Fixpoint alternation and the game quantifier

J. C. Bradfield

BRICS,[1] Department of Computer Science, Aarhus University, Bygning 540, Ny Munkegade, 8000 Århus C, Denmark
and
LFCS,[2] Division of Informatics, University of Edinburgh, Edinburgh, United Kingdom, EH9 3JZ (email: jcb@dcs.ed.ac.uk)

Abstract: Drawing on an analogy with temporal fixpoint logic, we relate the arithmetic fixpoint definable sets to the winning positions of certain games, namely games whose winning conditions lie in the difference hierarchy over Σ^0_2. This both provides a simple characterization of the fixpoint hierarchy, and refines existing results on the power of the game quantifier in descriptive set theory.

Key words: descriptive set theory, fixpoint, game quantifier, induction

1 Introduction

For several decades, games have been an essential tool for the study of logic, both in mathematical logic and more recently in computer science. Perhaps the most developed application in computer science logic is the use of Ehrenfeucht–Fraïssé games for first-order logic, and the refinements such as pebble games which correspond to finite variable fragments. However, games are also useful in temporal logic, and in particular for the modal mu-calculus. The ability to switch one's point of view between logics, automata and games facilitates many results. In particular, the semantics of the modal mu-calculus can be described by means of a *parity game*, that is, a game in which the winning condition concerns the parity of the highest *index* seen infinitely often in the game. This presentation is equivalent to a presentation via alternating Rabin automata, or via tableaux. In modal mu-calculus, a key issue is the alternation of minimal and maximal fixpoints; in automata, this corresponds to the Rabin index, and in normal form parity games it corresponds to the number of indices.

Games, in the form of Gale–Stewart games, also play an important role in descriptive set theory: they provide a tool with which many of the structure theorems of the classical and effective Borel and Lusin hierarchies can be obtained. The *game quantifier* $⅁$ takes

[1] Danish National Research Foundation Centre for Basic Research in Computer Science
[2] Laboratory for Foundations of Computer Science

a game, defined by its set of winning plays, and returns the set of winning positions; the power of this quantifier is the object of our study. Kechris and Moschovakis showed that a Σ^0_1 game has a Π^1_1 set of winning positions; and Robert Solovay showed that the set of winning positions of a Σ^0_2 game is Σ^1_1-inductive. Thomas John studied Σ^0_3 games, and the characterization is complex, involving higher-type recursion and certain levels of the constructible universe.

If one looks at the previous paragraph with fixpoint glasses on, one notices that Π^1_1 is Π^0_1-inductive, that is, on the first level Σ^μ_1 of the fixpoint alternation hierarchy of arithmetic with fixpoints; and that Σ^1_1-inductive is the second level. It is then natural to ask whether this is coincidence, or whether there is, in arithmetic, a nice relationship between fixpoint alternation and some hierarchy of games, mediated by the game quantifier. One might initially speculate that Σ^0_n games have Σ^μ_n winning positions, but alas this cannot be true. However, the world of Rabin automata and modal mu-calculus provides a suggestion: the complexity of the Rabin or parity condition corresponds to the complexity of the winning plays, and since the modal fixpoint alternation *is* correlated nicely to that, the next obvious thing to do is to try to find a notion of complexity in arithmetic that corresponds to the Rabin index, and then hope that the correlation still holds in the rather different world of arithmetic. The result of such an exploration is that fixpoint complexity of winning positions does indeed correspond to a natural fine hierarchy of arithmetic, in a way that matches well with the finite automata world; and although the result is pure descriptive set theory, the games used in its proof are natural analogues of games developed for the automata and temporal logic world. In fact, the games were formulated in order to establish the modal mu-calculus fixpoint alternation hierarchy with the assistance of descriptive set theory; so it is pleasing to go back and use them to obtain results in descriptive set theory.

2 Preliminaries

2.1 Notations and basic definitions

ω is the set of non-negative integers; variables $i, j, \ldots, n$ range over ω. The set of finite sequences of integers is denoted ω^*; finite sequences are identified with integers via standard codings; variables u, v range over ω^*. The set of infinite sequences of integers is ${}^\omega\omega$; variables α, β range over ${}^\omega\omega$. For $\alpha \in {}^\omega\omega$, $\alpha(i)$ is the i'th element of α, and $\alpha(<i)$ is the finite sequence $\langle \alpha(0), \ldots, \alpha(i-1)\rangle$. Concatenation of finite and infinite sequences is written with concatenation of symbols or with $\cdot$, and extended to sets pointwise.

We consider (only) spaces that are the product of copies of ${}^\omega\omega$ and ω; they are given the product topology, where ω carries the discrete topology and ${}^\omega\omega$ itself carries the infinite product topology (in which the basic open sets are the sets $u \cdot {}^\omega\omega$ for every finite sequence u). A *pointset* P is a subset of such a space $\mathfrak{X}$; we write variously $P(i, \alpha)$ or $(i, \alpha) \in P$ if $\mathfrak{X}$ is, for example, $\omega \times {}^\omega\omega$. A pointset is *semi-recursive or* Σ^0_1 iff it is a recursive union of basic opens, in other words given by a semi-recursive set of prefixes: in particular, a pointset $P \subseteq {}^\omega\omega$ is Σ^0_1 iff $P = \bigcup_i N_{\varepsilon(i)}$, where N_k denotes the basic neighbourhood $u_k \cdot {}^\omega\omega$ for some recursive enumeration $k \mapsto u_k$ of the finite sequences w^*, and ε is a recursive function $\omega \to \omega$, otherwise known as a recursive element of ${}^\omega\omega$. A *pointclass* is a set of pointsets; if Λ and Λ' are pointclasses, then $\Lambda \wedge \Lambda'$ is the pointclass $\{ P \cap P' \mid P' \in \Lambda, P' \in \Lambda' \}$ and similarly for $\vee$ and $\neg$; if Λ is a pointclass on $\omega \times \mathfrak{X}$, then $\exists^\omega \Lambda$ is the pointclass $\{ Q \subseteq \mathfrak{X} \mid \exists P \in \Lambda . x \in Q \Leftrightarrow \exists i.(i,x) \in P \}$; similarly for $\exists^{{}^\omega\omega}$. The *Kleene pointclasses* (the arithmetical and analytical hierarchies) are defined by $\Pi^i_j = \neg\Sigma^i_j$; $\Sigma^0_{j+1} = \exists^\omega \Pi^0_j$; $\Sigma^1_0 = \Sigma^0_1$; $\Sigma^1_{j+1} = \exists^{{}^\omega\omega}\Pi^1_j$; $\Delta^i_j = \Sigma^i_j \cap \Pi^i_j$. Pointsets in the Kleene pointclasses are definable by formulae of first- and second-order arithmetic in the usual prenex normal form.

For completeness we recall that the boldface classes are given by: $\boldsymbol{\Sigma}^0_1$ is the class of open sets, and then similarly to the lightface classes; however, we are here concerned mainly with the lightface classes.

Ordinals are ranged over by variables ζ, ξ.

An *(ω-)tree* is a prefix-closed subset of ω^*. If T is a tree, α is an *infinite branch* of T iff $\forall i . \alpha({<}i) \in T$. The *body* $[T]$ of T is the set of its infinite branches. T is *recursive (etc.)* iff it is recursive (etc.) as a subset of ω via the coding of sequences. The following standard fact will be useful:

Lemma 1 *If $P \subseteq {}^\omega\omega$ is Π^0_1, then there is a Π^0_1 tree T such that $\alpha \in P \Leftrightarrow \alpha \in [T]$.*

Proof. If P is Π^0_1, by definition it is $\neg \bigcup_j N_{\varepsilon(j)} = \bigcap_j \overline{N_{\varepsilon(j)}}$ for some recursive ε. Put $T = \{ v \mid \forall j . v \cdot {}^\omega\omega \subseteq \overline{N_{\varepsilon(j)}} \}$. If $\alpha \in P$, then $\forall i . \alpha({<}i) \in T$; conversely, if $\alpha \notin P$, then $\alpha \in N_k$ some k, and then there is a prefix $\alpha({<}i) \in N_k$. Finally, T is Π^0_1 since the test '$v \cdot {}^\omega\omega \subseteq \overline{N_{\varepsilon(j)}}$' is recursive, reducing to 'v does not have $u_{\varepsilon(j)}$ as a prefix', where $u_{\varepsilon(j)}$ is as above.

2.2 Gale–Stewart games

An infinite game of perfect information, or Gale–Stewart game, on ω, is played between two players, Abelard and Eloise. The players take turns, starting with Eloise, to choose a number, so defining a *play* as an infinite sequence $\alpha \in {}^\omega\omega$. The game is defined by a *winning condition* $P \subseteq {}^\omega\omega$, a set of sequences; if $\alpha \in P$ (we write also $P(\alpha)$), then Eloise wins the play, otherwise Abelard. A *strategy* for Eloise is a function from partial plays where she is due to move, i.e. finite sequences of even length, to integers, telling Eloise her next move. A *winning strategy* for Eloise is one such that if she follows it, she is guaranteed to win the game no matter how Abelard plays. If u is a partial play in the game P, then $P[u]$ denotes the game $\{ \alpha \mid u\alpha \in P \}$. A *winning position* for Eloise is a partial play u of even length from which Eloise has a winning strategy for $P[u]$; thus Eloise has a winning strategy for the game, or *wins the game*, iff $\langle\rangle$ is a winning position for her.

It is frequently convenient to relax the definition to allow games with *rules* which constrain the choices of the players, and games where the players' turns need not strictly alternate. This is harmless provided that the rules and the turn function are recursive in the partial plays.

For our purposes, it is also useful to permit finite plays where Eloise wins outright at a particular point. A game with such plays can easily be modified to a game with only infinite plays; an important point for us is that we shall always have recursive winning conditions on finite plays.

A game is *determined* if one or other player wins it. By a celebrated theorem of Martin, all games with $\mathbf{\Delta}^1_1$ winning conditions are determined. However, Wolfe much earlier proved determinacy for $\mathbf{\Sigma}^0_2$ games, and it is a generalization of this proof, far easier than Martin's theorem, that will give us our result.

It is convenient to define *cogames* in which Abelard moves first. Now we can extend the definition of Eloise winning position to partial plays u of odd length by saying that u is a winning position in the game P iff Eloise wins the cogame $P[u]$.

If $P \subseteq {}^{\omega}\omega \times \mathfrak{X}$, then for each $x \in \mathfrak{X}$ the set $\{\, \alpha \mid (\alpha, x) \in P \,\}$ defines a game (we call it $P(\alpha, x)$). The *game quantifier* is defined thus: $\Game\alpha.P(\alpha, x)$ is the set $\{\, x \mid$ Eloise wins the game $P(\alpha, x) \,\}$. Although formally defined in terms of strategies, it is intuitively understood as an infinite string of first-order quantifiers:

$$\exists a_0.\forall a_1.\exists a_2.\forall a_3.\ldots P(a_0 a_1 \ldots, x).$$

Let Γ be a pointclass on ${}^{\omega}\omega \times \mathfrak{X}$; then $\Game\Gamma$ is the pointclass

$$\{\, Q \subseteq \mathfrak{X} \mid \exists P \in \Gamma. Q = \Game\alpha.P(\alpha, x) \,\}.$$

The following standard fact (see [Mos80]) will be required:

Lemma 2 *If Γ is a determined pointclass closed under recursive substitution, then $\neg\Game\Gamma = \Game\neg\Gamma$.*

2.3 Mu-arithmetic

In [Lub93] Robert Lubarsky studies the logic given by adding fixpoint constructors to first-order arithmetic. This logic is also known as LFP in finite model theory, where it is most studied. The logic ('mu-arithmetic' for short) has as basic symbols the following: function symbols f, g, h; predicate symbols P, Q, R; first-order variables x, y, z; set variables X, Y, Z; and the symbols $\vee, \wedge, \exists, \forall, \mu, \nu, \neg, \in$. As with the modal mu-calculus, $\neg$ can be pushed inwards to apply only to atomic formulae, by De Morgan duality.

The language has expressions of three kinds, individual terms, set terms, and formulae. The individual terms comprise the usual terms of first-order logic. The set terms comprise set variables and expressions $\mu(x, X).\, \phi$ and $\nu(x, X).\, \phi$, where X occurs positively in ϕ. Here μ binds both an individual variable and a set variable; henceforth we shall often write just $\mu X.\, \phi$, and assume that the individual variable is the lower-case of the set variable. We also use $\mu\!\!\!\nu$ to mean 'μ or ν as appropriate'. The formulae are built by the usual first-order construction, together with the rule that if τ is an individual term and Ξ is a set term, then $\tau \in \Xi$ is a formula.

This language is interpreted over the structure ω of first-order arithmetic. The semantics of the first-order connectives is as usual; $\tau \in \Xi$ is interpreted naturally; and the set term $\mu X.\, \phi(x, X)$ is interpreted as the least fixpoint of the functional $X \mapsto \{\, m \in \omega \mid \phi(m, X) \,\}$ (where $X \subseteq \omega$).

Mu-arithmetic has a prenex normal form [Lub93,Bra97] of the following shape:

$$\tau_n \in \mu X_n.\tau_{n-1} \in \nu X_{n-1}.\tau_{n-2} \in \mu X_{n-2}.\dots\tau_1 \in \mu\nu X_1.\,\phi$$

where ϕ is first-order—that is, a string of alternating fixpoint quantifiers, and a first-order body. If we refer to a formula in normal form, we shall refer to its components by this notation.

We define levels of the fixpoint alternation hierarchy similarly to the Kleene hierarchies: first-order formulae are Σ^μ_0 and Π^μ_0, as are set variables. The Σ^μ_{n+1} formulae and set terms are formed from the $\Sigma^\mu_n \cup \Pi^\mu_n$ formulae and set terms by closing under (i) the first-order connectives and (ii) the formation of $\mu X.\,\phi$ for $\phi \in \Sigma^\mu_{n+1}$.

A set $X \in \omega$ is Σ^μ_n if $x \in X \Leftrightarrow \tau(x) \in \Xi(x)$ for some Σ^μ_n set term Ξ. Note that this does *not* mean that X is a fixpoint, only that X is definable via a fixpoint.

A Σ^μ_1 set corresponds to a set definable by an inductive definition over an arithmetic predicate; hence by Kleene's theorem, Σ^μ_1 is equal to Π^1_1. The higher levels of the fixpoint hierarchy have been characterized by Lubarsky [Lub93] in terms of large admissible ordinals involving a generalized reflection principle, devised for the purpose, and whose essential content is the iteration of the idea $\Pi^1_1 = \Pi^0_1\text{-}IND = (\Sigma_1 \text{ on } \Delta^1_1)$; however, there has not been a simple characterization in terms of existing notions.

2.4 Rabin conditions and parity games

Consider a (non-deterministic) finite automaton on which Eloise and Abelard play a Gale–Stewart style game by alternately choosing next states. A *Rabin condition* is a winning condition for the game of the following form:

$$\bigvee_{1\le i\le n} (\infty G_i \wedge \neg\infty R_i)$$

where G_i and R_i are subsets of states, and ∞X means that the set X is met infinitely often during the play. n is the *Rabin index*.

These *alternating Rabin automata* are important in temporal logic, as they are one characterization of modal mu-calculus. They are equivalent to alternating *parity* automata: a *parity* condition has the form 'for given sets X_i, $1 \le i \le n$, of states, the greatest j such that X_j occurs infinitely often in the play, is even'. The relationship between parity automata and modal mu-calculus is direct [EmJ91], as the parity condition corresponds to the statement 'the highest fixpoint variable regenerated infinitely often is a maximal fixpoint'.

The Rabin index, or the number of sets in a parity condition, correspond to the fixpoint alternation in modal mu-calculus, and it is this that inspires both our question and its solution. (Niwiński [Niw97] gives a survey of all the concepts mentioned here, as part of a study of fixpoint operators on trees.)

3 Games for mu-arithmetic

The familiar Ehrenfeucht–Fraïssé games for first-order logic are used to distinguish structures; one can also define *model-checking games* or semantic games where the object is to determine whether $\phi(x)$ holds for a formula ϕ and element x of a given structure. For first-order logic, the game may be defined thus: given a formula ϕ and a structure $\mathscr{T}$, a position in the game is a subformula ψ of ϕ and a valuation of the free variables of ψ by elements of $\mathscr{T}$. If ψ is a conjunction, then it is Abelard's turn, and he chooses a conjunct; if a disjunction, then Eloise chooses a disjunct. If $\psi = \exists x.\psi'$, then Eloise chooses a value for x and play moves to ψ'; and dually for $\forall x.\psi'$. Play terminates at atomic formulae $P(\vec{x})$; Eloise wins the play if $P(\vec{x})$ holds of the chosen values for $\vec{x}$, Abelard otherwise. It is a standard theorem that Eloise wins the game iff ϕ is true. (If ϕ itself has free variables, there is one game for each valuation of them.)

Remark 3 Traditionally, the valuation of variables is not explicitly encoded in the position, but read off from the history. This is a vital distinction in the finite model theory use of games, since we do not wish to assume that we can keep arbitrary amounts of information. The notion of pebble game was invented in order to encode the variable assignment in the position, and then the number of pebbles (variables) can be limited, so that we can ask 'is there a winning strategy using only the bounded history information available?'. However, we are working in arithmetic, so we have all the coding apparatus we want, and may as well carry the assignment with us; it is purely a matter of convenience.

Games have been extended to LFP in the world of finite model theory, by choosing a candidate fixpoint set X when one passes through a fixpoint operator. Uwe Bosse [Bos92] has used such games to obtain expressivity results on fragments of LFP. However, such a game is undesirable in arithmetic, since it has second-order positions. A more useful game for mu-arithmetic is defined by adapting the game for parity automata or modal mu-calculus: instead of finite plays, one now has infinite plays, and the winning condition is given by a parity condition.

Given a formula of mu-arithmetic (or indeed FOL with fixpoints in general), the model-checking game has moves as for first-order logic together with the following rules for the fixpoints: if the position is $\tau \in \mu\!\nu(x, X).\,\phi$, then play moves to the position ϕ with x valued at the current value of τ. It does not matter who moves, but for definiteness say Eloise moves for μ and Abelard for ν. If the position is $\tau \in X$, where X is bound by $\mu\!\nu(x, X).\,\phi$, then again play moves to ϕ with x valued at the current value of τ, and we say that we have seen X. It remains to define the winning conditions: if play terminates, the play is won as for first-order logic; if the play is infinite, then Eloise wins iff the outermost fixpoint variable seen infinitely often is maximal. If we start with a formula in normal form, then this is exactly a parity condition. We have

Theorem 4 *A formula $\phi(\vec{x})$ of mu-arithmetic holds for some valuation of $\vec{x}$ exactly if Eloise has a winning strategy for the model-checking game for ϕ with the given initial valuation.*

Proof. A full proof of this theorem is quite long; however it is strategically the same as the corresponding proof for modal mu-calculus and parity games. It is also the essential content of Theorem 5 of [Bra97]. Therefore I shall not give the proof again here. □

4 The power of the game quantifier

Suppose that a winning condition $P(\alpha, x)$ for a Gale–Stewart game has a given descriptive complexity: what is the descriptive complexity of $\Game\alpha.P(\alpha, x)$? If the winning conditions are in the analytical hierarchy, then this is a question intimately related to the structure theory of the hierarchy, and a question that depends on hypotheses outside *ZFC*, in particular, the hypothesis of Projective Determinacy (that all projective games are determined). Given *PD*, the answer is quite simple for analytical games: $\Game\Sigma^1_n = \Pi^1_{n+1}$ and dually, so $\Game$ is a 'hermaphrodite' second-order quantifier.

If the winning conditions are below $\mathbf{\Delta}^1_1$, then determinacy is not an issue, and one can expect unequivocal answers. However, the game quantifier turns out to be quite delicate. The first answer was:

Theorem 5 (Kechris–Moschovakis) $\Game\Sigma^0_1 = \Pi^1_1$

Later, Robert Solovay (unpublished, cited in [Mos80], *q.v.* also for the previous theorem) characterized Σ^0_2 games, based on Wolfe's proof of the determinacy of $\mathbf{\Sigma}^0_2$ games.

Theorem 6 (Solovay) $\Game\Sigma^0_2 = \Sigma^1_1$*-IND (that is, sets given via an inductive definition over a* Σ^1_1 *predicate).*

A decade or so later, the next step was taken by Thomas John, who studied Σ^0_3 games. Unfortunately, the characterization is complex: it involves capturing the levels of Gödel's L at which winning strategies can be found, and is given in terms of higher-type recursion. This appears to be inevitable: the proof of determinacy for $\mathbf{\Sigma}^0_3$ proceeds via games in which the positions are themselves games.

There are no published results on Σ^0_4 or beyond.

As was remarked in the introduction, Π^1_1 is also Π^0_1-*IND*; or in terms of the fixpoint hierarchy, Σ^μ_1. Then Σ^1_1-*IND* is just Σ^μ_2. We then naturally ask whether $\Sigma^\mu_n = \Game X$ for some natural class X. The conjecture that immediately comes to mind is $\Sigma^\mu_n = \Game\Sigma^0_n$; unfortunately, the complexities of higher-type recursion are not so easily banished. In fact, the game semantics of mu-arithmetic shows that

Theorem 7 *For all n,* $\Sigma^\mu_n \subseteq \Game\nabla^0_2$*, where* ∇^0_2 *denotes the boolean closure of* $\Sigma^0_2 \cup \Pi^0_2$.

Proof. Consider a mu-arithmetic formula in normal form; wlog, consider the case of odd n. The winning condition for the associated game comprises a recursive part dealing with the finite plays, and a parity condition dealing with the infinite plays. The parity condition says that in a play α, the highest X_i seen infinitely often is maximal, i.e. i is even. In other words, the condition is

$$\neg\infty X_n \wedge (\infty X_{n-1} \vee (\neg\infty X_{n-2} \wedge \infty X_{n-3} \vee (\ldots \infty X_2 \vee \neg\infty X_1)\ldots)).$$

Now, the statement 'X_i is seen at position j of the play α' is a recursive predicate of α; and the statement ∞X_i is just $\forall j.\exists k > j.$'X_i is seen at k', which is Π^0_2. Therefore the entire condition is a boolean combination of Π^0_2 and Σ^0_2 statements, Q.E.D. $\square$

One may equally well use a Rabin condition, although this is less natural.

At this point, it seems 'obvious' that the argument should also run backwards. However, model-checking games are a very restricted format of games, and the statement ∞X_i is apparently a rather restricted form of Π^0_2 statement about a play α; we wish to make a statement about arbitrary ∇^0_2 winning conditions. Thus the obvious statement requires some work to prove. The first step is to choose the appropriate fine hierarchy within ∇^0_2. One may here choose to follow the pattern of Rabin conditions: by using disjunctive normal form, it is trivial that formulae of the shape $\bigvee_i(\Sigma^0_2 \wedge \Pi^0_2)$ give a normal form for ∇^0_2. However, it is easier and more elegant to follow the pattern of parity conditions, and use a hierarchy known as the *difference hierarchy* (over Σ^0_2). Difference hierarchies over open sets have been studied long ago in classical descriptive set theory; more recently Victor Selivanov has, in a series of papers, made a study of an abstract fine hierarchy which subsumes, in a certain sense, difference hierarchies: applications include simpler proofs and refinements of Wagner's hierarchies of ω-regular languages [Sel98]. However, we shall not need any of this more general theory; let us just define the hierarchy we need.

The difference hierarchy over Σ^0_2 is defined thus: $\Sigma^\partial_1 = \Sigma^0_2$; $\Pi^\partial_n = \neg\Sigma^\partial_n$; $\Sigma^\partial_{n+1} = \Sigma^0_2 \wedge \Pi^\partial_n$. To provide a simpler base case, let us also define $\Sigma^\partial_0 = \Sigma^0_1$ (which fits into the induction, since $\Sigma^0_2 \wedge \Pi^0_1 = \Sigma^0_2$). The main result is now

Theorem 8 *$\Game\Sigma^\partial_n = \Sigma^\mu_{n+1}$ for $n \geq 0$.*

Proof. First consider the easier direction, that $\Sigma^\mu_{n+1} \subseteq \Game\Sigma^\partial_n$. This is not trivial: by inspection, the parity condition of rank n is in Σ^∂_n, but this is not tight enough. However, if we consider more carefully the winning condition for the game of a Σ^μ_1 formula $\tau \in \mu X.\phi$, it says simply 'X is seen finitely often'. Since the only way a play can be infinite is to pass infinitely often through X, this is equivalent to saying that the play is really finite (and therefore terminates on an outright Eloise win). Hence the winning condition is really just $\exists i.$'Eloise wins outright at $\alpha(i)$', and since the outright winning conditions are recursive, this is a Σ^0_1 statement. Hence $\Sigma^\mu_1 \subseteq \Game\Sigma^0_1 = \Game\Sigma^\partial_0$. Now an induction following the proof of Theorem 4 gives the rest. (Of course, we already know from Theorem 5 that $\Sigma^\mu_1 = \Game\Sigma^0_1$; however, the above direct argument of the base case has the advantage of being easy to fit directly into the induction.)

The harder direction is showing that $\Game\Sigma^\partial_n \subseteq \Sigma^\mu_{n+1}$. For convenience we shall let Theorem 5 deal with the base case; it is an easy exercise to write down the direct proof using a simplified version of the strategy here. The inductive step is a generalization of Solovay's result, using a generalization of Wolfe's determinacy proof. We shall follow, more or less, the presentation of Wolfe's proof by Moschovakis [Mos80], extending as necessary.

The approach is to define inductively 'easy winning positions', and show that all winning positions are easy. We then inspect the inductive definition, and see that it has the required form.

Suppose we have a Σ^∂_n winning condition $P(\alpha, \vec{x})$; for notational convenience we omit the parameters $\vec{x}$. Then it has the form

$$(\exists i. Q(i, \alpha)) \wedge R(\alpha)$$

where Q is Π^0_1 and R is Π^∂_{n-1}. In Solovay's result, the winning condition is $\Sigma^0_2 = \Sigma^\partial_1$, and so there is no R term; we have to show that the argument still goes through with this additional term, so allowing us to use the proof in an induction on n.

We start with a trivial but critical observation:

$$(\exists i. Q(i, \alpha)) \wedge R(\alpha) \Leftrightarrow \exists i.(Q(i, \alpha) \wedge R(\alpha)).$$

The second observation is that (by lemma 1) since, for a given i, $Q(i, \alpha)$ is a Π^0_1 predicate of α, there is a Π^0_1 tree $T_i \subseteq \omega^*$ such that $Q(i, \alpha) \Leftrightarrow \alpha \in [T_i]$.

We shall build the set of winning positions by a transfinite induction; to explain the technique let us first consider the base case on its own. We can define a set of *really* easy winning positions: let

$$W^0 = \{\, u \mid \exists i. \text{'Eloise wins the game } H^0_i[u] = (Q(i, \alpha) \wedge R(\alpha))[u]\text{'}\,\}$$

Strictly, if u is an odd length sequence, we mean the cogame $H^0_i[u]$; we will assume henceforth that 'game' means 'game' for even length u and 'cogame' for odd length u. Now, it is clear that if $u \in W^0$, then Eloise wins the game $P[u]$.

To extend this base case into an inductive step, we first reformulate this definition using the second observation:

$$W^0 = \{\, u \mid \exists i. \text{'Eloise wins the game } H^0_i[u] = (R(\alpha) \wedge \forall k. \alpha(<k) \in T_i)[u]\text{'}\,\}$$

So the 'really easy' winning positions can be thought of as the places where Eloise knows how to win R while also staying within T_i. Now the inductive step is to look at places where Eloise knows how to win R while staying within the winning positions for easier games. That is, if W^ξ is defined for $\xi < \zeta$, let $W^{<\zeta} = \bigcup_{\xi<\zeta} W^\xi$, and define the game

$$H^\zeta_i(\alpha) = R(\alpha) \wedge \forall k. \alpha(<k) \in W^{<\zeta} \cup T_i.$$

Then we define

$$W^\zeta = \{\, u \mid \exists i. \text{'Eloise wins the game } H^\zeta_i[u]\text{'}\,\}.$$

We show by induction that if $u \in W^\zeta$, then u is a winning position in the original game P. So, let $u \in W^\zeta$. Then for some i, Eloise wins $H^\zeta_i[u]$; let Eloise play according to her winning strategy to produce a play α. Then by definition, $R(\alpha) \wedge \forall k. \alpha(<k) \in W^{<\zeta} \cup T_i$. If play ever reached a position $v = \alpha(<k) \in W^{<\zeta}$, then by induction Eloise could have won $P[v]$, so by switching to her winning strategy there, instead of continuing with α, she can win $P[u]$. If not, then $R(\alpha) \wedge \forall k. \alpha(<k) \in T_i$; but then $\alpha \in [T_i]$, so $Q(i, \alpha) \wedge R(\alpha)$, so Eloise wins the play α in the game P.

Now, W^ζ is an increasing chain, and so (by cardinality) closes at some $W = W^\kappa = W^{<\kappa}$. We now show that if $u \notin W^\kappa$, then Abelard wins $P[u]$. So, let $u = a_0 \dots a_j \notin W^\kappa$. By definition, for all i, Abelard wins $H^\kappa_i[u]$. (Note: see Remark 9.) Let Abelard continue to play according to his winning strategy for $H^\kappa_0[u]$, generating a play $\alpha = ua_{j+1} \dots$. First suppose that $\forall k.\alpha(<k) \in W^{<\kappa} \cup T_o$; then we must have $\alpha \notin R$, so $\alpha \notin P$, and so Abelard has won P. On the other hand, suppose at some j_0 we have $u_0 = \alpha(<j_0) \notin W^{<\kappa} \cup T_0$. Then firstly $u_0 \notin T_0$, and since T_0 is a tree, any extension of u_0 is also $\notin T_0$. Secondly, $W^{<\kappa} = W^{\kappa}$, so $u_0 \notin W^{\kappa}$, so Abelard wins all $H_i^{\kappa}[u_0]$. So now let Abelard switch to his strategy for H_1^{κ}. Now repeat the argument: either Abelard plays and wins with R, or there is a $u_1 \notin W^{<\kappa} \cup T_1$. If the process of finding $u_0, u_1, \ldots$ continues for ever, then the final play α is not an infinite branch of any T_i, and so $\neg\exists i.Q(i, \alpha)$ and again Abelard has won the play.

We have now shown that $u \in W$ iff Eloise wins the game $P[u]$, in other words that $W = \Game\alpha.P(u \cdot \alpha)$. All that remains is to recast the inductive definition in terms of $\Game$ and mu-arithmetic:

$$W = \mu(w, W).\exists i.\Game\alpha.(R(w \cdot \alpha) \wedge \forall k.((w \cdot \alpha)(<k) \in W \vee (w \cdot \alpha)(<k) \in T_i))$$

Now, T_i is a Π_1^0 set, and therefore $\forall k. \ldots$ is also Π_1^0; R is Π_{n-1}^{∂}, and so the body of the game quantified expression is also Π_{n-1}^{∂}. Therefore by the induction hypothesis and duality, $\Game\alpha. \ldots$ is equal to some Π_n^{μ} expression ϕ, and so W is indeed Σ_{n+1}^{μ}; Q.E.D. $\square$

Remark 9 At the point referring to this Remark, we seem to be assuming the determinacy of H_i^{κ}. At first sight, this seems odd, since the set W^{κ} occurring in the definition is rather complex; however, the determinacy theorems involve the boldface classes, not the lightface, and any subset of the integers is $\boldsymbol{\Delta}_1^0$, so the determinacy theorems apply. In fact, as I have mentioned, this proof is mostly Wolfe's, and was devised to show determinacy. This works because at the end we have constructed a set W such that Eloise wins P iff $\langle\rangle \in W$, and Abelard wins P iff $\langle\rangle \notin W$. To show determinacy, we use $\boldsymbol{\Sigma}_2^0$ rather than Σ_2^0; the only difference this makes is that Q is $\boldsymbol{\Pi}_1^0$ instead of Π_1^0, so the trees T_i are not necessarily Π_1^0. The argument goes through to produce the set W; thus we have an inductive proof of the determinacy of $\boldsymbol{\Sigma}_n^{\partial}$ games, and then we look at the lightface version in order to obtain the complexity results we really want.

It is worth mentioning, as already noted in [EmJ91], that the use of fixpoint notation makes Wolfe's proof itself rather more transparent. The determinacy of (boldface) $\boldsymbol{\nabla}_2^0$ games was studied by Büchi [Büc77] in the context of monadic second-order logic; again the use of fixpoint notation allows a transparent presentation, since the formulation of Theorem 8 works also in the boldface case.

As was mentioned above, we could as well use the Rabin style hierarchy as the difference hierarchy; indeed, Σ_{2n}^{∂} is equal to the Rabin class $\bigvee_{1 \le i \le n}(\Sigma_2^0 \wedge \Pi_2^0)$, and the odd levels of the difference hierarchy correspond to Rabin conditions with one disjunct being simply Σ_2^0.

Since the above proof is also defining a winning strategy, it follows from Lubarsky's characterization that

Corollary 10 *A Σ_n^{∂} game has a winning strategy in the $\omega^{+(n+1)}$'th level of Gödel's L, where ω^{+n} is the first n-reflecting admissible [Lub93] after ω.*

A result that is already known, but which is much more easily seen from this approach, is

Corollary 11 *The fixpoint definable sets of integers are strictly contained in* Δ^1_2.

Proof. Because the game quantifier is self-dual for reasonable point classes (lemma 2) and because if $U(i, \alpha, x) \subseteq \omega \times {}^{\omega}\omega \times \mathfrak{X}$ is universal for Γ on ${}^{\omega}\omega \times \mathfrak{X}$ then $\Game\alpha.U(i, \alpha, x) \subseteq \omega \times \mathfrak{X}$ is universal for $\Game\Gamma$ on $\mathfrak{X}$, the game quantifier preserves the strictness of reasonable hierarchies. In particular, it preserves the arithmetic and hyperarithmetic hierarchies. By the main theorem, the fixpoint definable sets are contained in $\Game\nabla^0_2$; but Δ^1_2 contains $\Game\Delta^1_1$, a much larger set.

The fact that fixpoint definable sets are contained in Δ^1_2 follows from the classical closure of Δ^1_2 under inductive definitions; the strictness of the containment is established by Lubarsky's analysis, as the ordinals ω^{+n} are all less (and in some sense much less) than the first non-Δ^1_2 ordinal. However, the game characterization is technically simpler, and gives a more transparent meaning to 'much less': ∇^0_2 is 'much smaller' than Δ^1_1 in a well understood sense.

5 Final remarks

The characterization we have established here is interesting in both directions. The fact that $\Game\Sigma^\partial_n$ is characterized by a natural and useful pointclass extends a little the point at which games become inherently difficult. Perhaps the other direction is more interesting: fixpoint alternation is notoriously incomprehensible, so characterizing it in terms of a simple hierarchy of games is helpful—and arguably more useful than the admissible-recursion-theoretic characterization. It also reinforces a slightly different view on the traditional world of automata with Rabin and parity conditions: 'infinitely often' is a fundamental concept in temporal logics, but really it is a fundamental concept because it is Π^0_2. Indeed, within the framework of recursion theory, any Π^0_2 statement about α is of the form 'infinitely often something happens at $\alpha(i)$', where in general 'something' includes statements about the previous and future elements of α.

There are some obvious directions in which to extend this investigation. As I have mentioned, Selivanov has made a detailed study of difference-like hierarchies, and I hope that further results may emerge from applying his work. One question of particular interest is that of transfinite extensions. The fixpoint hierarchies can be extended into the transfinite in the usual way, and Selivanov's hierarchies are also transfinite. One can then ask whether our characterization here extends at all. A possibly interesting issue here is that there is something of a mis-match between the hierarchies: the first natural stopping point for a transfinite extension of the difference hierarchy is ω_1^{CK}, whereas the transfinite fixpoint hierarchy has no natural stopping point before an otherwise unknown and extremely large ordinal. I conjecture that an extension may be possible up to ω_1^{CK}; beyond there it is not clear that the question is meaningful.

Another natural question is whether this descriptive set theoretic approach can be used directly to establish results such as the non-collapse of the modal mu-calculus alternation hierarchy, avoiding the passage through arithmetic that was used in [Bra97]. Here again Selivanov's work is relevant: he has shown how an approach from descriptive set theory provides elegant proofs of results on ω-regular expressions.

6 Acknowledgements

I am supported by an Advanced Fellowship (AF/97/0322) from the United Kingdom Engineering and Physical Sciences Research Council; also BRICS, the Danish National Research Foundation Centre for Basic Research In Computer Science, is supporting my visit to Aarhus.

References

[Bos92] U. Bosse, An 'Ehrenfeucht-Fraïssé game' for fixpoint logic and stratified fixpoint logic. *Computer science logic* , LNCS **702** (San Miniato, 1992) 100–114.

[Bra97] J. C. Bradfield, The modal mu-calculus alternation hierarchy is strict, *Theor. Comput. Sci.* **195** 133–153 (1997).

[Büc77] J. R. Büchi, Using determinancy of games to eliminate quantifers, *Proc. FCT '77*, LNCS **56** 367–378 (1977).

[EmJ91] E. A. Emerson and C. S. Jutla, Tree automata, mu-calculus and determinacy, in: *Proc. FOCS 91.* (1991)

[Lub93] R. S. Lubarsky, μ-definable sets of integers, *J. Symbolic Logic* **58** (1993) 291–313.

[Mos80] Y. N. Moschovakis, *Descriptive Set Theory* (North-Holland, Amsterdam, 1980).

[Niw97] D. Niwiński, Fixed point characterization of infinite behavior of finite state systems. *Theoret. Comput. Sci.* **189** (1997) 1–69.

[Sel98] V. Selivanov, Fine hierarchy of regular ω-languages, *Theoret. Comput. Sci.* **191** (1998) 37–59.