μ-Calculus via Games
(Extended Abstract)

Damian Niwiński*

Institute of Informatics, Warsaw University
Banacha 2, 02-097 Warsaw, Poland
niwinski@mimuw.edu.pl

Introduction

In this survey I would like to present some connections between the μ-calculus and games, more specifically games of possibly infinite duration played on labeled graphs. A fundamental connection was established by Emerson and Jutla [12] and subsequently developed by several authors [13, 36, 42, 1]. Essentially, the result is that any formula of the μ-calculus expresses the existence of a strategy in a certain game. The idea of such a correspondence can be traced back to Büchi and McNaughton who observed a similar property of monadic second order arithmetic (see [10]).

The contribution is mainly conceptual: infinite games (specifically, parity games) can help us to understand the μ-calculus formulas, and eventually design more efficient algorithms for related algorithmic problems as, e.g., model-checking.

For a reader not well familiar with the μ-calculus, infinite games provide a good introductory example since they exhibit fixed points that are not necessarily least fixed points, and the μ-calculus arose from an observation that some properties can be expressed by fixed points, but not the least fixed points. Indeed, the expressive power of the μ-calculus results from the alternation of mutually dependent least and greatest fixed point operators. However the same feature makes the μ-calculus difficult to understand by a human. In fact, in contrast to first-order or temporal logic, the μ-calculus did not emerge by a formalization of the natural language.

Some insight to the μ-calculus has been previously given by reducing it to Rabin automata [37] (which, in fact have the same expressive power [30]). The close connection between automata and games was known already to Büchi, and the alternating automata of Muller and Schupp [27] are already games in disguise. When compared to automata, the combinatorial structure of parity games appears simpler and, in some sense, more symmetric, achieving perhaps the best trade-off between expressive power and simplicity.

For a future development, one may expect that a game-theoretic perspective will bring some new ideas which, in particular, will be helpful for solving the open problem of the complexity of the μ-calculus model-checking. Some steps in this direction have been already made [33, 18, 39, 32]. In a broader perspective, one can even think of a more general theory subsuming the μ-calculus. There seem to be analogies between fixed-point theorems of various kinds, but this direction remains to be investigated.

In this note I first recall the concept of general infinite games, and then the parity games relevant for the μ-calculus, drawing the reader's attention to the fixed point equations for the winning sets. The μ-calculus is subsequently presented in the style of [31, 3], that is as an algebraic system rather than a modal logic, which is slightly more general and allows a particularly simple presentation of the model-checking games, where the pairs of operations $\Diamond, \Box$ and $\vee, \wedge$ are treated in uniform way. Perhaps more importantly, the algebraic perspective reveals an algebraic aspect of the games: the set of winning strategies in a game designed for a term t turns out to be itself an interpretation of this term in a suitable poweralgebra of strategies. In the last section I briefly review the results on several aspects of complexity.

* Supported by Polish KBN grant No. 7 T11C 027 20

Parity games

A general *infinite game* is a game with perfect information of possibly infinite duration played by two players, whom we call *Eva* and *Adam*, to hint a connection to existential and universal quantifiers, respectively.

The game is specified by a set of *positions* Pos partitioned into two disjoint subsets Pos_e and Pos_a of positions of Eva and Adam, respectively, the move relation $Mov \subseteq Pos \times Pos$, a ranking function $rank : Pos \to \omega$, and a set $Win_e \subseteq \omega^\omega$. (Throughout the paper ω denotes the set of natural numbers.) A play can be viewed as moving a token along the edges of the (directed) graph $\langle Pos, Mov \rangle$. A token can be moved from a given position v to any w such that $(v,w) \in Mov$. The actual choice of w is made by Eva or by Adam, depending on whether $v \in Pos_e$ or $v \in Pos_a$. If there is no such w, the player who has to play, looses; in this case we call v a deadlock position.

A result of a play (which we can identify with the play itself) is a sequence $v_0, v_1, \ldots$, such that, for each $i > 0$, $(v_{i-1}, v_i) \in Mov$, and this sequence is either infinite, or ends in a deadlock position. If the play is infinite, the win depends on the sequence of ranks: $rank(v_0)\, rank(v_1)\, rank(v_2) \ldots$. If this sequence belongs to Win_e, Eva wins the game, otherwise Adam is the winner.

A strategy for Eva is a mapping that tells Eva her next move in the play, depending on the current history. That is, a strategy maps $\alpha \in (Pos)^* Pos_e$ to $p \in Pos$ such that $(last(\alpha), p) \in Mov$.

If we fix an initial position v then we need not require a strategy be defined for all histories in $(Pos)^* Pos_e$ but only those that can be reached if Eva actually plays according to the strategy. Thus a strategy from v can be viewed as a tree. We find it convenient to present it as a labeled tree, i.e., a mapping $s : \mathrm{dom}\ s \to Pos$, where $\mathrm{dom}\ s \subseteq \omega^*$ is a set of finite strings over natural numbers, say[1], closed under initial segments. The following properties are required of s.

- $s(\varepsilon) = v$ (since s is a strategy from v).
- If $s(\alpha) \in Pos_a$ then, for each p such that $(s(\alpha), p) \in Mov$, α has a successor α' in dom s labeled $s(\alpha') = p$ (and no other successors).
- If $s(\alpha) \in Pos_e$ then α has exactly one[2] successor α' in dom s such that $(s(\alpha), s(\alpha')) \in Mov$.

Clearly, the labeled paths in the tree s correspond to some possible plays of the game. (Note that any two different paths are labeled in different way.) Observe that by definition no finite play consistent with a path in s is lost by Eva. A strategy s is *winning* for Eva if, additionally, any infinite path $(\varepsilon = \alpha_0, \alpha_1, \alpha_2, \ldots)$ in dom s satisfies the winning condition, that is, $(rank(s(\alpha_0))\, rank(s(\alpha_1))\, rank(s(\alpha_2)) \ldots)$ belongs to Win_e. In other words, a strategy is winning if Eva wins any play played according to the strategy.

We say that v is a *winning position* of Eva if there exists a winning strategy for Eva from this position. (Note that v itself need not be a position of Eva.)

The concepts of strategy and winning position of Adam are defined analogously.

Let V_e and V_a be the sets of winning positions of Eva and Adam respectively. Clearly $V_e \cap V_a = \varnothing$. It would be plausible to think that $V_e \cup V_a = Pos$, but it is not always the case. If it happens, we say that a game is *determinate*. It follows from Axiom of Choice that there exist games that are not determinate [14] (see [28, 15]). But the realm of determinate games is large: the celebrated theorem by D. A. Martin says that any game with a Borel winning set $W_e \subseteq \omega^\omega$ is determinate [23] (see also [28]).

In reality, a strategy need not always depend on the whole history of the game played so far. An important case is if it depends only on the current position. We call a strategy of Eva s *positional* (or memoryless) if, whenever $s(\alpha) = s(\beta) \in Pos_e$, then we have also $s(\alpha') = s(\beta')$ for the successors of α and β, respectively.

We call a game *positionally determinate* if it is determinate, and, moreover, the winner has always a positional strategy.

[1] If the set Pos is uncountable, a larger cardinal should be used.

[2] The uniqueness requirement is not very essential; by relaxing it, we get non-deterministic strategies.

Before we focus on parity games which are most relevant for the μ-calculus, we would like to consider a more general situation that allows to present the winning sets by fixed points. For $(a_n)_n$ and $(b_n)_n$ elements of ω^ω, we let $(a_n)_n \sim (b_n)_n$ if there exist i, j such that the sequences $(a_n)_{n\geq i}$ and $(b_n)_{n\geq j}$ are identical. We call a game *eventuality game* if the winning sets W_e and $W_a = \omega^\omega - W_e$ are saturated by relation $\sim$. (That is, if $(a_n)_n$ is winning for Eva then so are all the sequences equivalent to it; similarly for Adam.) Note that in an eventuality game the win of an infinite play is not altered if we change a finite segment of the play. In particular, if Eva has a winning strategy s from a position v then any position occurring as $s(\alpha)$ is also winning. Also, if at position v Eva has a strategy to reach some winning position w then the position v is also winning. Analogous properties hold for Adam by symmetry.

Now we observe that the set of winning positions of Eva in an eventuality game can be viewed as a fixed point. For, let us adopt a notation from modal logic

$$\Diamond X = \{v : \exists w,\ w \in X \ \wedge\ Mov(v, w)\}$$
$$\Box X = \{v : \forall w,\ Mov(v, w) \Rightarrow w \in X\}$$

and let for brevity $Pos_e = E$ and $Pos_a = A$. Then the set V_e of all positions winning for Eva satisfies the equation

$$X = (E \cap \Diamond X) \cup (A \cap \Box X) \tag{1}$$

Indeed, let $v \in V_e$. If v is a position of Eva then a winning strategy tells Eva a move $(v, w) \in Mov$ to a position which, by eventuality, is again winning. Hence $v \in \Diamond V_e$. A similar argument shows that if v is a position of Adam then $v \in \Box V_e$. Conversely, if $v \in (E \cap \Diamond V_e)$ then Eva has a move (v, w) to a position w from where she already has a winning strategy. By eventuality, Eva has also a winning strategy from v. A similar argument shows that $A \cap \Box V_e \subseteq V_e$.

In other words, the set V_e is a fixed point of an operator

$$Eva(X) = (E \cap \Diamond X) \cup (A \cap \Box X)\,.$$

Note that by the Knaster–Tarski Theorem, this operator has a least fixed point $\mu X.Eva(X)$ and a greatest fixed point $\nu X.Eva(X)$, and hence V_e lies somewhere between them.

The reader should be warned here that a (positional) strategy: *stay in* V_e is no guarantee to win.[3]

On the other hand, we can observe that Eva has always a positional strategy from the positions belonging to the least fixed point of $Eva(X)$. (But note that this set may be empty !) Indeed, we have

$$\mu X.Eva(X) = \bigcup_{\xi<\eta} Eva^{\xi}(\varnothing)$$

for some ordinal η, where $Eva^{\xi+1}(\varnothing) = Eva(Eva^{\xi}(\varnothing))$, and $Eva^{\gamma}(\varnothing) = \bigcup_{\xi<\gamma} Eva^{\xi}(\varnothing)$ for a limit ordinal γ. Now, let the *order* of a position $v \in \mu X.Eva(X)$ (not to be confused with $rank(v)$) be the least ordinal ξ such that $v \in Eva^{\xi}(\varnothing)$. Then it is easy to see that if v is a position of Eva then there exists a position w in $\mu X.Eva(X)$ of a strictly lower order such that $Mov(v, w)$. Similarly, if v is position of Adam then any w such that $Mov(v, w)$ has strictly lower order. Hence a strategy *decrease the order* is well defined and winning for Eva for all positions in $\mu X.Eva(X)$.

[3] It is easy to see that this strategy is winning only if the set of rank labellings of all the winning plays is closed in the standard (Baire) topology on ω^ω. Related questions about maximal nondeterministic strategies are investigated in [5].

Now, in a similar mode, the set V_a of winning positions of Adam is a fixed point of an operator

$$Adam(X) = (A \cap \Diamond X) \cup (E \cap \Box X).$$

It can be further observed that the complement of $Eva(X)$ satisfies

$$\overline{Eva(X)} = Adam(\overline{X})$$

and hence

$$X = Eva(X) \Longrightarrow \overline{X} = Adam(\overline{X}). \tag{2}$$

That is, if X is a fixed point of *Eva* then its complement is a fixed point of *Adam*.

Nevertheless, an eventuality game need not always be determinate. In fact, it is not difficult to adapt the Gale and Stewart diagonal construction [14] (see [15]) to obtain an indeterminate game over the binary tree[4].

But the situation is not always so puzzling. In particular, if $Win_e = \omega^\omega$ (that is, *any* infinite play is won by Eva) then it is easy to see that the winning set of Eva is the *greatest* fixed point of the equation (1) and a positional strategy "stay in V_e " is winning in this case. Indeed, if we have *any* L such that $L = (E \cap \Diamond L) \cup (A \cap \Box L)$ and a position v in L, then Eva can preserve this situation after her move, and also any move of Adam remains in L (or perhaps Adam is stuck). Hence $\nu X.Eva(X) \subseteq V_e$, but since V_e is a fixed point, we get equality.

By (2), the complement $\overline{V_e}$ is the *least* fixed point of $Adam(X)$ and by remark above, Adam has a winning strategy on this set. Hence we get determinacy in this very simple case.

The class of *parity games* is a subclass of general games obtained by assuming that the image $rank(Pos)$ of the function $rank : Pos \to \omega$ is *finite* and Win_e consists of those sequences $(a_0, a_1, \ldots)$ in $(rank(Pos))^\omega$ that the maximal a_i appearing infinitely often is even. That is, Eva wins an infinite play $(v_0, v_1, \ldots)$ if $\limsup_{i\to\infty} rank(v_i)$ is even, otherwise Adam is the winner. Clearly this condition is Borel and hence parity games are always determinate by Davis' theorem. But the parity games turn out to enjoy the positional determinacy: the winner has always a positional strategy which, moreover, need not depend on an initial position of the play.

That is, for any parity game we have $V_e \cup V_a = Pos$, and there exist two functions, say $s_e : V_e \to Pos$ and $s_a : V_a \to Pos$, which induce the winning strategies for Eva and Adam whenever some winning strategies exist.

The property of positional determinacy of parity games was discovered independently by Emerson and Jutla [12], Mostowski [25], and McNaughton [24] (the last with restriction to finite graphs). These authors were partly motivated by searching for a simpler proof of the Rabin Tree Theorem (see [38]). Proof via determinacy of certain games was given previously by Gurevich and Harrington [16] and, independently, by A. A. Muchnik [26]. McNaughton [24] analyzed a general class of winning conditions presented by a family of sets $\mathcal{F} \subseteq \mathcal{P}(rank(Pos))$ (assuming $rank(Pos)$ finite), where Eva wins an infinite play iff the set of ranks occurring infinitely often belongs to $\mathcal{F}$. Clearly a parity condition can be also presented in this way; McNaughton [24] (who presented this condition differently) showed that this is the only winning condition of the analyzed class that always guarantee the positional determinacy.

It might be an interesting direction of research to investigate what class of Borel conditions lead to the positional determinacy and, in particular, if it is somehow possible to extend the concept of parity condition to the case when $rank(Pos)$ is infinite.

What especially concerns us here is that the sets of winning positions in parity games can be presented and (at least in principle) computed as fixed points. Clearly, the parity games are eventuality games and, by considerations above the sets of wining positions of Eva and Adam

[4] In the construction of the winning set W ([15], page 94), instead of requiring that $X_\alpha \notin \{Y_\beta : \beta < \alpha\}$, we require that $X_\alpha \notin \bigcup_{\beta<\alpha}[Y_\beta]_\sim$, and similarly for the Y_α's. There is no problem since the sets in question are countable.

are fixed points of the operators *Eva* and *Adam*, respectively. It turns out that these sets can be characterized by means of mutually dependent fixed point operators μ and ν. (This kind of expressions gives rise to a μ-calculus which we formalize in the next section.)

If we let E_i (resp. A_i) be the set of positions of Eva (resp. Adam) with $rank(v) = i$ then, assuming $rank(Pos) = \{0, 1, 2, \ldots, m\}$, the set of winning positions of Eva is precisely

$$V_e = \theta X_m. \ldots .\nu X_2.\mu X_1.\nu X_0. \bigcup_i (E_i \cap \Diamond X_i) \cup \bigcup_i (A_i \cap \Box X_i) \tag{3}$$

where θ stands for μ or ν depending on whether m is odd or even. The reader may easily verify that this set is a fixed point of the equation (1). The characterization (3) was established by Emerson and Jutla [12] and perhaps this result has inspired their currently common definition of parity games which we use also here. (The formulations in [25, 24] were different.) Another, very transparent proof of this property was given later by Walukiewicz [41], by a refinement of the argument of ordinal orders that we have used above for ω^ω games (which, by the way, are parity games with all ranks equal 0). Note that the equation (3) immediately implies the determinacy of parity games (although not yet the positional determinacy). Indeed, the winning set for Adam is characterized by a symmetric formula which turns out to define the complement of (3).

A proof of the positional determinacy of parity games based on the Boolean μ-calculus was given by Arnold [1] (see also [4]).

The μ-calculus

Let *Sig* be a signature, i.e., a finite set of function symbols, each $f \in Sig$ given with arity $ar(f)$. In view of subsequent applications, it is convenient to extend *Sig* by a disjoint copy of dual symbols; let $Sig_\sim = Sig \cup \{\tilde{f} : f \in Sig\}$, with $ar(\tilde{f}) = ar(f)$. We will write f^* to denote f or $\tilde{f}$. We fix a countable set of variables *Var* and define the set of *fixed-point terms* (with dualities) as the least set of expressions $\mathcal{F}_{Sig_\sim}$ such that

- $Var \subseteq \mathcal{F}_{Sig_\sim}$,
- if $f^* \in Sig_\sim$ with $ar(f^*) = k$ and $t_1, \ldots, t_k \in \mathcal{F}_{Sig_\sim}$ then $f(t_1, \ldots, t_k) \in \mathcal{F}_{Sig_\sim}$,
- if $x \in Var$ and $t \in \mathcal{F}_{Sig_\sim}$ then $\mu x.t \in \mathcal{F}_{Sig_\sim}$ and $\nu x.t \in \mathcal{F}_{Sig_\sim}$.

Fixed-point terms can be interpreted by an abstract interpretation

$$\mathcal{D} = \langle \langle D, \leq_{\mathcal{D}} \rangle, \{f^{*\mathcal{D}} : f^* \in Sig_\sim\} \rangle$$

where $\langle \langle D, \leq_{\mathcal{D}} \rangle$ is a complete lattice and, for $f^* \in Sig$ of arity k, $f^{*\mathcal{D}}$ is a k-ary operation on D monotonic in all variables. Given a valuation of variables $v : Var \to D$, the interpretation proceeds by induction on the structure of term t:

$$\begin{aligned}
[\![x]\!]^v_{\mathcal{D}} &= v(x) \\
[\![f^*(t_1, \ldots, t_k)]\!]^v_{\mathcal{D}} &= f^{*\mathcal{D}}([\![t_1]\!]^v_{\mathcal{D}}, \ldots, [\![t_k]\!]^v_{\mathcal{D}}) \\
[\![\mu x.t]\!]^v_{\mathcal{D}} &= \bigwedge \{d : [\![t]\!]^{v[d/x]}_{\mathcal{D}} \leq_{\mathcal{D}} d\} \\
[\![\nu x.t]\!]^v_{\mathcal{D}} &= \bigwedge \{d : d \leq_{\mathcal{D}} [\![t]\!]^{v[d/x]}_{\mathcal{D}}\}
\end{aligned}$$

where $v[d/x]$ is a valuation that sends x to d and otherwise coincides with v.

The concept of a free variable is defined as usual (see, e.g., [4]). If a term t is *closed* (without free variables), we omit valuation and write simply $[\![t]\!]_{\mathcal{D}}$.

Note that we do not *a priori* require any connection between the interpretation of f and $\tilde{f}$, but in concrete realizations these operations often appear coupled by some duality connections.

A useful class of concrete interpretations arises by applying the powerset functor to model-theoretic structures.

A *semi-algebra* [31] over *Sig* can be presented by $\mathcal{B} = \langle B, \{f^{\mathcal{B}} \mid f \in \mathit{Sig}\}\rangle$, where B is the universe of $\mathcal{B}$ and, for each $f \in \mathit{Sig}$ of arity $ar(f) = k$, $f^{\mathcal{B}}$ is a relation over B of arity $k+1$, that is, $f^{\mathcal{B}} \subseteq B^{k+1}$. We call $f^{\mathcal{B}}$ a k-ary *basic operation* from B^k to B (although it need not be a function), and write $b \doteq f^{\mathcal{B}}(a_1, \ldots, a_k)$ to mean $(a_1, \ldots, a_k, b) \in f^{\mathcal{B}}$.

Any semi-algebra over *Sig* naturally induces an interpretation of the fixed-point terms in $\mathcal{F}_{\mathit{Sig}_\sim}$, which we call the *powerset algebra* of $\mathcal{B}$,

$$\wp\mathcal{B} = \langle\langle \mathcal{P}(B), \subseteq\rangle \{f^{\wp\mathcal{B}} \mid f \in \mathit{Sig}\} \cup \{\tilde{f}^{\wp\mathcal{B}} \mid f \in \mathit{Sig}\}\rangle.$$

Here $\langle \mathcal{P}(B), \subseteq\rangle$ is the complete lattice of all subsets of B ordered by the subset ordering and, for each $f \in \mathit{Sig}$ of arity k, and for $L_1, \ldots, L_k \subseteq B$,

$$\begin{aligned} f^{\wp\mathcal{B}}(L_1, \ldots, L_k) &= \{b : \exists a_1 \in L_1, \ldots, \exists a_k \in L_k, f^{\mathcal{B}}(a_1, \ldots, a_k) \doteq b\}, \\ \tilde{f}^{\wp\mathcal{B}}(L_1, \ldots, L_k) &= \{b : \forall a_1, \ldots, a_k \in B,\ f^{\mathcal{B}}(a_1, \ldots, a_k) \doteq b \Rightarrow a_i \in L_i \text{ for some } i\}. \end{aligned}$$

Note that we have

$$\tilde{f}^{\wp\mathcal{B}}(L_1, \ldots, L_k) = \overline{f^{\wp\mathcal{B}}(\overline{L_1}, \ldots, \overline{L_k})}$$

where $\overline{X} = B - X$.

An important example is that of the Kozen's *modal μ-calculus* ([20], see also a recent survey [8]). The syntax and semantics of this calculus are well known, so we recall them here only very briefly. In the usual setting, the formulas are constructed from variables $x \in \mathit{Var}$, propositions $p \in \mathit{Prop}$, and their duals $\bar{p}$, by means of logical connectives $\wedge$ and $\vee$, the fixed-point operators μx and νx, and the modalities $\langle a\rangle$ and $[a]$, for a ranging over some set of *actions Act*. Semantics is provided by a *Kripke structure*, or a labeled transition system, given by a set S of worlds (or states), and an interpretation of each proposition p by a set $[\![p]\!] \subseteq S$, and each action a by a relation $[\![a]\!] \subseteq S \times S$. In our framework, this structure can be presented as a semi-algebra, say $\mathcal{S}$, over signature $\mathit{Sig} = \mathit{Prop} \cup \mathit{Act} \cup \{\mathit{eq}\}$, where the propositions are considered of arity 0, the actions of arity 1, and *eq* is a special symbol of arity 2 (added in order to handle the logical connectives $\vee$ and $\wedge$). The universe of $\mathcal{S}$ is S. For $p \in \mathit{Prop}$ and $s \in S$, we let $p^{\mathcal{S}} \doteq s$ if and only if $s \in [\![p]\!]$. For $a \in \mathit{Act}$ and $s_1, s_2 \in S$, we let $s_1 \doteq a^{\mathcal{S}}(s_2)$ if and only if $(s_1, s_2) \in [\![a]\!]$. Finally, we let $\mathit{eq}^{\mathcal{S}}(s_1, s_2) \doteq s_3$ if and only if $s_1 = s_2 = s_3$. Then the classical semantics of the modal μ-calculus coincides with the one given by the following translation ρ of the formulas into fixed-point terms.

$$\begin{array}{ll} \rho : x \mapsto x & \\ \quad : p \mapsto p & : \bar{p} \mapsto \tilde{p} \\ \quad : (\varphi \wedge \psi) \mapsto \mathit{eq}(\rho(\varphi), \rho(\psi)) & : (\varphi \vee \psi) \mapsto \tilde{\mathit{eq}}(\rho(\varphi), \rho(\psi)) \\ \quad : \langle a\rangle\varphi \mapsto a(\rho(\varphi)) & : [a]\varphi \mapsto \tilde{a}(\rho(\varphi)) \\ \quad : \mu x.\varphi \mapsto \mu x.\rho(\varphi) & : \nu x.\varphi \mapsto \nu x.\rho(\varphi) \end{array}$$

The converse transformation is also easy to provide (see [4]).

Substantially, the framework of semi-algebras is not much more general than that of the modal μ-calculus. However, it is not more complex either, and we believe that an algebraic perspective may give some insight into the connections between various realizations of the μ-calculus.

Game semantics for μ-calculus

Let $\mathcal{B}$ be a semi-algebra over *Sig* and t a closed fixed-point term in $\mathcal{F}_{\mathit{Sig}_\sim}$. We are going to define a parity game such that winning this game by Eva from a position (somehow encoding an element) $b_{start} \in B$ is equivalent to $b_{start} \in [\![t]\!]_{\mathcal{D}}$.

To this end it will be convenient to make a *proviso* that variables occurring in t are indexed in some special way. Without loss of generality we may assume that each variable is bound only once and thus can be qualified as a μ- or ν-variable. We assume that each variable appears in the form $x_{i,j}$, where i is even for ν-variables and odd for μ-variables. Furthermore we require that the first indices decrease top–down. More precisely, we assume that if a variable $x_{k,\ell}$ appears *free* in the scope of $\theta x_{i,j}.\ldots$ (where θ stands for μ or ν) then $k \geq i$. Note that the inequality must be strict for variables bound by different fixed-point operators. For example, a term $\mu x.\nu y.f(x, y, \mu z.\nu w.f(x, z, w))$ can be presented $\mu x_{11}.\nu x_{01}.f(x_{11}, x_{01}, \mu x_{12}.\nu x_{02}.f(x_{11}, x_{12}, x_{02}))$. (Note that, e.g., x_{11} appears free in the scope of μx_{12} but x_{01} does not.)

Let $Sub(t)$ be the set of all subterms of t.

We call a term of the form $f^*(\tau_1, \ldots, \tau_k)$ *guarded*. It will be further convenient to have an operation which reduces a term $\tau \in Sub(t)$ to its "outermost guarded" form, or, if it is not possible, to one of two special symbols $\bot$ or $\top$. It can be defined in the following way. Consider a graph of subterms of t where, for each subterm $\theta x.\sigma$, we draw an edge from $\theta x.\sigma$ to σ and also from x to $\theta x.\sigma$ if x is a subterm of t. Then, for each $\tau \in Sub(t)$, one of the following possibilities occurs.

- There is a path from τ to a unique guarded term σ; in this case we let $red(\tau) = \sigma$. (The path is trivial if τ itself is guarded.)
- The path originating in τ has a loop involving a unique edge of the form $(x, \theta x.\sigma)$; in this case we let $red(\tau) = \bot$ if $\theta = \mu$ and $red(\tau) = \top$ if $\theta = \nu$.

For example, $red(\mu x.\nu y.f(x, y)) = f(x, y)$ and $red(\mu x.\nu y.\mu z.y) = \top$.

We additionally fix two symbols *head* and *arg*.

We are ready to define our game

$$\mathcal{G}(\mathcal{B}, t) = \langle Pos_e, Pos_a, Mov, rank \rangle .$$

The set of positions $Pos = Pos_e \cup Pos_a$ will be a disjoint union of two sets:

- the set of *head positions* $B \times Sub(t) \times \{head\}$, and
- the set of *argument positions*, that is a subset of $B^* \times Sub(t) \times \{arg\}$ defined by

$$\{\langle \langle a_1, \ldots, a_k \rangle, \tau, arg \rangle : red(\tau) \text{ is } f^*(\ldots) \text{ and } ar(f^*) = k\ \} .$$

(We allow an empty sequence $\langle\rangle$ if $k = 0$.) A head position $\langle b, \tau, head \rangle$ belongs to Eva (i.e., to Pos_e) if $red(\tau)$ is of the form $f(\ldots)$ or $\bot$. The remaining head positions belong to Adam. For the argument positions, the qualification is dual. That is, an argument position $\langle \langle a_1, \ldots, a_k \rangle, \tau, arg \rangle$ belongs to Eva if and only if $red(\tau)$ is of the form $\tilde{f}(\ldots)$.

The *Mov* relation consists of edges defined as follows. Suppose $red(\tau) = f^*(\tau_1, \ldots, \tau_k)$ and $b \doteq f^{\mathcal{B}}(a_1, \ldots, a_k)$. Then

$$(\langle b, \tau, head \rangle, \langle \langle a_1, \ldots, a_k \rangle, \tau, arg \rangle) \in Mov$$
$$(\langle \langle a_1, \ldots, a_k \rangle, \tau, arg \rangle, \langle a_i, \tau_i, head \rangle) \in Mov$$

for $i = 1, \ldots, k$.

To complete the definition of the game, we need to define the function *rank*. For a head position $p = \langle b, x_{i,j}, head \rangle$ we let $rank(p) = i$, and we let $rank(p) = 0$ for all remaining positions.

A fundamental connection between the μ-calculus and parity games is the following.

Theorem 1 *Eva has a winning strategy in the game $\mathcal{G}(\mathcal{B}, t)$ from a position $\langle b_{start}, t, head \rangle$ if and only if $b_{start} \in [\![t]\!]_{\wp\mathcal{B}}$.*

This result was essentially established by Emerson and Jutla [12], although the games considered there characterize a more general problem of the *satisfiability* of the μ-calculus formulas rather than problem precisely corresponding to our computing $[\![t]\!]_{\wp\mathcal{B}}$, which is the celebrated *model checking* problem. Games for model-checking appeared subsequently in [13] in less direct way (via automata); an explicit construction has been given by C. Stirling [36]. The proof in [4] exploits a representation of the problem within the Boolean μ-calculus. It is possible to give a direct proof of Theorem 1, by induction on the structure of t.

We will give in the next section (an idea of) a yet another argument, emphasizing an algebraic structure of the powerset of strategies.

An algebraic view of strategies

Again, let $\mathcal{B}$ be a semi-algebra over Sig. We will now define the concept of pre-strategy which is, roughly speaking, a "hint" for Eva in the game of the previous section, although it does not depend on a particular fixed-point term. We first define a set $X_{\mathcal{B}}$, very similar to the set of positions in the game $\mathcal{G}(\mathcal{B},t)$, but with terms in $Sub(t)$ replaced by symbols in $Sig_\sim \cup \{\top\}$. Let

$$X_{\mathcal{B}} = B \times (Sig_\sim \cup \{\top\}) \times \{head\} \cup \bigcup_{f^* \in Sig_\sim} \{\langle\langle a_1,\ldots,a_k\rangle, f^*, arg\rangle : a_1,\ldots,a_k \in B\}.$$

A tree $s : \mathrm{dom}\ s \to X_{\mathcal{B}}$ is a *pre-strategy* if the following conditions hold.

- $s(\varepsilon) = \langle b, f^*, head\rangle$, for some b, f^*.
- If $s(\alpha) = \langle b, f, head\rangle$ then there exist $a_1,\ldots,a_k \in B$ such that $b \doteq f^{\mathcal{B}}(a_1,\ldots,a_k)$ and α has a unique[5] successor α' in dom s labeled $s(\alpha') = \langle\langle a_1,\ldots,a_k\rangle, f, arg\rangle$.
- If $s(\alpha) = \langle b, \tilde{f}, head\rangle$ then, for all $a_1,\ldots,a_k \in B$ such that $b \doteq f^{\mathcal{B}}(a_1,\ldots,a_k)$, α has a successor α' in dom s labeled $s(\alpha') = \langle\langle a_1,\ldots,a_k\rangle, \tilde{f}, arg\rangle$.
- If $s(\alpha) = \langle\langle a_1,\ldots,a_k\rangle, f, arg\rangle$ then, for any $i = 1,\ldots,k$, α has a successor αi in dom s labeled $s(\alpha i) = \langle a_i, g_i^*, head\rangle$, for some $g_i^* \in Sig_\sim \cup \{\top\}$.
- If $s(\alpha) = \langle\langle a_1,\ldots,a_k\rangle, \tilde{f}, arg\rangle$ then, for some i, α has a (unique) successor αi labeled $s(\alpha i) = \langle a_i, g_i^*, head\rangle$, for some $g_i^* \in Sig_\sim \cup \{\top\}$.
- If $s(\alpha) = \langle b, \top, head\rangle$, for some b, then α is a leaf.

We will write $s = s_b$ to indicate that $s(\varepsilon) = \langle b, f^*, head\rangle$, for some f^* (clearly, there can be many such s_b's).

Let $S_{\mathcal{B}}$ be the set of all pre-strategies. The powerset $\mathcal{P}(S_{\mathcal{B}})$ can be organized into an interpretation of the fixed-point terms in $\mathcal{F}_{Sig_\sim}$,

$$\wp S_{\mathcal{B}} = \langle\langle \mathcal{P}(S_{\mathcal{B}}), \subseteq\rangle, \{f^{\wp S_{\mathcal{B}}} : f \in Sig\} \cup \{\tilde{f}^{\wp S_{\mathcal{B}}} : f \in Sig\}\rangle.$$

Here, for $L_1,\ldots,L_k \subseteq S_{\mathcal{B}}$,

- $f^{\wp S_{\mathcal{B}}}(L_1,\ldots,L_k)$ consists of all pre-strategies s such that $s(\varepsilon) = \langle b, f, head\rangle$, for some b, the unique successor of ε (say, 1) is labeled $\langle\langle a_1,\ldots,a_k\rangle, f, arg\rangle$ (for some $a_1,\ldots,a_k \in B$ such that $b \doteq f^{\mathcal{B}}(a_1,\ldots,a_k)$), and the subtrees rooted in the k successors of 1, that are $11,\ldots,1k$, are some pre-strategies $s_{a_1} \in L_1,\ldots,s_{a_k} \in L_k$, respectively. More formally,

$$\begin{aligned} \mathrm{dom}\ s &= \{\varepsilon, 1\} \cup \bigcup_{i=1,\ldots,k} 1i\ \mathrm{dom}\ s_{a_i} \\ s(\varepsilon) &= \langle b, f, head\rangle \\ s(1) &= \langle\langle a_1,\ldots,a_k\rangle, f, arg\rangle \\ s(1iw) &= s_{a_i}(w) \text{ for } w \in \mathrm{dom}\ s_{a_i},\ i = 1,\ldots,k. \end{aligned}$$

[5] Again, uniqueness is not essential, c.f. footnote 2.

– Similarly, $\tilde{f}^{\wp\mathcal{S}_\mathcal{B}}(L_1,\ldots,L_k)$ consists of all pre-strategies s such that $s(\varepsilon)=\langle b,\tilde{f},\mathit{head}\rangle$, for some b, and, for each $a_1,\ldots,a_k$ satisfying $b \doteq f^{\mathcal{B}}(a_1,\ldots,a_k)$, there is a successor of ε, say ξ, labeled $s(\xi)=\langle\langle a_1,\ldots,a_k\rangle, f, \mathit{arg}\rangle$, and this ξ has a unique successor ξi, for some $i\in\{1,\ldots,k\}$, which is a root of some pre-strategy s_{a_i} in L_i.

Now, let π be a projection which maps a set of pre-strategies $L\subseteq S_\mathcal{B}$ to a subset $\{b:\exists s_b\in L\}$ of B. The interesting feature stemming directly from the definition of $\wp\mathcal{S}_\mathcal{B}$ is that this mapping is a homomorphism with respect to the interpretation of $f^*\in \mathit{Sig}_\sim$, that is,

$$\pi(f^{\wp\mathcal{S}_\mathcal{B}}(L_1,\ldots,L_k)) = f^{\wp\mathcal{B}}(\pi(L_1),\ldots,\pi(L_k))$$
$$\pi(\tilde{f}^{\wp\mathcal{S}_\mathcal{B}}(L_1,\ldots,L_k)) = \tilde{f}^{\wp\mathcal{B}}(\pi(L_1),\ldots,\pi(L_k))$$

Moreover, π clearly preserves arbitrary joints and meets, i.e., for any $\mathcal{X}\subseteq\mathcal{P}(S_\mathcal{B})$,

$$\pi(\bigcup\mathcal{X}) = \bigcup\{\pi(X): X\in\mathcal{X}\}$$
$$\pi(\bigcap\mathcal{X}) = \bigcap\{\pi(X): X\in\mathcal{X}\}.$$

These two conditions imply that π is a homomorphism of the fixed-point interpretations $\wp\mathcal{S}_\mathcal{B}$ and $\wp\mathcal{B}$, i.e., preserves all fixed-point terms (see, e.g., [3]).

Proposition 2 *For any closed term $t\in\mathcal{F}_{\mathit{Sig}_\sim}$, $\pi([\![t]\!]_{\wp\mathcal{S}_\mathcal{B}}) = [\![t]\!]_{\wp\mathcal{B}}$. In particular, $b\in[\![t]\!]_{\wp\mathcal{B}}$ if and only if there exists a pre-strategy s_b in $[\![t]\!]_{\wp\mathcal{S}_\mathcal{B}}$.*

We will infer Theorem 1 from the above proposition if we can show that a pre-strategy s_b in $[\![t]\!]_{\wp\mathcal{S}_\mathcal{B}}$ can be used by Eva as a strategy in the game $\mathcal{G}(\mathcal{B},t)$, from the position $\langle b,t,\mathit{head}\rangle$. This can be achieved as follows.

Let us fix a closed term t. It is easy to see that any strategy s for Eva in the game $\mathcal{G}(\mathcal{B},t)$ can be modified to a pre-strategy s' just by extracting the head function symbol of each current term. More specifically, dom s' = dom s, and s coincides with s' except for that whenever the second component of $s(v)$ is τ with $\mathit{red}(\tau)=f^*$, then the value of the analogous component of $s'(v)$ is f^*, and if $\mathit{red}(\tau)=\top$ then this value is $\top$. (The case of $\mathit{red}(\tau)=\bot$ will not occur since such a position would be losing for Eva, contradicting our definition of a strategy.) In this case we say that the pre-strategy s' *induces* the strategy s. Then Theorem 1 relies on the following more technical fact.

Proposition 3 *$[\![t]\!]_{\wp\mathcal{S}_\mathcal{B}}$ consists of those pre-strategies in $S_\mathcal{B}$ that induce winning strategies in the game $\mathcal{G}(\mathcal{B},t)$.*

Sketch of proof. To use induction on structure of t, we need to extend the concept of our game to terms with free variables. For a term $t=t(z_1,\ldots,z_m)$ and a valuation $v:\mathit{Var}\to\mathcal{P}(B)$, the game $\mathcal{G}(\mathcal{B},t,v)$ is defined similarly as for closed terms, but there are new head positions of the form (b,z_i,head). Such a position is always deadlock and it belongs to Eva iff $b\notin v(z_i)$. (Thus Eva wins iff $b\in v(z_i)$.) The concept of inducing a strategy s by a pre-strategy s' is slightly modified: We now only require that dom $s\subseteq$ dom s' and the coincidence of s and s' holds for all $v\in$ dom s, except possibly if $s(v)=\langle b,z_i,\mathit{head}\rangle$, when we only require $s'(v)=\langle b,\varphi,\mathit{head}\rangle$, for some $\varphi\in\mathit{Sig}_\sim\cup\{\top\}$.

Now the claim which is to be established by induction, is the following.

> For any valuation $\mathit{val}:\mathit{Var}\to\mathcal{P}(S_\mathcal{B})$, $[\![t]\!]^{\mathit{val}}_{\wp\mathcal{S}_\mathcal{B}}$ consists of those pre-strategies in $S_\mathcal{B}$ that induce winning strategies in the game $\mathcal{G}(\mathcal{B},t,\pi\circ\mathit{val})$.

The proof can proceed along standard lines, similarly as, e.g., the proof of the equivalence of fixed-point terms and automata (see, e.g., [4]). Here we will only sketch the induction step for $t=\mu x.\tau(x,z_1,\ldots,z_m)$. Let K be the set of pre-strategies inducing the winning strategies in the game $\mathcal{G}(\mathcal{B},t,\pi\circ\mathit{val})$, let us call them *good* for t and val, for short. By Knaster–Tarski Theorem, it is enough to show two things.

(i) $[\![\tau]\!]_{\wp\mathcal{S}_\mathcal{B}}^{val[K/x]} \subseteq K$.

(ii) If $[\![\tau]\!]_{\wp\mathcal{S}_\mathcal{B}}^{val[L/x]} \subseteq L$ then $K \subseteq L$.

To show (i) it is enough to modify a pre-strategy good for τ and $val[K/x]$ to a pre-strategy good for $\mu x.\tau$ and val.

To show (ii), consider a strategy s good for $\mu x.\tau$ and val. Note that if the induced strategy never calls the variable x (as the value of the second component) then s is also good for τ and $val[L/x]$ and hence, by assumption, $s \in L$. In general the set of occurrences of x need not be empty, however, it is well-founded by the parity condition (since x is the outermost μ-variable in t). Then the proof proceeds by the Principle of Tree Induction. □

Remarks on complexity

Definability in arithmetic. Theorem 1 gives us an upper bound on the macro-scale of complexity, which can be relevant for infinite structures. If we can interpret our semi-algebra within the standard model of arithmetic (expressing the basic operations by first-order formulas) then the sets of integers definable by fixed-point terms remain on the level Δ^1_2, although by a direct translation we would obtain an arbitrary nesting of second-order quantifiers here. Indeed, a Δ^1_2 formula is sufficient to express the existence of a winning strategy in a parity game. However, this fact has been essentially known in the inductive definability theory, as well as the fact that the $\Sigma^1_1 \cup \Pi^1_1$ level can indeed be achieved (see Theorems 3.2 and 3.10 in [17]).

Relation to alternation depth. The *alternating depth* measures the number of alternations of μ and ν in a term. The hierarchy can be conveniently presented by organizing the set $\mathcal{F}_{Sig_\sim}$ in classes Σ_n, Π_n, where

- $\Sigma_0 = \Pi_0$ are terms without fixed-point operators,
- Σ_{n+1} is the closure of Π_n under composition and operator μ,
- Π_{n+1} is the closure of Σ_n under composition and operator ν.

The reader may notice a direct connection between the alternation depth and our special indexing of variables in the construction of games.

It is perhaps worth to note that an idea of the parity condition historically preceded parity games themselves (although it was not called that way). In particular, Wagner [40] showed that the sets of strings $L_n \subseteq \{0, 1, 2, \ldots, n\}^\omega$ satisfying the parity condition are hard for deterministic Rabin automata with n pairs (see [38]), and the author used [29] the sets M_n of binary trees with all branches in L_n as a first example of an interpretation of the μ-calculus with a strict alternation–depth hierarchy.

For the modal μ-calculus, the problem if there exist formulas whose alternation depth cannot be reduced was open for more than a decade.

It would nicely fit in our story if we could say that it was eventually solved by showing the hardness of the game formulas (3) (formalized in the signature of the modal μ-calculus), but it was not quite the case. Indeed, Bradfield [7] solved the problem in 1996 by a sophisticated reduction of some hierarchy in arithmetic. A closely related result was shown independently by Lenzi [21]. But the hardness of game formulas was eventually established by Bradfield [7]. Another proof of this result was given later by Arnold [2], by an ingenious application of the Banach Fixed-Point Theorem.

In any case, the game formulas (3) seem to be a simplest example of formulas with irreducible alternation depth.

Relation to model–checking problem. In our setting, the problem can be stated as follows.

Problem: Model checking.
Given: A finite semi-algebra $\mathcal{B}$, and a closed fixed-point term t.
Compute: $[\![t]\!]_{\wp\mathcal{B}}$.

By Theorem 1 and the positional determinacy of parity games, the above problem is polynomial time reducible to the following.

Problem: Parity game solving.
Given: A finite parity game $\langle Pos_a, Pos_e, Mov, rank\rangle$.
Compute: The partition of Pos to the sets of winning positions $V_e \cup V_a$.

The converse reduction is easy via formulas (3), so the problems are polynomial time equivalent. The game solving problem can be stated more ambitiously.

Problem: Constructive game solving.
Given: A finite parity game $\langle Pos_a, Pos_e, Mov, rank\rangle$.
Compute: (Some) positional winning strategies for both players.

This formulation can be viewed as an advantage gained from the game concept, since a strategy can give us an evidence why an element b is, or is not, in the interpretation $[\![t]\!]_{\wp\mathcal{B}}$.

Given a positional strategy for Eva (say, a "candidate" for a winning strategy), it is easy to check in polynomial time if this strategy is winning [12]. This observation places the decision problem: *is a position v winning for Eva ?* in the class NP. Since by determinacy any position is winning for one of the players, the problem is also in co-NP.

It is worth to note that no polynomial time algorithm is known to verify if a candidate partition $V_e' \cup V_a'$ is indeed a desired partition to the winning sets. Also, there is no evidence that knowing exactly all winning positions can help us to find a winning strategy. (As we have already noted, a simple strategy: *stay in* V_e is usually not winning.)

Deterministic algorithms that have been proposed for the game solving, in some sense carry the problem back to the μ-calculus, as they all boil down to computing some fixed points.

In fact, from an algorithmic point of view, the problem is better presented in a yet another polynomially equivalent form (see, e.g., [4]).

Problem: Evaluation of vectorial Boolean fixed-point terms.
Given:

$$\theta \begin{bmatrix} x_1^{(k)} \\ \vdots \\ x_n^{(k)} \end{bmatrix} \cdots\cdot \nu \begin{bmatrix} x_1^{(2)} \\ \vdots \\ x_n^{(2)} \end{bmatrix} .\mu \begin{bmatrix} x_1^{(1)} \\ \vdots \\ x_n^{(1)} \end{bmatrix} .\nu \begin{bmatrix} x_1^{(0)} \\ \vdots \\ x_n^{(0)} \end{bmatrix} . \begin{bmatrix} f_1 \\ \vdots \\ f_n \end{bmatrix}$$

Compute: The value of this fixed point in $\{0,1\}^n$.

Here $f_1, \ldots, f_n$ are monotonic Boolean functions in $n \cdot k$ variables $x_j^{(i)}$.

We can see that the above problem depends on two parameters: n and k which in the game formulation correspond to the number of positions and the number of different ranks, respectively. A standard method of computing the fixed point by subsequent iterations will evaluate this expression in space $\mathcal{O}(n \cdot k)$ and time $n^k \cdot n^{\mathcal{O}(1)}$ (where the polynomial factor $n^{\mathcal{O}(1)}$ is responsible for the computation of functions f_i). This can be improved to $(n/k+1)^k \cdot n^{\mathcal{O}(1)}$ [11] (which implies a bound of $2^n \cdot n^{\mathcal{O}(1)}$). A significant improvement in the iteration algorithm was made by Browne *et al.* [9], decreasing the exponent by half, which was subsequently still improved by Seidl [35], who achieved time $(n/k+1)^{k/2} \cdot n^{\mathcal{O}(1)}$. Both algorithms however use exponential space.

Jurdziński [19] proposed a new algorithm matching the best previous time bound and using only space $\mathcal{O}(n \cdot k)$. The algorithm relies on the game version of the problem and boils down to a computation of a single least fixed point (but in a larger state space).

Some authors approached the problem by making a connection with other games. Stirling [36] reports an observation by Mark Jerrum on a reduction of parity games to simple stochastic games. Puri [33] found a polynomial time reduction of parity games to mean payoff games. Jurdziński independently discovered a similar reduction and used it to show [18] that the problem *is a position v winning for Eva ?* is not only in NP, but even in UP (unambiguous nondeterministic polynomial time), and hence by symmetry also in UP $\cap$ co-UP.

Vöge and Jurdziński [39] embarked on a discretization of a Hoffman and Karp's *strategy–improvement algorithm* originally proposed for stochastic games (in 1966) and using the arithmetic of real numbers. The resulted discrete algorithm for parity games provides a deep insight to the problem, but the complexity of the algorithm is unknown.

Petersson and Vorobyov [32] used a reduction to discounted mean payoff games to obtain a randomized algorithm working in subexponential time in the case if k grows faster then $\sqrt{n}$.

One may expect that further investigations of the game solving problem will not only lead to better algorithms for the μ-calculus model-checking, but perhaps will also open new paths in the μ-calculus itself.

Acknowledgments. I wish to thank Marcin Jurdziński for enlightening remarks on game solving algorithms. Special thanks go to André Arnold for very helpful remarks on the preliminary version of this note. I also thank Igor Walukiewicz for last minute comments.

References

[1] A. Arnold. A selection property of the boolean μ-calculus and some of its applications. *RAIRO–Theoretical Informatics and Applications*, 31:371–384, 1997.

[2] A. Arnold. The μ-calculus alternation-depth hierarchy is strict on binary trees. *RAIRO-Theoretical Informatics and Applications*, 33:329–339, 1999.

[3] A. Arnold and D. Niwiński. Fixed point characterization of weak monadic logic definable sets of trees. In M. Nivat and A. Podelski, editors, *Tree automata and Languages*, pages 159–188. Elsevier, 1992.

[4] A. Arnold and D. Niwiński. Rudiments of μ-Calculus. Elsevier Science, Studies in Logic and the Foundations of Mathematics, 146, North–Holland, Amsterdam, 2001.

[5] J. Bernet, D. Janin and I.Walukiewicz. Permissive strategies: from parity games to safety games. Theoretical Informatics and Applications (RAIRO), to appear.

[6] J. C. Bradfield. The modal mu-calculus alternation hierarchy is strict. *Theoretical Computer Science*, 195:133–153, 1997.

[7] J. C. Bradfield. Simplifying the modal mu-calculus alternation hierarchy. In M. Morvan, C. Meinel, and D. Krob, editors, *Proc. STACS '98*, pages 39–49. Lect. Notes Comput. Sci. 1373, 1998.

[8] J. Bradfield and C. Stirling. Modal logics and mu-calculi: an introduction. In J. Bergstra, A. Ponse and S. Smolka, editors, *Handbook of Process Algebra*, pages 293-332. Elsevier, North-Holland, 2001.

[9] A. Browne, E. M. Clarke, S. Jha, D. E. Long, and W. Marrero. An improved algorithm for the evaluation of fixpoint expressions. *Theoretical Computer Science*, 178:237–255, 1997.

[10] J. R. Büchi. Using determinacy to eliminate quantifiers. In M. Karpinski, editor, *Fundamentals of Computation Theory*, volume 56, pages 367–378. Lect. Notes Comput. Sci., 1977.

[11] R. Cleaveland, M. Klein and B. Steffen. Faster model checking for the modal mu-calculus. In G.v. Bochmann and D. Probst, editors, *Computer-Aided Verification (CAV '92)*, volume 663, pages 410–422, Lect. Notes Comput. Sci., 1992.

[12] E. A. Emerson and C. S. Jutla. Tree automata, mu-calculus and determinacy. In *Proceedings 32th Annual IEEE Symp. on Foundations of Comput. Sci.*, pages 368–377. IEEE Computer Society Press, 1991.

[13] E. A. Emerson, C. S. Jutla, and A. P. Sistla. On model-checking for fragments of the μ-calculus. In C. Courcoubetis, editor, *Computer Aided Verification*, pages 385–396. Lect. Notes Comput. Sci. 697, 1993.

[14] D. Gale and F. M. Stewart. Infinite games with perfect information. Ann. of Math. Studies 28 (Contribution to the Theory of Games II), pages 245–266, Princeton, 1953.

[15] Y. Gurevich. The logic in computer science column. *Bull. EATCS*, 38, pages 93–100, 1989.

[16] Y. Gurevich and L. Harrington. Trees, automata and games. In *Proc. 14th ACM Symp. on the Theory of Computing*, pages 60–65, 1982.

[17] P. G. Hinman. Recursion–Theoretic Hierarchies. Perspective in Mathematical Logic, Springer–Verlag, 1978.

[18] M. Jurdziński. Deciding the winner in parity games is in UP $\cap$ co-UP. *Information Processing Letters*, 68:119–124, 1998.

[19] M. Jurdziński. Small progress measures for solving parity games. In *Proc. 17th Symp. Theoretical Aspects of Computer Science*, volume 1770, pages 290–301. Lect. Notes Comput. Sci., 2000.

[20] D. Kozen. Results on the propositional μ-calculus. *Theoretical Computer Science*, 27:333–354, 1983.

[21] G. Lenzi. A hierarchy theorem for the mu-calculus. In F. Meyer auf der Heide and B. Monien, editors, *Proc. ICALP '96*, pages 87–109. Lect. Notes Comput. Sci. 1099, 1996.

[22] R. McNaughton. Testing and generating infinite sequences by a finite automaton. *Information and Control*, 9:521–530, 1966.

[23] D. A. Martin. Borel determinacy. Annals of Math. 102, pages 363–371, 1975.

[24] R. McNaughton. Infinite games played on finite graphs. *Annals of Pure and Applied Logic*, 65:149–184, 1993.

[25] A. W. Mostowski. Games with forbidden positions. Technical Report Technical Report 78, Instytut Matematyki, University of Gdansk, 1991.

[26] A. A. Muchnik. Games on infinite trees and automata with dead-ends: a new proof of the decidability of the monadic theory of two successors (in Russian). *Semiotics and Information*, 24:17–40, 1984.

[27] D. E. Muller and P. E. Schupp. Alternating automata on infinite trees. *Theoretical Computer Science*, 54:267–276, 1987.

[28] J. Mycielski. Games with perfect information. In R. J. Aumann and S. Hart, editors, *Handbook of Game Theory with Economic Applications*, volume 1, pages 41–70. North-Holland, 1992.

[29] D. Niwiński. On fixed point clones. In L. Kott, editor, *Proc. 13th ICALP*, pages 464–473. Lect. Notes Comput. Sci. 226, 1986.

[30] D. Niwiński. Fixed points vs. infinite generation. In *Proc. 3rd IEEE Symp. on Logic in Comput. Sci.*, pages 402–409, 1988.

[31] D. Niwiński. Fixed points characterization of infinite behaviour of finite state systems. *Theoretical Computer Science*, 189:1–69, 1997.

[32] V. Petersson and S. Vorobyov. A randomized subexponential algorithm for parity games. *Nordic Journal of Computing*, 8:324–345, 2001.

[33] A. Puri. Theory of Hybrid Systems and Discrete Event Systems. PhD thesis, College of Engineering, University of California, Berkeley, 1995.

[34] M. O. Rabin. Decidability of second-order theories and automata on infinite trees. *Trans. Amer. Soc*, 141:1–35, 1969.

[35] H. Seidl. Fast and simple nested fixpoints. *Information Processing Letters*, 59:303–308, 1996.

[36] C. Stirling. Local model checking games. In *Proc. Concur'95*, volume 962, pages 1-11, Lect. Notes Comput. Sci., 1995.

[37] R. S. Streett and E. A. Emerson. An automaton theoretic decision procedure for the propositional μ-calculus. *Information and Computation*, 81:249–264, 1989.

[38] W. Thomas. Languages, automata, and logic. In G. Rozenberg and A. Salomaa, editors, *Handbook of Formal Languages*, volume 3, pages 389–455. Springer-Verlag, 1997.

[39] J. Vöge and M. Jurdziński. A discrete strategy improvement algorithm for solving parity games. In *Proc. CAV 2000*, volume 1855, pages 202-215, Lect. Notes Comput. Sci., 2000.

[40] K. Wagner. Eine topologische Charakterisierung einiger Klassen regulärer Folgenmengen. *J. Inf. Process. Cybern. EIK*, 13:473–487, 1977.

[41] I. Walukiewicz. Monadic second-order logic on tree-like structures. In C. Puech and R. Reischuk, editors, *Proc. STACS '96*, pages 401–414. Lect. Notes Comput. Sci. 1046, 1996.

[42] I. Walukiewicz. Pushdown processes: Games and model checking. *Information and Computation*, 164(2):234–263, 2001.