A Dichotomy Theorem for Ordinal Ranks in MSO*

Damian Niwiński
Institute of Informatics, University of Warsaw, Poland
niwinski@mimuw.edu.pl

Paweł Parys
Institute of Informatics, University of Warsaw, Poland
parys@mimuw.edu.pl

Michał Skrzypczak
Institute of Informatics, University of Warsaw, Poland
mskrzypczak@mimuw.edu.pl

Recall that a least fixed point $\mu X.F(X)$ in a given structure can be computed by starting from the empty set and then iteratively applying F; in general, a transfinite number of iterations is needed, where for limit ordinals we take the union of all sets obtained so far. The necessary number of iterations depends both on a formula and on a structure. We may however ask if a formula F admits some upper bound: an ordinal $\mathrm{rank}(F)$, called the closure ordinal of $\mu X.F(X)$, such that in every structure the fixed-point is reached in at most $\mathrm{rank}(F)$ steps. Assuming that the variable X does not occur in F in the scope of any further fixed-point operations, we prove that either the closure ordinal of $\mu X.F(X)$ is uncountable, or strictly smaller than ω^2. We also provide a way of deciding which of these cases holds. The same result (with minor differences) was shown by Afshari, Barlucchi, and Leigh at FICS 2024; we provide an alternative proof, by reducing to an analogous dichotomy for MSO-definable relations over the infinite binary tree. The main part of the proof is then done in the framework of finite automata over the infinite binary tree, and involves a game-based technique.

A full version of our result was presented at STACS 2025 [19].

In the context of μ-calculus, one asks how many iterations are needed to reach a fixed point; this aspect concerns complexity of model checking (cf. [7, 13]), as well as expressive power of the logic (cf. [7, 6]). Recall that, in an infinite structure, a least fixed point $\mu X.F(X)$ is in general reached in a transfinite number of iterations: $\emptyset, F(\emptyset), \dots, F^{\gamma}(\emptyset), \dots$, where, for a limit ordinal γ, $F^{\gamma}(\emptyset) = \bigcup_{\xi<\gamma} F^{\xi}(\emptyset)$. It is therefore natural to ask if, for a given formula, one can effectively find a *closure ordinal* $\mathrm{rank}(F)$, such that, in any model, the fixed point can be reached in $\mathrm{rank}(F)$, but in general not less, iterations. This notion generalises in a natural way to a vectorial least fixed point $\mu\vec{X}.\vec{F}(\vec{X})$, where $\vec{F}(\vec{X})$ is a tuple of formulae $F_1(\vec{X}),\dots,F_k(\vec{X})$ and $\vec{X}$ is a tuple of variables $X_1,\dots,X_k$.

For example, consider $F(X) = \Box^5\bot \wedge (\Diamond X \vee \Box\bot)$. It holds in a node x if there are no paths of length 5 from x, and additionally either some successor of x is in X or x has no successors. Then $\mu X.F(X)$ holds in a node x if there are no paths of length 5 from x. Syntactically, the formula also requires that a node without successors is reachable from x; this automatically follows from the first part. Let us see how the fixed point $\mu X.F(X)$ is reached by iterating F. The first iteration $F(\emptyset)$ holds in nodes without successors. The second iteration $F(F(\emptyset))$ holds also in nodes having a successor without successors (but restricted to nodes from which there are no paths of length 5). Then $F(F(F(\emptyset)))$ adds the next level: nodes from which in two steps you can reach a node without successors; likewise $F^4(\emptyset)$. Finally $F^5(\emptyset)$ contains already all nodes from which there are no paths of length 5, because from every such node you can reach a node without successors in at most 4 steps. Thus the closure ordinal of $\mu X.F(X)$ is 5.

Note that the same set can be expressed without any fixed point, writing just $\Box^5\bot$, or with a trivial fixed point $\mu X.\Box^5\bot$, stabilizing after a single step. It follows that the closure ordinal is a property of a formula, not of defined sets; equivalent formulae may have different closure ordinals.

*All authors are supported by the National Science Centre, Poland (grant no. 2021/41/B/ST6/03914).

Submitted to:
FICS 2026

Fontaine [15] effectively characterised the formulae such that in each model the fixed point is reached in a finite number of steps (i.e., having closure ordinal at most ω). Czarnecki [12] observed that some formulae have no closure ordinals but, for each ordinal $\eta < \omega^2$, there is a formula whose closure ordinal is η. He also raised the following question.

Question (Czarnecki [12])**.** Is there a μ-calculus formula of the form $\mu X.F(X)$ that has a countable closure ordinal $\mathrm{rank}(F) \geq \omega^2$?

Gouveia and Santocanale [16] exhibited an example of a formula (with an essential alternation of the least and greatest fixed-point operators) whose closure ordinal is ω_1; clearly this limit can be achieved only in uncountable models. In general, it remains open whether a formula $\mu X.F(X)$ of the μ-calculus may have a *countable* closure ordinal $\mathrm{rank}(F) \geq \omega^2$. Results are known for fragments of μ-calculus. Afshari and Leigh [2] claimed a negative answer for formulae of the alternation-free fragment of μ-calculus; however, as the authors have later admitted [1], the proof contained some gaps. In a recent paper [1], Afshari, Barlucchi, and Leigh update the proof and extend the result to formulae of the so-called Σ-*fragment* of the μ-calculus. More specifically, the authors consider vectorial fixed points $\mu\vec{X}.\vec{F}(\vec{X})$, where the formulae in $\vec{F}$ may contain closed sub-formulae of the full μ-calculus, but the variables of $\vec{X}$ do not fall in the scope of any fixed-point operators. The authors show that if a countable number of iterations η suffices to reach the least fixed point of such a system in any countable structure then $\eta < \omega^2$.

In our work, we obtain an alternative proof of the same theorem, slightly generalised. One difference is that Afshari, Barlucchi, and Leigh consider only countable models. We prove that this does not change anything: if a formula requires some countable number of iteration in some (uncountable) model, then the same holds in some countable model. In other words, while considering only countable closure ordinals, it does not matter if in their definition we quantify only over countable models, or over all models. Second, Afshari, Barlucchi, and Leigh require that the considered formulae are guarded, meaning that all occurrences of the variables $X_1,\dots,X_k$ should be in the scope of some modal operator. While it is known that every formula is equivalent to a guarded one, it was not clear whether the translation preserves the closure ordinal (as we have seen, equivalent formulae may have different closure ordinals). We show how to perform this conversion so that it does not affect the obtained theorem (the closure ordinal may change slightly, but only by a finite number). Third, we provide a decision procedure, which was not at all present in their paper.

Theorem 1. *Let $\vec{F}(\vec{X}) = (F_1(\vec{X}),\dots,F_k(\vec{X}))$ with $\vec{X} = (X_1,\dots,X_k)$ be a tuple of μ-calculus formulae in which the variables $X_1,\dots,X_k$ do not occur in the scope of any fixed-point operator. Then, the closure ordinal of $\mu\vec{X}.\vec{F}(\vec{X})$ is either strictly smaller than ω^2, or at least ω_1, and it can be effectively decided which of the cases holds. In the former case, it is possible to compute a number $N \in \mathbb{N}$ such that the closure ordinal is bounded by $\omega \cdot N$.*

Note that we allow arbitrary closed formulae of μ-calculus to be nested in $\vec{F}$; however, we do not cover the whole μ-calculus, because of the restriction on occurrences of the variables of $\vec{X}$. This stays in line with the fragment considered by Afshari, Barlucchi, and Leigh [1].

While the proof of Afshari, Barlucchi, and Leigh [1] analyses the μ-calculus formulae, we proceed in a completely different way. Namely, we first (in a rather standard way) prove that it is enough to consider models which are countable trees. Then, we switch to the monadic second-order logic (MSO), taking into account the fact that over such models MSO and μ-calculus are equiexpressible, and we further move from trees of arbitrary branching to the infinite binary tree. The question about closure ordinals is translated to a question about MSO-definable relations of well-founded sets, which we now explain. The later question is interesting on its own.

The concept of a well-founded relation plays a central role in foundations of mathematics. It gives rise to ordinal numbers, which underlie the basic results in set theory, for example that any two sets can be compared in cardinality. Well-foundedness is no less important in the realm of computer science, where it often underlies the proofs of *termination* of non-deterministic processes, especially when no efficient bound on the length of a computation is known. In such cases, the complexity of possible executions is usually measured using an ordinal called *rank*. Such a rank can be seen as a measure of the *depth* of the considered partial order, taking into account suprema of lengths of possible descending chains. Estimates on a rank can provide upper-bounds on the computational complexity of the considered problem [22].

In this work, we adopt the perspective of mathematical foundations of program verification and model-checking. We focus on the monadic second-order logic (MSO) interpreted in the infinite binary tree (with the left and right successors as the only non-logical predicates), which is one of the reference formalisms in the area. The famous Rabin Tree Theorem [21] established its decidability, but—half a century after its introduction—the theory is still an object of study. On the one hand, it has led to numerous extensions, often shifting the decidability result far beyond the original theory (see e.g. [4, 20]). On the other hand, a number of natural questions regarding Rabin's theory remain still open, including a large spectrum of *simplification* problems. For example, we still do not know whether we can decide if a given formula admits an equivalent form with all quantifiers restricted to finite sets. Similar questions have been studied in related formalisms like μ-calculus or automata; for example if we can effectively minimise the Mostowski index of a parity tree automaton [11, 14], or the $\mu\nu$-alternation depth of a μ-calculus formula [5].

On the positive side, some decidability questions have been solved by reductions to the original theory. For example, it has been observed [17] that for a given formula $\varphi(\vec{X})$, the *cardinality* of the family of tuples of sets $\vec{X}$ satisfying $\varphi(\vec{X})$ can be computed; this cardinality can be either finite, $\aleph_0$, or $\mathfrak{c}$. Later on, Bárány, Kaiser, and Rabinovich [3] proved a more general result: they studied *cardinality quantifiers* $\exists^{\geq\kappa} X.\, \varphi(\vec{Y}, X)$, stating that there are at least κ distinct sets X satisfying $\varphi(\vec{Y}, X)$, and showed that these quantifiers can be expressed in the standard syntax of MSO; thus the extended theory remains decidable.

In the present work, instead of asking *how many* sets X witness to the formula $\exists X.\, \varphi(\vec{Y}, X)$, we ask how *complex* these witnesses must be in terms of their depth-and-branching structure. A set X of nodes of a tree is *well-founded* if it contains no infinite chain with respect to the descendant order. Recall that an ordinal η can be seen as the linearly ordered set of all ordinals smaller than η. Given a set X, a *counting function* is a function $\mathfrak{C}\colon X \to \eta$ such that $\mathfrak{C}(u) < \mathfrak{C}(v)$ whenever $u \in X$ is a proper descendant of $v \in X$ (such a function exists for some η whenever X is well-founded). The *rank* of a well-founded set X is the least ordinal η for which a counting function from X to η exists. Intuitively, the smaller $\mathrm{rank}(X)$, the simpler the set X is, in terms of its *branching structure*.

We consider formulae of the form $\exists X.\, \varphi(\vec{Y}, X)$, where φ is an arbitrary formula of MSO (it may contain quantifiers). We assume that, whenever the formula is satisfied for some valuation of variables $\vec{Y}$, the value of X witnessing the formula is a well-founded set. Note that well-foundedness of a set is expressible in MSO (it suffices to say that each branch contains only finitely many nodes in X), hence the requirement can be expressed within φ. For a fixed formula as above, we search for a minimal ordinal $\mathrm{rank}(\varphi)$, such that the rank of a witness can be bounded by $\mathrm{rank}(\varphi)$,

$$\mathrm{rank}(\varphi) \stackrel{\text{def}}{=} \sup_{\vec{Y}} \min_{X} \; \mathrm{rank}(X), \tag{1}$$

where $\vec{Y}$ and X range, as expected, over the values satisfying $\varphi(\vec{Y}, X)$.

Since $\mathrm{rank}(\varphi)$ is a supremum of countable ordinals, its value is at most ω_1 (the least uncountable ordinal). This value can be achieved, for example, by the formula "*$X = Y$ and Y is a well-founded set*", as there are well-founded sets of arbitrarily large countable ranks. On the other hand, for each pair of natural numbers (k,l), one can construct a formula φ with $\mathrm{rank}(\varphi) = \omega \cdot k + l$ in an analogous way to the μ-calculus formulae of Czarnecki [12]. Our main theorem for MSO (which implies Theorem 1) shows that no other ordinals can be obtained:

Theorem 2. *For any formula $\exists X.\, \varphi(\vec{Y}, X)$ as above, the ordinal $\mathrm{rank}(\varphi)$ is either strictly smaller than ω^2 or equal to ω_1. Moreover, it can be effectively decided which of the cases holds. In the former case, it is possible to compute a number $N \in \mathbb{N}$ such that $\mathrm{rank}(\varphi) < \omega \cdot N$.*

We also show that, in contrast to the aforementioned cardinality quantifiers, the property that $\mathrm{rank}(X)$ is smaller than ω^2 cannot be expressed directly in MSO.

The proof of Theorem 2 develops the game-based technique used previously to characterise certain properties of MSO-definable tree languages (see e.g. [9, 24]). Each application of this technique requires a specific game, designed in a way which reflects the studied property. The game should have a finite arena and be played between two perfectly-informed players $\exists$ and $\forall$. The winning condition of the game is given by a certain ω-regular set. Then, the seminal result of Büchi and Landweber [8] yields that the game is determined and the winner of this game can be effectively decided. The construction of the game is such that a winning strategy of each of the players provides a witness of either of the considered possibilities; in our case: if $\exists$ wins then $\mathrm{rank}(\varphi) = \omega_1$ and if $\forall$ wins then $\mathrm{rank}(\varphi) < \omega \cdot N$ for some computable N. The heart of this approach lies in a proper definition of the game, so that both these implications actually hold.

Similar dichotomies have been discovered in the studies of decision problems related to topological complexity of MSO-definable tree languages. An open problem related to the aforementioned question of definability with finite-set quantifiers, is whether we can decide if a tree language belongs to the Borel hierarchy (in general, it needs not). A positive answer is known if a tree language is given by a deterministic parity automaton [18], based on the following dichotomy: such a language is either Π^1_1-complete (very hard), or on the level Π^0_3 (relatively low) of the Borel hierarchy. Skrzypczak and Walukiewicz [24] gave a proof in the case when a tree language is given by a non-deterministic Büchi tree automaton, inspired by a rank-related dichotomy conjectured in the previous work of the first author [23, Conjecture 4.4]. There, an ordinal has been associated with each Büchi tree automaton, and it turns out (in view of [24]) that this ordinal either equals ω_1 or is smaller than ω^2, in which case the tree language is Borel. It should be mentioned that a procedure to decide whether a Büchi definable tree language is weakly definable (without the topological counterpart) was given earlier by Colcombet et al. [10].

References

[1] Bahareh Afshari, Giacomo Barlucchi & Graham E. Leigh (2025): *The Limit of Recursion in State-based Systems*. In: *FICS*, *Electronic Proceedings in Theoretical Computer Science* 435, Open Publishing Association, pp. 1–12, doi:10.4204/EPTCS.435.1.

[2] Bahareh Afshari & Graham E. Leigh (2013): *On Closure Ordinals for the Modal mu-Calculus*. In: *CSL*, *LIPIcs* 23, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 30–44, doi:10.4230/LIPICS.CSL.2013.30.

[3] Vince Bárány, Łukasz Kaiser & Alexander M. Rabinovich (2010): *Expressing Cardinality Quantifiers in Monadic Second-Order Logic over Trees*. *Fundam. Inform.* 100(1-4), pp. 1–17, doi:10.3233/FI-2010-260.

[4] Mikołaj Bojańczyk (2014): *Weak MSO+U with Path Quantifiers over Infinite Trees*. In: *ICALP (2)*, *Lecture Notes in Computer Science* 8573, Springer, pp. 38–49, doi:10.1007/978-3-662-43951-7_4.

[5] Julian C. Bradfield (1998): *Simplifying the Modal mu-Calculus Alternation Hierarchy*. In: *STACS*, *Lecture Notes in Computer Science* 1373, Springer, pp. 39–49, doi:10.1007/BFB0028547.

[6] Julian C. Bradfield, Jacques Duparc & Sandra Quickert (2005): *Transfinite Extension of the mu-Calculus*. In: *CSL*, *Lecture Notes in Computer Science* 3634, Springer, pp. 384–396, doi:10.1007/11538363_27.

[7] Julian C. Bradfield & Igor Walukiewicz (2018): *The mu-Calculus and Model Checking*. In: *Handbook of Model Checking*, Springer, pp. 871–919, doi:10.1007/978-3-319-10575-8_26.

[8] Julius R. Büchi & Lawrence H. Landweber (1969): *Solving Sequential Conditions by Finite-State Strategies*. *Transactions of the American Mathematical Society* 138, pp. 295–311.

[9] Lorenzo Clemente & Michał Skrzypczak (2021): *Deterministic and Game Separability for Regular Languages of Infinite Trees*. In: *ICALP*, *LIPIcs* 198, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 126:1–126:15, doi:10.4230/LIPICS.ICALP.2021.126.

[10] Thomas Colcombet, Denis Kuperberg, Christof Löding & Michael Vanden Boom (2013): *Deciding the Weak Definability of Büchi Definable Tree Languages*. In: *CSL*, *LIPIcs* 23, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 215–230, doi:10.4230/LIPICS.CSL.2013.215.

[11] Thomas Colcombet & Christof Löding (2008): *The Non-deterministic Mostowski Hierarchy and Distance-Parity Automata*. In: *ICALP (2)*, *Lecture Notes in Computer Science* 5126, Springer, pp. 398–409, doi:10.1007/978-3-540-70583-3_33.

[12] Marek Czarnecki (2010): *How Fast Can the Fixpoints in Modal mu-Calculus Be Reached?* In: *FICS*, Laboratoire d'Informatique Fondamentale de Marseille, pp. 35–39. Available at `https://hal.archives-ouvertes.fr/hal-00512377/document#page=36`.

[13] Javier Esparza, Stefan Kiefer & Michael Luttenberger (2011): *Derivation Tree Analysis for Accelerated Fixed-Point Computation*. *Theor. Comput. Sci.* 412(28), pp. 3226–3241, doi:10.1016/J.TCS.2011.03.020.

[14] Alessandro Facchini, Filip Murlak & Michał Skrzypczak (2013): *Rabin-Mostowski Index Problem: A Step beyond Deterministic Automata*. In: *LICS*, IEEE Computer Society, pp. 499–508, doi:10.1109/LICS.2013.56.

[15] Gaëlle Fontaine (2008): *Continuous Fragment of the mu-Calculus*. In: *CSL*, *Lecture Notes in Computer Science* 5213, Springer, pp. 139–153, doi:10.1007/978-3-540-87531-4_12.

[16] Maria J. Gouveia & Luigi Santocanale (2019): *$\aleph_1$ and the Modal μ-Calculus*. *Log. Methods Comput. Sci.* 15(4), doi:10.23638/LMCS-15(4:1)2019.

[17] Damian Niwiński (1991): *On the Cardinality of Sets of Infinite Trees Recognizable by Finite Automata*. In: *MFCS*, *Lecture Notes in Computer Science* 520, Springer, pp. 367–376, doi:10.1007/3-540-54345-7_80.

[18] Damian Niwiński & Igor Walukiewicz (2003): *A Gap Property of Deterministic Tree Languages*. *Theor. Comput. Sci.* 1(303), pp. 215–231, doi:10.1016/S0304-3975(02)00452-8.

[19] Damian Niwiński, Paweł Parys & Michał Skrzypczak (2025): *A Dichotomy Theorem for Ordinal Ranks in MSO*. In: *STACS*, *LIPIcs* 327, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 69:1–69:18, doi:10.4230/LIPICS.STACS.2025.69.

[20] C.-H. Luke Ong (2006): *On Model-Checking Trees Generated by Higher-Order Recursion Schemes*. In: *LICS*, IEEE Computer Society, pp. 81–90, doi:10.1109/LICS.2006.38.

[21] Michael O. Rabin (1969): *Decidability of Second-Order Theories and Automata on Infinite Trees*. *Transactions of the American Mathematical Society* 141, pp. 1–35.

[22] Sylvain Schmitz (2017): *Algorithmic Complexity of Well-Quasi-Orders. (Complexité algorithmique des beaux pré-ordres)*. Habilitation thesis, Université Paris Diderot (Paris 7). Available at `https://tel.archives-ouvertes.fr/tel-01663266`.

[23] Michał Skrzypczak (2016): *Descriptive Set Theoretic Methods in Automata Theory – Decidability and Topological Complexity*. *Lecture Notes in Computer Science* 9802, Springer, doi:10.1007/978-3-662-52947-8.

[24] Michał Skrzypczak & Igor Walukiewicz (2016): *Deciding the Topological Complexity of Büchi Languages*. In: *ICALP*, *LIPIcs* 55, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 99:1–99:13, doi:10.4230/LIPICS.ICALP.2016.99.