A Note on the Fragment of Modal μ-Calculus Equivalent to First Order Logic over Infinite Trees

Massimo Benerecetti[1], Dario Della Monica[2], Angelo Matteo[2], Fabio Mogavero[1] and Gabriele Puppis[2]

[1]University of Napoli "Federico II", Italy
[2]University of Udine, Italy

Abstract

The expressive power of First-Order Logic (FO) over infinite trees remains poorly understood. In particular, it is still an open problem whether one can decide if a given regular language of infinite (or finite) trees is definable in FO. A common approach to better understand such formalisms is to provide alternative characterizations using, for example, automata or regular expressions; indeed, two automaton-based characterizations of FO have recently been introduced. In this note, we return to a purely logical perspective to look for syntactic variants of the well-known modal μ-calculus able to capture the expressiveness of FO. More precisely, we provably identify the exact fragment of the two-way modal μ-calculus that corresponds to FO. As a by-product, we obtain tight complexity results for the extension of the well-known logic CTL with past operators. Conversely, eliminating past operators yields a significantly more intricate landscape. We present a specific fragment of the standard μ-calculus that we conjecture to be expressively equivalent to FO, and discuss the main challenges in establishing a formal proof for this claim. This underscores the complex nature of FO over infinite trees and highlights how—unlike in word logics—past operators are crucial for expressiveness in tree logics.

Keywords

First Order Logic, Infinite trees, Modal μ-calculus

1. Introduction

First-Order Logic (FO) over infinite words is one of the most thoroughly understood formalisms in theoretical computer science. Its expressive power has been captured by a wide variety of equivalent characterizations, including temporal logics, automata classes and monoid classes. These equivalences, alongside several others, are extensively surveyed in [1].

The theoretical and practical impact of these characterizations over the past decades cannot be overstated. They have deeply enhanced our understanding of the expressive power of FO over infinite words, and the culmination of these efforts is an algorithm that decides, given an arbitrary regular language of infinite words L, whether L is definable in FO. This algorithm is a consequence of a variant of the celebrated Schützenberger Theorem [2].

The situation regarding the expressiveness of FO over infinite trees is completely different. While two tree temporal logics have been proven equivalent to FO through the decades [3, 4], until recently [5] there was no automaton-based characterization, and still there is no algebraic one. The lack of understanding regarding FO is a primary reason why the expressiveness and the behavior of logics of trees remain elusive. This note addresses this issue, trying to tackle the problem of precisely capturing FO expressiveness with respect to the *modal μ-calculus* ($\mathcal{L}_\mu$). The modal μ-calculus is the most expressive "temporal logic" in the context of regular tree languages, and it can be used as a benchmark to understand the expressive power of a given formalism. To this day, it is known that FO can be translated into the *alternation-free* fragment of $\mathcal{L}_\mu$, but no other information about the fragment of the modal μ-calculus equivalent to FO is available. Here, we aim to bridge this gap. We demonstrate that by employing past operators, one can readily isolate a fragment of $\mathcal{L}_\mu$ that is expressively equivalent to FO. Conversely,

8th International Workshop on Artificial Intelligence and fOrmal VERification, Logic, Automata, and sYnthesis. OVERLAY 2026 @ FLoC, July 18–19, Lisbon, Portugal

massimo.benerecetti@unina.it (M. Benerecetti); dario.dellamonica@uniud.it (D. Della Monica); angeo.matteo@uniud.it (A. Matteo); fabio.mogavero@unina.it (F. Mogavero); gabriele.puppis@uniud.it (G. Puppis)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073
published 2026-11-09

CEUR-WS.org/Vol-4324/paper13.pdf

we discuss how the situation changes in the absence of such operators and present a fragment of the standard $\mathcal{L}_\mu$ that we conjecture to be equivalent to FO, while making explicit the importance and the technical challenges of proving this claim.

2. The fragment of two-way modal μ-calculus equivalent to FO

In this section we show that, employing the variant of the modal μ-calculus ($\mathcal{L}_\mu$) that features *past* and *graded* modalities, called *two-way graded* modal μ-calculus, one can find quite easily a syntactic fragment expressively equivalent to FO over infinite trees. We consider unranked, unordered, and leafless trees (from now on, trees); specifically, trees where each node has a finite but arbitrary number of children, sibling order is irrelevant, and every branch extends infinitely. Let Σ be a finite alphabet. We define a Σ-tree as a tree whose nodes are labeled with symbols from Σ. We study *first-order logic* (FO) over these structures, specifically, monadic first-order logic extended with identity and ancestor predicates $=$ and $<$. In this framework, monadic predicates act as label tests, e.g., for a symbol $p \in \Sigma$, the formula $\exists x(p(x))$ states that there exists a node labeled p, while $x < y$ indicates that node x is an ancestor of node y. Hence, a formula such as $\exists x \exists y(x < y \wedge p(y))$ says that there is a node with a descendant labeled p. The identity is defined as usual. Thanks to [3, Theorem 4.5], there is a tree temporal logic equivalent to FO over trees. It is the following, called *Graded Polarized Computation Tree Logic with Past* (GPCTL$_\pm$):

$$\varphi \ ::= p \ \mid \ \neg\varphi \ \mid \ \varphi \vee \varphi \ \mid \ \mathsf{D}^n\varphi \ \mid \ \mathsf{EX}\varphi \ \mid \ \mathsf{E}\varphi\mathsf{U}\varphi \ \mid \ \mathsf{Y}\varphi \ \mid \ \varphi\mathsf{S}\varphi$$

where, as usual, p ranges over a set AP of atomic propositions, D^n is a counting modality, used to express the number of children of a given node that satisfy a property, while E is a path quantifier, X, U are future temporal operators, and Y, S are past temporal operators. For the semantics, see, e.g., [5].

We also consider *two-way graded* modal μ-calculus (2G$\mathcal{L}_\mu$), whose syntax is given by:

$$\varphi ::= p \mid X \mid \neg\varphi \mid \varphi \vee \varphi \mid \diamond_k\varphi \mid \Box_k\varphi \mid \ominus\varphi \mid \mu X.\varphi \mid \nu X.\varphi$$

The semantics, as well as the definition of a sentence, are standard, see, e.g., [6]. We restrict our past operator exclusively to the symbol $\ominus$: indeed, since our models are trees, every node has a *unique* past, rendering the distinction between $\diamond$ and $\Box$ in the past unnecessary. We denote by G$\mathcal{L}_\mu$ the fragment of the above syntax in which $\ominus$ is disallowed and by 2$\mathcal{L}_\mu$ the fragment in which we only allow $\diamond_k$ and $\Box_k$ with $k = 1$. We will show that a syntactic restriction of this logic yields a fragment expressively equivalent to FO. We call this fragment FO-2G$\mathcal{L}_\mu$.

Definition 1 (FO-2G$\mathcal{L}_\mu$ fragment)**.** *Let AP be a set of atomic propositions and let $X, Y, Z, \ldots$ be monadic second-order variables from a set* Var_2*. Formulas of* FO-2G$\mathcal{L}_\mu$ *are generated according to the following syntax, where $p \in AP$ and ψ is a formula of* FO-2G$\mathcal{L}_\mu$ *in which the variable X does not appear:*

$$\begin{aligned}\varphi ::= \ & p \mid X \mid \neg\varphi \mid \varphi \vee \varphi \mid \diamond_k\varphi \mid \Box_k\varphi \mid \ominus\varphi \mid \mu X.\psi \vee (\psi \wedge \ominus X) \mid \\ & \mu X.\psi \vee (\psi \wedge \diamond X)\end{aligned}$$

This fragment is evidently a direct rewriting of GPCTL$_\pm$, which naturally explains its equivalence to FO. In particular, we will show that GPCTL$_\pm$ is translatable into FO-2G$\mathcal{L}_\mu$ and that FO-2G$\mathcal{L}_\mu$ is instead translatable into a class of automata recently shown to be equivalent to FO. This yields the expressive equivalence. Actually, to obtain new (to the best of our knowledge) complexity results, we prove something more. Let *Graded Computation Tree Logic with Past* (GPCTL) be GPCTL$_\pm$ in which one allows in the syntax also formulas of the form $\mathsf{A}(\varphi\mathsf{U}\varphi)$, where A is the universal path quantifier. Then, PCTL and PCTL$_\pm$ are the restrictions of GPCTL and GPCTL$_\pm$ in which D^n is disallowed.

Lemma 1. *For any* GPCTL *formula φ, there is an equivalent* 2G$\mathcal{L}_\mu$ *formula $f(\varphi)$ whose size is linear with respect to φ. More precisely, the translation $f(\varphi)$ is in the alternation-free fragment of* 2G$\mathcal{L}_\mu$*.*

Proof sketch. By induction on the structure of the formulas in GPCTL, e.g., $f(\mathsf{D}^n\psi) = \diamond_n f(\psi)$. □

Proposition 1. *The satisfiability problem for* PCTL$_\pm$*,* PCTL*,* GPCTL$_\pm$ *and* GPCTL *is EXPTIME-complete.*

Proof. The lower bound comes from [7, Theorem 3.2]. The upper bound comes from the linear translation of Lemma 1 which applies not only to GPCTL, but also to all the other logics considered. Then, the satisfiability problem for 2G$\mathcal{L}_\mu$ can be solved in EXPTIME, according to [6, Theorem 5.2]. □

The following is implied by the translation above and the PSPACE-hardness of PCTL model checking [8]. The implication is striking since model checking for the alternation-free $\mathcal{L}_\mu$ is in linear time [9], further showing how the use of past can fundamentally alter properties of branching-time formalisms.

Proposition 2. *The model checking problem for the alternation-free 2$\mathcal{L}_\mu$ over trees is PSPACE-hard.*

Lemma 2. *For every formula φ of* GPCTL$_\pm$*, there is an equivalent* FO-2G$\mathcal{L}_\mu$ *sentence $f(\varphi)$ whose size is linear with respect to φ.*

Proof. This follows directly from Lemma 1. Specifically, because formulas of the form $\mathsf{A}(\varphi \mathsf{U} \varphi)$ are excluded from the GPCTL$_\pm$ grammar, we avoid needing the least fixpoint operator μ with its bound variable in the scope of a $\Box$ modality. □

For the other direction, we show that FO-2G$\mathcal{L}_\mu$ can be translated into *two-way graded linear hesitant tree automata*, whose exact definition can be found in [5], where it was shown that this automata class is equivalent to FO. We introduce the definition of these automata contextually within the proof below.

Lemma 3. *For every* FO-2G$\mathcal{L}_\mu$ *sentence, there is an equivalent two-way graded linear hesitant tree automaton.*

Proof. Fix a set of atomic propositions AP. First, we put our fragment in negation normal form, i.e., we consider formulas generated by the grammar:

$$\begin{aligned}\varphi ::= \; & p \mid \neg p \mid X \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \diamond_k \varphi \mid \Box_k \varphi \mid \ominus \varphi \mid \circ \varphi \mid \mu X.\psi \vee (\psi \wedge \ominus X) \mid \\ & \mu X.\psi \vee (\psi \wedge \diamond X) \mid \nu X.\psi \wedge (\psi \vee \Box X)\end{aligned}$$

Recall that ψ is a formula in the fragment not containing the variable X. We have $\circ\varphi = \neg \ominus \neg\varphi$. Moreover, $\mu X.\psi \vee (\psi \wedge \diamond X)$ and $\nu X.\psi \wedge (\psi \vee \Box X)$ are self-dual. Finally, we do not need to introduce a symbol for $\neg\mu X.\psi_1 \vee (\psi_2 \wedge \ominus X)$, since it is equivalent to $\mu X.(\neg\psi_1 \wedge (\neg\psi_2 \vee \circ\bot)) \vee (\neg\psi_1 \wedge \ominus X)$.

Consider a sentence φ of FO-2G$\mathcal{L}_\mu$ and let $sub(\varphi)$ be its set of subformulas, as it is usually defined. Let $\mathcal{A}_\varphi = \langle sub(\varphi), \Sigma, \delta, \varphi, F\rangle$. We let $\Sigma = 2^{AP}$ and define $F = \{\gamma \mid \gamma \text{ is a formula in } sub(\varphi) \text{ of the form } \nu X.\psi_1 \wedge (\psi_2 \vee \Box X)\}$. It remains to define δ. For the atomic and boolean cases, it is easy, so we account only for the remaining cases

- $\delta(\diamond_k\varphi, \sigma) = (\diamond_k, \varphi)$
- $\delta(\Box_k\varphi, \sigma) = (\Box_k, \varphi)$
- $\delta(\ominus\varphi, \sigma) = (\ominus, \varphi)$, if $\mathcal{A}_\varphi$ is not at the root,
- $\delta(\ominus\varphi, \sigma) = \bot$, if $\mathcal{A}_\varphi$ is at the root,
- $\delta(\circ\varphi, \sigma) = (\ominus, \varphi)$, if $\mathcal{A}_\varphi$ is not at the root,
- $\delta(\circ\varphi, \sigma) = \top$, if $\mathcal{A}_\varphi$ is at the root,
- $\delta(\mu X.\psi_1 \vee (\psi_2 \wedge \diamond X), \sigma) = \delta(\psi_1, \sigma) \vee (\delta(\psi_2, \sigma) \wedge (\diamond, \mu X.\psi_1 \vee (\psi_2 \wedge \diamond X)))$
- $\delta(\nu X.\psi_1 \wedge (\psi_2 \vee \Box X), \sigma) = \delta(\psi_1, \sigma) \wedge (\delta(\psi_2, \sigma) \vee (\Box, \nu X.\psi_1 \wedge (\psi_2 \vee \Box X)))$
- $\delta(\mu X.\psi_1 \vee (\psi_2 \wedge \ominus X), \sigma) = \delta(\psi_1, \sigma) \vee (\delta(\psi_2, \sigma) \wedge (\ominus, \mu X.\psi_1 \vee (\psi_2 \wedge \ominus X)))$

We now examine the structural constraints of a two-way graded linear hesitant tree automaton. Fundamentally, this is a standard graded alternating automaton [10], characterized by four additional properties:

- two-way: the automaton is allowed to navigate upward along the input tree;
- linear: the state-set can be partitioned into singleton components ordered by the subformula relation, and for any state $q \in Q$, the transition function $\delta(q, \sigma)$ only contains q itself or states strictly lower in the ordering;
- hesitant: every state is strictly classified as either transient, existential, universal, or upward. Roughly speaking, a state q is *existential* (resp., *universal*, *upward*), if, whenever $\delta(q, \sigma)$ contains q, it is in the form $(\diamond, q)$ (resp., $(\Box, q)$, $(\ominus, q)$), while it is transient if $\delta(q, \sigma)$ does not contain q at all;

- polarized: the set of accepting states is the set of universal states.

It is easy to see that the automaton $\mathcal{A}_\varphi$ is two-way and linear. In particular, linearity is possible because, in the fixpoint cases, formulas ψ_1 and ψ_2 do not contain the variable X, preventing any structural loops back to the initial fixpoint. Moreover, it is hesitant (again, see [5]) and polarized. □

By Lemmas 2 and 3 above, and the fact that GPCTL$_\pm$, FO, and two-way graded linear hesitant tree automata are equivalent [3, 5], it follows that:

Theorem 1. FO *and* FO-2G$\mathcal{L}_\mu$ *are equivalent formalisms over infinite trees.*

Next, we explore the challenge of identifying a G$\mathcal{L}_\mu$ fragment expressively equivalent to FO.

3. The fragment of standard modal μ-calculus equivalent to FO

It is relatively unusual to work with the two-way modal μ-calculus; indeed, it was studied only sporadically until it began receiving significant attention from the proof theory community over the last few years [11, 12, 13, 14]. Given this context, it is natural to investigate the fragment of *standard one-way* G$\mathcal{L}_\mu$ equivalent to FO. Interestingly, in the literature there is already a characterization of the G$\mathcal{L}_\mu$ fragment equivalent to *Weak Chain Logic* (WCL), an extension of FO featuring second-order quantification over *finite totally ordered sets* [15]. Moreover, there is a fragment of G$\mathcal{L}_\mu$ which has been shown to be at most as expressive as *Weak Tree Logic* (WTL), an extension of FO in which second-order quantification is restricted to *finite subtrees* of a tree [4]. We conjecture that the fragment of G$\mathcal{L}_\mu$ equivalent to FO is the intersection of these two fragments, produced by the following syntax.

Definition 2 (Conjectured FO-G$\mathcal{L}_\mu$)**.** *Let* AP *be a set of atomic propositions. Formulas of* FO-G$\mathcal{L}_\mu$ *are generated according to the following context-sensitive grammar:*

$$\begin{aligned}
\varphi_{Z,O} &::= p \mid \neg\varphi_{Z,O} \mid \varphi_{Z,O} \vee \varphi_{Z,O} \mid \Diamond_k\varphi_{O,\emptyset} \mid \Box_k\varphi_{O,\emptyset} \mid \theta_{\emptyset,\emptyset} \\
\theta_{Z,O} &::= \mu W.\alpha_{Z\cup\{W\},O\cup\{W\}} \mid \varphi_{Z,O} \\
\alpha_{Z,O} &::= Y \mid \varphi_{Z,O} \mid \alpha_{Z,O} \vee \alpha_{Z,O} \mid \alpha_{Z,O} \wedge \varphi_{\emptyset,O} \mid \Diamond\alpha_{O,\emptyset} \mid \mu X.\alpha_{Z\cup\{X\},O\cup\{X\}}
\end{aligned}$$

where $Y \in Z$.

Intuitively, the free second-order variables available are restricted to exactly those provided by the outer context Z. The primary purpose of the structural restrictions on $\alpha_{Z,O}$ (such as the asymmetric conjunction) is to ensure that second-order quantification is strictly confined to finite paths—that is, sets of nodes that are linearly ordered and connected. Furthermore, the Z and O parameter sets act as a syntactic starvation mechanism that strictly limits the modal depth between fixpoint unfoldings. By allowing at most one modality to occur before a fixpoint variable is reached, the grammar enforces what [4] refers to as "1-stepness". This guarantees that the denotation of any fixpoint variable forms a continuous, connected set of nodes. Without the Z and O mechanism, the logic could evaluate over disconnected regions of the tree, which requires expressive power strictly beyond that of FO. These are the main reasons why we conjecture this fragment to be equivalent to FO. Clearly, this syntax is far more intricate than that of FO-2G$\mathcal{L}_\mu$. However, proving our conjecture would yield profound theoretical insights—most notably, it would imply that FO *is the semantic intersection of* WCL *and* WTL, significantly deepening our understanding of FO's expressive power. Despite its importance, achieving this proof is far from trivial. If one were to apply the approach taken in the previous section, one would encounter a series of technical challenges: first, translating the FO-equivalent "one-way" temporal logic [5] into G$\mathcal{L}_\mu$ results in an exponential blow-up, yielding formulas too complex for meaningful analysis; second, the one-way automaton proven equivalent to FO in [5] is very complex, and translating FO-G$\mathcal{L}_\mu$ directly into it seems, for now, out of reach. Nevertheless, we are currently addressing the first problem by adapting the translation from CTL* to the modal μ-calculus introduced in [16] to our specific context. We believe that establishing such a characterization would significantly enhance the current understanding of FO's expressiveness over trees. Ideally, it would also make the long-standing problem of deciding whether a regular tree language is FO-definable more approachable.

Acknowledgments

This work was partially supported by the INdAM-GNCS Project "Neurosymbolic semantic contract for rigorous reachability and certificates in cyber-physical systems (CPS)" (CUP: E53C25002010001), Unione europea-Next Generation EU, Missione 4 Componente 1 (CUP: C63C22000770006), and Unione europea-Next Generation EU, Missione 4 Componente 1 (CUP: D33C22001300002).

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

References

[1] V. Diekert, P. Gastin, First-order definable languages., Logic and automata 2 (2008) 261–306.
[2] M. P. Schützenberger, On finite monoids having only trivial subgroups, Inf. Control. 8 (1965) 190–194.
[3] B.-H. Schlingloff, Expressive completeness of temporal logic of trees, Journal of Applied Non-Classical Logics 2 (1992) 157–180.
[4] M. Benerecetti, L. Bozzelli, F. Mogavero, A. Peron, Quantifying over trees in monadic second-order logic, in: 2023 38th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), IEEE, 2023, pp. 1–13.
[5] M. Benerecetti, D. Della Monica, A. Matteo, F. Mogavero, G. Puppis, An automaton-based characterisation of first-order logic over infinite trees, arXiv preprint arXiv:2509.14090 (2025).
[6] P. A. Bonatti, C. Lutz, A. Murano, M. Y. Vardi, The complexity of enriched mu-calculi, Logical Methods in Computer Science 4 (2008).
[7] A. Meier, M. Thomas, H. Vollmer, M. Mundhenk, The complexity of satisfiability for fragments of CTL and CTL*, International Journal of Foundations of Computer Science 20 (2009) 901–918.
[8] O. Kupferman, A. Pnueli, M. Y. Vardi, Once and for all, Journal of Computer and System Sciences 78 (2012) 981–996.
[9] A.-F. M. Mu-Calculus, A linear-time model-checking algorithm for the, in: Computer Aided Verification: 3rd International Workshop, CAV'91, Aalborg, Denmark, July 1-4, 1991. Proceedings, volume 575, Springer Science & Business Media, 1992, p. 48.
[10] O. Kupferman, U. Sattler, M. Y. Vardi, The complexity of the graded μ-calculus, in: International Conference on Automated Deduction, Springer, 2002, pp. 423–437.
[11] J. Kloibhofer, Y. Venema, Interpolation for the two-way modal μ-calculus, in: 2025 40th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), IEEE, 2025, pp. 155–168.
[12] B. Afshari, S. Enqvist, G. E. Leigh, J. Marti, Y. Venema, Proof systems for two-way modal m-calculus, The Journal of Symbolic Logic 90 (2025) 1211–1260.
[13] J. Rooduijn, Y. Venema, Focus-style proofs for the two-way alternation-free μ-calculus, in: International Workshop on Logic, Language, Information, and Computation, Springer, 2023, pp. 318–335.
[14] Y. Tanabe, K. Takahashi, M. Yamamoto, A. Tozawa, M. Hagiya, A decision procedure for the alternation-free two-way modal μ-calculus, in: International Conference on Automated Reasoning with Analytic Tableaux and Related Methods, Springer, 2005, pp. 277–291.
[15] F. Carreiro, Y. Venema, et al., Pdl inside the μ-calculus: a syntactic and an automata-theoretic characterization, Advances in Modal Logic 10 (2014) 74–93.
[16] M. Dam, CTL* and ECTL* as fragments of the modal μ-calculus, Theoretical Computer Science 126 (1994) 77–96.