{24}(e \leq N \wedge temp \geq 20).$$

Verifying TEMPCTRL satisfiability determines if the system can achieve the specified property. The resulting model provides an energy-efficient heating schedule. A binary search on N showed that at least 10 hours of heating is needed. All cases were resolved within seconds, with BLACK being slightly faster.

Formulas from Geatti et al. We analyzed all formulas from Table 2 in [27], which includes two families with constraints in linear integer arithmetic (LIA), denoted by LIA_1 and LIA_2 , and a family with constraints in linear real arithmetic (LRA), denoted LRA_1 . All formulas are parameterized by an integer N . Table 1 shows that our approach handles all instances, showing similar or superior performance. Notably, we can detect the unsatisfiability of LIA_1 with parameter -1 . Moreover, [27] contains two families of formulas that our approach cannot currently handle, because the theories employed are not supported by CHC solvers: one with a non-linear real constraint;³ and another with an uninterpreted integer function. Lastly, the authors of [27] point out that the unsatisfiability of the formula

$$\text{GANDF} \stackrel{\text{def}}{=} G(x > 3) \wedge F(x < 2)$$

cannot be proved by their method, whereas our approach promptly identifies the contradiction.

³ In that paper, that family is erroneously tagged with LRA , despite containing the non-linear term $\frac{1}{g}$.

Formula	N	Sat	# of CHCs	Time (sec.)	
				This paper	BLACK [27]
TEMPCTRL	6	No		4.3	< 1
	9	No		6.9	< 1
	10	Yes	202	2.3	< 1
	12	Yes		1.7	< 1
	24	Yes		< 1	< 1
LIA ₁	-1	No		< 1	timeout
	10	Yes	7	< 1	< 1
	100	Yes		< 1	< 1
	1000	Yes		6.6	5.2
LIA ₂	10	No	14	< 1	< 1
	50	No	54	2.7	4.8
	100	No	104	59	167.4
LRA ₁	10	Yes	18	< 1	< 1
	100	Yes	108	< 1	6.2
	1000	Yes	1008	6.9	timeout
GANDF	-	No	5	< 1	timeout

Table 1. Satisfiability experiments. Running times for this paper refer to the CHC solving time only, as the time to generate the CHCs is negligible. Timeout is 10 minutes.

8 Related Work

Our work is related to many works in the literature in different ways. In addition to the works discussed in the introduction, here we focus on those that seem to be closest to the results presented in this paper.

Data-aware Logics. Various proposals extend temporal logics with data-oriented features. Temporal Stream Logic [25] is an undecidable logic aimed at reactive synthesis, able to model the dynamic input-output relation of a reactive system while abstracting from specific data domains and operations. Further developments support specific data theories [24], develop a counterexample-guided abstraction refinement (CEGAR) loop involving LTL synthesis and SMT queries [42], and exploring connections with syntax-guided synthesis [8].

Other LTL extensions aim at retaining the decidability of satisfiability by focusing on simple data theories [16] or by avoiding comparisons between data at different positions of the word [17, 2, 29].⁴ Additionally, Geatti et al. [28] describe various decidable fragments of LTL_f^{MT} . In contrast, our logic (derived from [27]) is data-theory-agnostic and allows comparisons between data values at different positions using the next term constructor. These features radically raise the expressive power of the specifications and renders the satisfiability problem undecidable. Thus, the above extensions of LTL do not subsume our results, and our contributions do not subsume the state-of-the-art decidability results.

Symbolic Automata. There is a variety of models for symbolic automata, the most prominent being the symbolic finite automata of D’Antoni and Veanes [9]. In these automata, only the alphabet is possibly infinite, whereas the state space is finite. In particular, they do not allow to compare data values in different positions of the data word, which is a main concern of our investigation. Such symbolic finite automata form a strict and decidable subclass of our SDWAs. In contrast, SDWAs support infinite state spaces and their emptiness is undecidable, similarly to programs.

Several other models feature a hybrid state-space with a finite component that has a limited interaction with the infinite domain of

symbols being read. For instance, the Theory Mealy and Moore machines of Maderbacher and Bloem [42] pair a finite state-space with the possibility of storing data values and comparing them with predicates from the given data theory. Other models of finite automata on infinite alphabets include the register and pebble machines of Neven et al. [46] and the variable automata of Grumberg et al. [31].

CHCs for Program Verification. Our work builds upon the recent surge of interest in constrained Horn clauses for automated software model checking [6, 4]. Algorithms that efficiently solve systems of CHCs have been developed, often by adapting or extending techniques from automatic program verification [30, 32, 6]. CHCs offer a unique and elegant way to construct model checkers entirely through logic rules [20, 7, 26, 33, 35, 37, 39, 45]. This approach holds promise for efficient verification of temporal logic with data, especially considering the significant advancements made in CHC satisfiability solvers in recent years [22].

To leverage this potential, we introduce a novel automata-theoretic approach [40, 50]. We employ a translation technique that transforms SDWAs (or equivalently, transition systems) into CHCs. This translation aligns with existing applications of CHCs in software verification [41, 30]. Furthermore, our approach follows the growing trend of using automata theory in automated software model checking [34].

Runtime Monitors. Related monitoring problems have been studied in [15], focusing on comparing data values at arbitrary distances, and [23], with an emphasis on linear arithmetic and properties conducive to decidable solutions. Unlike [23], our approach, while possibly non-terminating, operates on more general data constraints. While our method might not always find a solution, it can handle a wider range of data constraints compared to these prior works.

9 Conclusions

In summary, we presented a novel automata-based approach supporting linear temporal logic modulo theory as a specification language for data words. Our framework translates LTL_f^{MT} into symbolic data-word automata, demonstrating their efficacy in capturing intricate temporal properties of executions of infinite-state systems.

Notwithstanding the undecidability of both LTL_f^{MT} satisfiability and SDWA emptiness, we showed that both problems can be reduced to solving CHCs, which allows us to capitalize on modern solvers. Empirical experiments substantiate the effectiveness of our approach, sometimes outperforming a previous custom solution and tool. These results highlight the practicality and broad applicability of our automata-based framework. Our approach seamlessly extends beyond satisfiability to encompass model checking and runtime monitoring in a unified way, effectively bridging three traditionally separate problems in prior research.

Future work. A promising direction for future work is integrating results from [23] and [27], into the domain of solving CHCs, to improve the effectiveness of CHC solvers in general and our approach in particular. Additionally, we see potential in extending our framework to accommodate other LTL dialects, such as CaRet logic [1] with data and context-free systems that involve data. This expansion could advance formal verification and monitoring by handling a wider range of temporal specifications.

Funding Information. Supported by INDAM-GNCS 2023-24, the MUR project SOP (Securing sOftware Platforms - CUP: H73C22000890001) part of the SERICS project (n. PE00000014 - CUP: B43C22000750006), *Verifica di proprietà di sicurezza nello sviluppo del software* under the Start-up 2022 program funded by the

⁴ Note that, in the context of model checking, comparing data in different positions can generally be achieved within the system, using extra storage. Hence, a more detailed comparison with the concurrent findings of Gianola et al. [29] merits further investigation.

Computer Science Division UNIMOL, and the MUR project Future AI Research (FAIR) Spoke 3.

References

- [1] R. Alur, K. Etessami, and P. Madhusudan. A temporal logic of nested calls and returns. In *TACAS'04, Barcelona, Spain, Mar 29 - Apr 2, volume 2988 of LNCS*, pages 467–481, 2004.
- [2] F. Baader, S. Borgwardt, and M. Lippmann. Temporal query entailment in the description logic shq. *Journal of Web Semantics*, 33:71–93, 2015.
- [3] A. Bauer, M. Leucker, and C. Schallhart. Comparing LTL semantics for runtime verification. *J. Log. Comput.*, 20(3):651–674, 2010.
- [4] T. A. Beyene, C. Popeea, and A. Rybalchenko. Efficient CTL verification via Horn constraints solving. In *HCVS'16, Eindhoven, The Netherlands, 3rd Apr 2016*, volume 219 of *EPTCS*, pages 1–14, 2016.
- [5] N. Bjørner, A. Gurfinkel, K. McMillan, and A. Rybalchenko. *Horn Clause Solvers for Program Verification*, pages 24–51. Springer International Publishing, 2015.
- [6] N. Bjørner, A. Gurfinkel, K. L. McMillan, and A. Rybalchenko. Horn clause solvers for program verification. In *Fields of Logic and Computation II*, volume 9300 of *LNCS*, pages 24–51, 2015.
- [7] A. Champion, T. Chiba, N. Kobayashi, and R. Sato. Ice-based refinement type discovery for higher-order functional programs. *J. Autom. Reason.*, 64(7):1393–1418, 2020.
- [8] W. Choi, B. Finkbeiner, R. Piskac, and M. Santolucito. Can reactive synthesis and syntax-guided synthesis be friends? In *PLDI'22, San Diego, CA, USA, June 13 - 17, 2022*, pages 229–243. ACM, 2022.
- [9] L. D'Antoni and M. Veanes. Automata modulo theories. *Communications of the ACM*, 64(5):86–95, 2021.
- [10] M. Davis. Recursive unsolvability of Post's problem of "Tag" and other topics in the theory of Turing machines. *annals of mathematics*, second series, vol. 74 (1961), pp. 437–455. *Journal of Symbolic Logic*, 31(4):654–655, 1966.
- [11] G. De Giacomo and M. Favorito. Compositional approach to translate LTL_f/LDL_f into deterministic finite automata. In *ICAPS'21, Guangzhou, China, Aug 2-13, 2021*, pages 122–130. AAAI Press, 2021.
- [12] G. De Giacomo and M. Y. Vardi. Linear temporal logic and linear dynamic logic on finite traces. In *IJCAI'13, Beijing, China, Aug 3-9, 2013*, pages 854–860. IJCAI/AAAI, 2013.
- [13] G. De Giacomo, R. De Masellis, M. Grasso, F. M. Maggi, and M. Montali. Monitoring business metaconstraints based on LTL and LDL for finite traces. In *BPM'14*, volume 8659 of *LNCS*, pages 1–17, 2014.
- [14] L. M. de Moura and N. S. Bjørner. Z3: an efficient SMT solver. In *TACAS'08*, volume 4963 of *LNCS*, pages 337–340, 2008.
- [15] N. Decker, M. Leucker, and D. Thoma. Monitoring modulo theories. *Int. J. Softw. Technol. Transf.*, 18(2):205–225, 2016.
- [16] S. Demri and D. D'Souza. An automata-theoretic approach to constraint LTL. *Inf. Comput.*, 205(3):380–415, 2007.
- [17] D. Dietsch, M. Heizmann, V. Langenfeld, and A. Podelski. Fairness modulo theory: A new approach to LTL software model checking. In *CAV'15, Proceedings, Part I*, volume 9206 of *LNCS*, pages 49–66, 2015.
- [18] A. Duret-Lutz, E. Renault, M. Colange, F. Renkin, A. G. Aisse, P. Schlehuber-Caissier, T. Medioni, A. Martin, J. Dubois, C. Gillard, and H. Lauko. From spot 2.0 to spot 2.10: What's new? In *CAV'22, Proceedings, Part II*, volume 13372 of *LNCS*, pages 174–187, 2022.
- [19] M. Faella and G. Parlato. Reasoning about data trees using CHCs. In *CAV'22, Haifa, Israel, Aug 7-10, 2022, Proceedings, Part II*, volume 13372 of *LNCS*, pages 249–271, 2022.
- [20] M. Faella and G. Parlato. Reachability games modulo theories with a bounded safety player. In *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI'23*, pages 6330–6337. AAAI Press, 2023. doi:10.1609/AAAI.V37I5.25779. URL <https://doi.org/10.1609/aaai.v37i5.25779>.
- [21] M. Faella and G. Parlato. A unified automata-theoretic approach to LTL_f modulo theories (Extended version), 2024. URL <https://arxiv.org/abs/2408.08817>.
- [22] G. Fedukovich and P. Rümmer. Competition report: CHC-COMP-21. In *HCVS'21, 28 Mar 2021*, volume 344 of *EPTCS*, pages 91–108.
- [23] P. Felli, M. Montali, F. Patrizi, and S. Winkler. Monitoring arithmetic temporal properties on finite traces. In *AAAI'23, Washington, DC, USA, Feb 7-14, 2023*, pages 6346–6354. AAAI Press, 2023.
- [24] B. Finkbeiner, P. Heim, and N. Passing. Temporal stream logic modulo theories. In *FOSSACS'22*, volume 13242 of *LNCS*, pages 325–346.
- [25] B. Finkbeiner, F. Klein, R. Piskac, and M. Santolucito. Temporal stream logic: Synthesis beyond the booleans. In *CAV'19, Proceedings, Part I*, volume 11561 of *LNCS*, pages 609–629, 2019.
- [26] P. Garoche, T. Kahsai, and X. Thirioux. Hierarchical state machines as modular horn clauses. In *HCV'16, Eindhoven, The Netherlands, 3rd Apr 2016*, volume 219 of *EPTCS*, pages 15–28, 2016.
- [27] L. Geatti, A. Gianola, and N. Gigante. Linear temporal logic modulo theories over finite traces. In *IJCAI'22, Vienna, Austria, 23-29 July 2022*, pages 2641–2647. ijcai.org, 2022.
- [28] L. Geatti, A. Gianola, N. Gigante, and S. Winkler. Decidable fragments of LTL_f modulo theories. In *ECAI'23, Sept 30 - Oct 4, 2023, Kraków, Poland*, volume 372 of *Frontiers in Artificial Intelligence and Applications*, pages 811–818. IOS Press, 2023.
- [29] A. Gianola, M. Montali, and S. Winkler. Linear-time verification of data-aware processes modulo theories via covers and automata. In *AAAI'24, Feb 20-27, 2024, Vancouver, Canada*. AAAI Press, 2024.
- [30] S. Grebenshchikov, N. P. Lopes, C. Popeea, and A. Rybalchenko. Synthesizing software verifiers from proof rules. In *PLDI'12, Beijing, China - June 11 - 16, 2012*, pages 405–416. ACM, 2012.
- [31] O. Grumberg, O. Kupferman, and S. Sheinvald. Variable automata over infinite alphabets. In *LATA'10, Trier, Germany, May 24-28, 2010. Proceedings 4*, pages 561–572.
- [32] A. Gurfinkel and N. Bjørner. The science, art, and magic of constrained horn clauses. In *SYNAS'19*, pages 6–10. IEEE, 2019.
- [33] A. Gurfinkel, T. Kahsai, A. Komuravelli, and J. A. Navas. The SeaHorn verification framework. In *CAV'15, San Francisco, CA, USA, July 18-24, 2015, Proceedings, Part I*, volume 9206 of *LNCS*, pages 343–361, 2015.
- [34] M. Heizmann, J. Hoenicke, and A. Podelski. Software model checking for people who love automata. In *CAV'13, Saint Petersburg, Russia, July 13-19, 2013. Proceedings*, volume 8044 of *LNCS*, pages 36–52, 2013.
- [35] H. Hojjat, F. Konečný, F. Garnier, R. Iosif, V. Kuncak, and P. Rümmer. A verification toolkit for numerical transition systems - tool paper. In *FM'12, Paris, France, Aug 27-31, 2012. Proceedings*, volume 7436 of *LNCS*, pages 247–251.
- [36] J. Jaffar and M. J. Maher. Constraint logic programming: A survey. *J. Log. Program.*, 19/20:503–581, 1994.
- [37] T. Kahsai, P. Rümmer, H. Sanchez, and M. Schäf. Jayhorn: A framework for verifying java programs. In *CAV'16, Toronto, ON, Canada, July 17-23, 2016, Proc., Part I*, volume 9779 of *LNCS*, pages 352–358.
- [38] H. Kallwies, M. Leucker, and C. Sánchez. General anticipatory monitoring for temporal logics on finite traces. In *Runtime Verification*, pages 106–125, Cham, 2023. Springer Nature Switzerland.
- [39] N. Kobayashi, R. Sato, and H. Unno. Predicate abstraction and CEGAR for higher-order model checking. In *PLDI'11, San Jose, CA, USA, June 4-8, 2011*, pages 222–233. ACM.
- [40] R. P. Kurshan. *Computer-aided Verification of Coordinating Processes: The Automata-theoretic Approach*. Princeton series in computer science. Princeton University Press, 1994.
- [41] S. La Torre, P. Madhusudan, and G. Parlato. Analyzing recursive programs using a fixed-point calculus. In *PLDI'09, Dublin, Ireland, June 15-21, 2009*, pages 211–222. ACM.
- [42] B. Maderbacher and R. Bloem. Reactive synthesis modulo theories using abstraction refinement. In *FMCAD'22, Trento, Italy, Oct 17-21, 2022*, pages 315–324. IEEE.
- [43] F. M. Maggi, M. Montali, M. Westergaard, and W. M. P. van der Aalst. Monitoring business constraints with linear temporal logic: An approach based on colored automata. In *BPM'11*, volume 6896 of *LNCS*, pages 132–147.
- [44] Z. Manna and C. G. Zarba. Combining decision procedures. In *Formal Methods at the Crossroads. From Panacea to Foundational Support, 10th Anniversary Colloquium of UNU/IIST, Lisbon, Portugal, Mar 18-20, 2002, Revised Papers*, volume 2757 of *LNCS*, pages 381–422, 2002.
- [45] Y. Matsushita, T. Tsukada, and N. Kobayashi. Rusthorn: CHC-based verification for rust programs. In *ESOP'20, Dublin, Ireland, Apr 25-30, 2020, Proceedings*, volume 12075 of *LNCS*, pages 484–514.
- [46] F. Neven, T. Schwentick, and V. Vianu. Finite state machines for strings over infinite alphabets. *ACM Transactions on Computational Logic (TOCL)*, 5(3):403–435, 2004.
- [47] A. Pnueli. The temporal logic of programs. In *FOCS'77, Providence, Rhode Island, USA, 31 Oct - 1 Nov 1977*, pages 46–57. IEEE.
- [48] M. Reynolds. A new rule for LTL tableaux. In *GandALF'16*, volume 226 of *EPTCS*, pages 287–301. Open Publishing Association.
- [49] M. H. van Emden and R. A. Kowalski. The semantics of predicate logic as a programming language. *J. ACM*, 23(4):733–742, 1976.
- [50] M. Y. Vardi and P. Wolper. An automata-theoretic approach to automatic program verification (preliminary report). In *LICS '86, Cambridge, Massachusetts, USA, June 16-18, 1986*, pages 332–344. IEEE.