Transfinite Extension of the Mu-Calculus

Julian Bradfield[1], Jacques Duparc[2], Sandra Quickert[1]

[1] Laboratory for Foundations of Computer Science, University of Edinburgh, {jcb,squicke1}@inf.ed.ac.uk

[2] Ecole des HEC, Université de Lausanne, jacques.duparc@unil.ch

Abstract. In [1] Bradfield found a link between finite differences formed by Σ_2^0 sets and the mu-arithmetic introduced by Lubarski [7]. We extend this approach into the transfinite: in allowing countable disjunctions we show that this kind of extended mu-calculus matches neatly to the transfinite difference hierarchy of Σ_2^0 sets. The difference hierarchy is intimately related to parity games. When passing to infinitely many priorities, it might not longer be true that there is a positional winning strategy. However, if such games are derived from the difference hierarchy, this property still holds true.

1 Introduction

Modal mu-calculus, the logic obtained by adding least and greatest fixpoint operators to modal logic, has long been of great practical and theoretical interest in systems verification. The problem of understanding alternating least and greatest fixpoints gave rise to a powerful and elegant theory relating them to alternating parity automata and to parity games, developed by many people including particularly Emerson, Lei, Jutla and Streett. Meanwhile, mu-arithmetic, the logic obtained by adding fixpoints to first-order arithmetic, made a brief appearance in the early 90s when Lubarsky studied its ordinal-defining capabilities – curiously, the logic had not previously been studied *per se* even by logicians. Then Bradfield used mu-arithmetic as a meta-language for modal mu-calculus, in which to prove a theorem on alternating fixpoints. Subsequently, Bradfield looked further into the analogies between mu-arithmetic and modal mu-calculus, and showed a natural equation between arithmetic fixpoints and the finite difference hierarchy over Σ_2^0, corresponding to the equation between modal fixpoints and parity games. Once in the world of arithmetic, it becomes natural to think about transfinite hierarchies. In this paper, we study the transfinite extension of the connection between mu-arithmetic and the difference hierarchy, and connect it to the Wadge hierarchy.

2 The Transfinite Mu-Calculus

2.1 Syntax and Semantics of the Transfinite Mu-Calculus

The logic we are considering is an extension of the usual mu-arithmetic, as introduced by Lubarski [7]. First, let us establish basic notation and conventions.

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; the length of a sequence s is denoted $\mathrm{lh}(s)$. The set of infinite sequences of integers is ${}^\omega\omega$. 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 $\frown$, and extended to sets pointwise. The usual Kleene lightface hierarchy is defined on ω, ${}^\omega\omega$ and their products: $\Sigma^0_1 = \Sigma^1_0$ is the semi-recursive sets, $\Sigma^0_{n_1} = \exists x \in \omega.\Pi^0_n$, $\Pi^i_n = \neg\Sigma^i_n$ and $\Sigma^1_{n+1} = \exists \alpha \in {}^\omega\omega.\Pi^1_n$. The corresponding boldface hierarchy is similar, but starts with $\boldsymbol{\Sigma}^0_1 = \boldsymbol{\Sigma}^1_0$ being the open sets.

Mu-arithmetic 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 $\vee, \wedge, \exists, \forall, \mu, \nu, \neg, \in$. 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.

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 $X \subseteq \omega$).

To produce a transfinite extension, we add the following symbols and formulae. If we have countably many recursively given Φ_i, $i \in \omega$, whose free set variables are contained in the same finite set of set variables, then we allow infinite countable disjunction $\bigvee_{i<\omega} \Phi_i$ and conjunction $\bigwedge_{i<\omega} \Phi_i$. The restriction on free variables means that we can transform any formula to a closed formula by adding finitely many fixpoint operators. The semantics is obvious.

Any formula in the mu-calculus can be rewritten in a prenex normal form:

$$\tau_n \in \mu X_n.\tau_{n-1} \in \nu X_{n-1}.\tau_{n-2} \in \mu X_{n-2} \ldots \tau_1 \in \mu\!\!\!\nu X_1.\Phi$$

For the transfinite mu-arithmetic we need an extension of this formulation.

Definition 1. By induction on the construction of the formula we say that a formula in the transfinite mu-calculus is written in extended prenex normal form

- if it is a formula in the finite mu-calculus and written in prenex normal form, or
- if the formula is an infinite disjunction or conjunction of extended prenex normal form formulae, or
- if it is some $\mu\!\!\!\nu X.\Phi$ where Φ is in extended prenex normal form.

Given formulae Φ_i for $i < \omega$ in the mu-arithmetic, we observe that the formula $\bigvee_{i<\omega} \Phi_i$ can be written in extended prenex normal form, simply by writing each

Φ_i in prenex normal form. Given an arbitrary formula of the extended arithmetic mu-calculus, an easy proof on induction by the formula's construction shows that it can be written in extended prenex normal form. Furthermore, we can unfold its complexity and represent it by a wellfounded tree on ω^*.

2.2 A Hierarchy of the Transfinite Mu-Calculus

The fixpoint alternation hierarchy of of mu-arithmetic is thus: the first order formulae and all set variables form the class Σ_0^μ which is the same as Π_0^μ. For any natural number n let Σ_{n+1}^μ be generated from $\Sigma_n^\mu \cup \Pi_n^\mu$ by closing it under $\vee, \wedge$ and the operation $\mu X.\Phi$ for $\Phi \in \Sigma_{n+1}^\mu$. Π_{n+1}^μ contains all negations of formulae and set terms in Σ_{n+1}^μ. In order to extend the hierarchy we need to describe the limit step. We allow recursively countable disjunctions and conjunctions, but we want to stay in the lightface hierarchy. Therefore we extend the hierarchy to ω_1^{ck}, the first non-recursive ordinal. Let λ be a recursive limit ordinal. In Σ_λ^μ we collect all formulae of earlier stages and close it under $\bigvee_{i<\omega}, \vee$ and $\wedge$. Observe that a formula in Σ_λ^μ is equivalent to a formula $\bigvee_{i<\omega} \Phi_i$ where each $\Phi_i \in \Sigma_{\alpha_i}^\mu$ with $\alpha_i < \lambda$. Finally, we let $\Pi_\lambda^\mu = \neg\Sigma_\lambda^\mu$. The transfinite successor stages are built in the same way as the finite successor stages.

Later, this hierarchy will be linked to the effective version of the Hausdorff–Kuratowski difference hierarchy of Σ_2^0-sets: a set is in Σ_α^∂ iff it is of the form

$$\bigcup_{\xi \in \mathrm{Opp}(\alpha)} A_\xi \setminus \textstyle\bigcup_{\zeta<\xi} A_\zeta$$

where $(A_\xi)_{\xi<\alpha}$ is an effective enumeration of a $\subseteq$-increasing sequence of Σ_2^0-sets, $\alpha < \omega_1^{ck}$, and $\mathrm{Opp}(\alpha)$ is the set of ordinals $< \alpha$ and of opposite parity to α, where the parity of a limit ordinal is even.

3 Model-Checking for the Transfinite Mu-Calculus

Fixpoints are often calculated by iteration, computing successive approximants until convergence. (Recall that the αth approximant of a least fixpoint $\mu X.\phi$ is defined as $\mu X^\alpha.\phi = \phi((\bigcup_{\beta<\alpha} \mu X^\beta.\phi)/X)$, and dually for greatest fixpoints.) In the finite case, this is the straightforward 'global' algorithm used for model-checking modal fixpoint logics.

The main focus of this paper is the relation between infinite parity games and fixpoint calculation. The games are infinite, and have somewhat complex payoff sets. This also has an analogue in the world of finite modal fixpoint logics, where it corresponds to the use of parity automata. Of course, in the finite world, it is well known that one does not have to play infinite games – repeats can be detected. It is perhaps of some interest to see that even in this infinite world of infinite formulae, it is possible to extend techniques well known from modal mu-calculus, and characterize truth of transfinite mu-arithmetic by a game in which all plays are finite (and therefore the payoff sets are clopen). Of course, there is a small catch – the moves involve playing ordinals, which amounts (for countable structures) to having to make second-order moves. The techniques being extended have somewhat intricate full definitions and proofs, and this is not the main focus of our paper, so given space restrictions we will just outline the game for those who have some familiarity with the modal mu-calculus work (as described for example in [2]).

A concept used in many basic theorems about modal mu-calculus, and in soundness proofs for techniques such as local model-checking with tableaus, is that of *(μ-)signature*. Consider a least fixpoint variable X_1 somewhere inside a formula, and suppose that $n \in X_1$ when X_1 has its actual value. We can consider at which approximant of X_1 the value n enters X_1; suppose this is α_1. However, X_1 itself may be defined in terms of some outer least fixpoint variable X_2; so we also need to know what approximant α_2 is currently being used as the value of X_2. Then a μ-signature of n at some subformula is an assignment of ordinals α_i to all the enclosing least fixpoint operators X_i that makes the given subformula true of n. When one is doing local model-checking, which effectively means exploring the proof tree that justifies $n \in \Phi$, signatures can be thought of as sets of 'clocks' for the verifier: every time verifier passes through a least fixpoint variable X, she has to decrement the clock for that variable, and if the clock ever hits zero, she loses. This ensures that she only passes through least fixpoints finitely often.

Dually, if (as done for example in completeness proofs) we are arguing about the falsity of a formula, we can consider the ν-signatures: now we look at the approximants of the greatest fixpoints at which a formula becomes false – and again, the signatures give a clock to bound the time by which refuter must establish the falsity.

It is possible to combine the clock intuitions for both μ-signatures and ν-signatures into a single game, which works as well for the transfinite logic as for the normal logic. Consider the usual rules for the model-checking game, with the obvious rule for the infinite disjunctions and conjunctions. Now extend it thus: whenever play enters a least fixpoint formula $\mu X.\phi$, verifier gets to choose an ordinal α_X. This amounts to a promise that she will 'bottom out' of the inductive definition in finite time, measured by α. Then any time play passes through a formula $\tau \in X$, verifier must decrease α_X, and if she can't she loses. Similarly, refuter chooses clocks for the greatest fixpoint formulae, and decreases them when passing through the variables.

This amendment to the model-checking game rules forces all plays to be finite, since it can be shown by standard arguments that signatures are well-ordered with respect to the game moves. Thus the payoff sets are simply sets of finite plays (i.e. clopen subsets of ${}^{\omega}\omega$) rather than parity conditions. Then one can show by minor extensions of the standard mu-calculus arguments that indeed verifier/refuter has a winning strategy iff the initial formula is true/false.

4 Parity Games

The use of parity games in model checking has been described by many authors. A very detailed survey is given by Niwiński [9]. Let us mention that we follow the convention that if the maximal priority seen infinitely often is odd, then player I wins. When looking at a formula in the transfinite mu-calculus, we need to play a parity game with infinitely many priorities: for each set variable we need a distinct priority. If we take the binary tree and attach to each node a priority in an arbitrary fashion, then, when playing a parity game on this tree, we might end up having a "wild" payoff set for player I, and we might also lose the nice property of having a memoryless winning strategy [4]. Furthermore, it might be that there is no maximum among the priorities seen infinitely often, and infinite runs might even meet each priority only finitely many times. However, as we will see, a labelling derived from a model checking game of a transfinite mu-calculus formula avoids all these undesired effects. Moreover, such a labelling describes some set of the transfinite difference hierarchy $\bigcup_{\alpha<\omega_1^{ck}} \Sigma^\partial_\alpha$ and vice versa.

5 Connecting the Transfinite Difference Hierarchy and the Transfinite Mu-Calculus

Our aim is to extend Bradfield's following theorem [1]:

Theorem 2. *For every natural number n the equality $\Game\Sigma^\partial_n = \Sigma^\mu_{n+1}$ holds true.*

The extension into the transfinite is our main result.

Theorem 3. *For every recursive ordinal α the equality $\Game\Sigma^\partial_\alpha = \Sigma^\mu_{\alpha+1}$ holds true. Thus, $\bigcup_{\alpha<\omega_1^{ck}} \Sigma^\mu_\alpha = \Game\Delta^0_3$.*

Proof. Let α be a recursive limit ordinal, and let $\mu X_{\alpha+1}. \bigvee_{i<\omega} \Phi_i \in \Sigma^\mu_{\alpha+1}$, so in particular each Φ_i is in some Σ^μ_β for some $\beta < \alpha$. We need to find a game with payoff set in Σ^∂_α whose winning positions for player I are calculated by this formula.

Assume that the formula $\mu X_{\alpha+1}. \bigvee_{i<\omega} \Phi_i$ describes a nonempty subset of ω, and choose some witness n for this nonemptyness. Now consider the game tree which results from the parity game played as a model checking game. We might think of it as a subtree in ω^*, each node labelled with the position in the model checking game. In extending the tree in an appropriate way we may assume that it does not contain finite maximal branches, and in further simplifying the tree we may assume that each node marks a loop-back, i.e. we see some X_β at each node. This can be done because any infinite branch must hit such nodes infinitely many times. In omitting which element $n' \in \omega$ is inspected in the model checking game, we get a tree which is simply labelled by the indices of the set variables, i.e. by countable ordinals up to $\alpha + 1$ without limit ordinals. Observe that the labelling has the structure of a set in $\Sigma^\partial_{\alpha+1}$. Let us describe the payoff set for player I.

Since the outmost variable is under the scope of a minimal fixpoint operator, player I wins the model checking game iff at some point, the game gets captured inside some subformula Φ_i and player I wins the subgame for Φ_i, where $X_{\alpha+1}$ is replaced by $\perp$. By the induction hypothesis, the payoff set of such a subgame has the complexity of some set in $\Sigma^{\partial}_{\beta}$ with $\beta < \alpha$. Since the payoff set of the whole game is the effective union of all these subsets, we obtain some set of complexity $\Sigma^{\partial}_{\alpha}$, which shows $\Sigma^{\mu}_{\alpha+1} \subseteq \partial\Sigma^{\partial}_{\alpha}$ for α limit. The successor case is an induction over the construction of formulae as in [1].

It remains to show $\partial\Sigma^{\partial}_{\alpha} \subseteq \Sigma^{\mu}_{\alpha+1}$. As before, we only need to consider the limit step, the successor case is done as in [1].

Let α be a countable limit ordinal, and let $A \in \Sigma^{\partial}_{\alpha}$. We may assume that $A = \bigcup_{i<\omega} A_i$ with $A_i \in \Sigma^{\partial}_{\alpha_i}$, $\alpha_i < \alpha$. We let A be the payoff set for player I and calculate her winning positions. By induction hypothesis, for each A_i we have a formula $\Phi_i \in \Sigma^{\mu}_{\beta}$, $\beta < \alpha$, describing the winning positions for player I with the payoff set A_i. Let

$$H_0 = \{s \in \omega^* \mid \exists i \; s \in \Phi_i\}$$

be the set of all nodes s.t. player I wins if some A_i is the payoff set. By recursion we define H_β for $\beta < \omega_1$: If β is a limit ordinal, then let $H_\beta = \bigcup_{\gamma<\beta} H_\gamma$, and if $\beta = \gamma + 1$, then let

$$H_\beta = \{s \in \omega^* \mid \exists i \text{ player } I \text{ wins with payoff set } A_i \text{ or she can reach } H_\gamma\}$$

To reach a certain set of nodes is an open condition, thus, H_β can be viewed as the set of winning conditions of a set of complexity less than $\Sigma^{\partial}_{\alpha}$. It is immediate from the definition that the H_β's describe an increasing sequence of subsets of a countable set. Therefore, at some countable stage ξ the process stabilizes, we have reached the minimal fixpoint H_ξ of this process. Thus, we can express the calculation of the winning positions within the transfinite mu-calculus, within complexity $\Sigma^{\mu}_{\alpha+1}$:

$$\begin{aligned}\mu X_{\alpha+1}.\quad & (\exists n \exists m \; \mathrm{lh}(x_{\alpha+1}) = 2n \wedge x_{\alpha+1}{}^\frown m \in X_{\alpha+1}) \\ & \vee (\exists n \forall m \; \mathrm{lh}(x_{\alpha+1}) = 2n+1 \wedge x_{\alpha+1}{}^\frown m \in X_{\alpha+1}) \\ & \vee (\textstyle\bigvee_{i<\omega} \Phi_i)\end{aligned}$$

By the same arguments used in [1], as a result that Borel games are determined [8] H_ξ describes exactly the winning positions for player I.

6 Nicely Behaving Labellings

When extending the mu-arithmetic into the transfinite we need to check whether we keep key properties, namely the existence of positional winning strategies. This leads to

Definition 4. Let P be a parity game with priorities in some $\alpha < \omega_1$. P is called max-closed iff for every infinite run the set of all labels seen infinitely often is non-empty and contains a maximum.

Clearly, the rules of the model checking game ensure that the parity game derived from a model checking game of a transfinite mu-calculus formula is max-closed.

Theorem 5. *Each max-closed parity game admits a positional winning strategy for one of the players.*

Proof. We proceed by induction on the set of labels α. Of course, any set of countably many labels can be relabelled by natural numbers, but max-closedness is not preserved in general. In the sequel l will always denote the labelling function, $l : \alpha \to V$ where V is the set of vertices in the considered game graph.

Let us first consider the easier case, i.e. α is a limit ordinal. Assume player I has a winning strategy f, we need to find a positional winning strategy.

Let T be the tree of all possible plays. We define

$$A_0 = \{s \in T \mid \exists \beta < \alpha . \forall t \in T[s] . l(t) \leq \beta\}$$

i.e. the cone of T above s is labelled with values up to β. Observe that by max-closedness, A_0 is dense in T. Otherwise, we could select a cone $T[t]$ having an empty intersection with A_0, meaning that every subcone of this cone is labelled with values cofinal in α. Since α is countable, there is a sequence $(\alpha_i)_{i<\omega}$ with each $\alpha_i < \alpha$ and $\bigcup_{i<\omega} \alpha_i = \alpha$. In the cone $T[t]$ it is easy to construct an infinite path x s.t. for each i there is some n_i with $l(x(n_i)) > \alpha_i$, contradicting the max-closedness.

Although A_0 is dense, it might be that the complement still contains infinite paths. Thus, we define by recursion:

$$A_\beta = \{s \in T \setminus \bigcup_{\alpha<\beta} A_\alpha \mid \exists \gamma < \alpha . \forall t \in T[s] \setminus \bigcup_{\alpha<\beta} A_\alpha . l(t) \leq \gamma\}$$

The process stops at some countable γ. From these sets we can easily determine the set of winning positions for player I. We let H_0 be the set of all elements in A_0 such that player I can win the game starting at that position. Since the labels in the cone of the game tree are bounded by some $\beta < \alpha$, by induction hypothesis player I has a positional winning strategy within H_0. In particular, the game stays within H_0. In general we let H_β be the subset of A_β such that player I has a winning strategy as long as the game stays within A_β, and as soon as the game leaves A_β, some $H_{\beta'}$ is entered with $\beta' < \beta$. Again, within H_β player I has a positional winning strategy. Analogously to Section 5 the process stabilizes at some countable γ, and $H_\gamma = \bigcup_{\beta<\gamma} H_\beta$ is the set of all winning positions of player I, and it can be described by a formula of the transfinite mu-arithmetic provided the set of labels does not exceed ω_1^{ck}. It is fairly easy to describe a positional winning strategy for player I: as long as the game takes place in some H_β, she follows the positional winning strategy within H_β. It might be that player I cannot force the game to stay inside H_β, but if this set is left, then some $H_{\beta'}$ is entered with $\beta' < \beta$, and from that moment on player I follows the positional winning strategy for $H_{\beta'}$. Since all the H_β are pairwise disjoint, the concatenation of all positional winning strategies for the H_β gives a positional winning strategy for all her winning positions.

Analogously, if player *II* has a winning strategy, then he has a positional winning strategy as well.

Now let us consider the successor case. Assume $\alpha = \beta+1$ is odd, thus player *I* needs to make sure that β is seen only finitely many times. Assume player *I* has a winning strategy, we need to find a positional winning strategy.

Let H_0 consist of those vertices s such that, starting from s, player *I* has a winning strategy which never leads to any vertex labelled with β. Such vertices must exist, otherwise player *II* has a winning strategy. In particular, being at such a node player *I* can force the game to stay in H_0. We claim that within H_0 player *I* has a positional winning strategy. Consider the game tree starting from $s \in H_0$ and remove all nodes outside H_0 together with the cones above those nodes. The remaining tree is labelled with values smaller than β, and by induction hypothesis on this subtree (and the corresponding subgraph) player *I* has a positional winning strategy. This positional winning strategy is clearly a winning strategy for the whole game. Constructing H_γ analogously to the limit case yields a positional winning strategy for the whole game.

Now assume player *II* has a winning strategy. This means that he either can manage to see β infinitely often, or, if player *I* keeps the occurrence of β finite, he wins the induced subgame. A positional strategy is described as follows: if a vertex belongs to player *II*'s winning region, and if he has a winning strategy which guarantees him to reach some vertex labelled with β, then he plays in a way that he never leaves his winning region and after finitely many steps he will reach β. Clearly, to reach some node within the winning region labelled with β is an open condition, thus there exists a positional strategy for achieving that goal. If, after reaching β, player *II* can still reach another β within his winning region, he goes for it. At some point it might be that he still has a winning strategy, but he cannot make sure that β is seen again. At this stage consider the subgraph S wich consists of all nodes in player *II*'s winning region with labels smaller than β, the edge relation restricted to S stays the same. Observe that by being a winning region player *I* can only leave the subgraph S in moving to a vertex which is still in player *II*'s winning region, but from where player *II* can reach β memoryless again while remaining in his winning region. As long as the run stays in S, by induction hypothesis player *II* has a positional winning strategy. Thus, in concatenating the positional winning strategies for the different regions we obtain a positional winning strategy for the whole game.

The remaining case, $\alpha = \beta + 1$ even, is handled similarly. We can construct a positional winning strategy for player *I* as in the odd case for player *II* and vice versa. □

Corollary 6. *For any formula in the transfinite mu-arithmetic, model checking with parity games admits positional winning strategies.*

7 A Descriptive Set Theoretical Approach

In this section we step into Descriptive Set Theory, and see that the effective version of the very refined Wadge Hierarchy of sets of infinite words unveils a clue to understanding what goes on in the result $\partial\Sigma^{\partial}_{n} = \Sigma^{\mu}_{n+1}$. The reader may find most of the basic material in [5]. The effective Wadge hierarchy is studied in [6].

7.1 The Difference Hierarchy of $\mathbf{\Sigma}^0_2$ Sets

This result admits a first step: $\Sigma^{\mu}_{1} = \partial\Sigma^0_1$. Recall that Σ^0_1 is the class of all semi-recursive sets, Σ^{∂}_{n} being the class of n differences of effectively countable unions of complements of semi-recursive sets. To be more precise, the space of infinite sequences over an alphabet Σ is equipped with the usual topology, that is the product topology of the discrete topology over the alphabet. So, Σ^0_1 is the class of all sets of the form $W\Sigma^{\omega}$ where $W \subseteq \Sigma^*$ is a recursively enumerable set of finite words (possibly empty in which case $W\Sigma^{\omega} = \varnothing$). And Σ^{∂}_{n} stands for the the class of sets of the form $A = A_n \setminus A_{n-1} \ \cup \ A_{n-2} \setminus A_{n-3} \ \cup \ldots$, where $A_1 \subseteq A_2 \subseteq \ldots \subseteq A_n$ is a sequence of sets in Σ^0_2 - the class of effectively countable unions of complements of semi-recursive sets.

So, as we see $\Sigma^{\mu}_{1} = \partial\Sigma^0_1$, $\Sigma^{\mu}_{2} = \partial\Sigma^0_2$, but then, $\Sigma^{\mu}_{n} = \partial\Sigma^0_n$ fails for $n > 2$, and must be replaced with $\Sigma^{\mu}_{n} = \partial\Sigma^{\partial}_{n-1}$. At first glance it seems there is no logic behind this. However, the effective Wadge hierarchy, a refinement of the effective difference hierarchy, gives the solution.

The Wadge Ordering A natural improvement of the Hausdorff–Kuratowski hierarchy was induced by Wadge's work based on a reduction relation defined in terms of continuous functions. This means, a natural way to compare the topological complexity of sets A and B was to say $A \leq_W B$ – intuitively meaning A is topologically less complicated than B – if the problem of knowing whether x belongs to A *reduces* to knowing whether $f(x)$ belongs to B for some *simple* function, where simple meant continuous. The effective version deals with recursive functions instead, and in the sequel we will concentrate only on the effective version:

$$A \leq_W B \quad \text{iff} \quad \exists \text{ recursive } f : \Sigma^{\omega}_{A} \to \Sigma^{\omega}_{B}.\, f^{-1}B = A$$

The Wadge ordering ($\leq_W$) induces the strict ordering ($<_W$) and the Wadge equivalence ($\equiv_W$):

$$A <_W B \quad \text{iff} \quad A \leq_W B \wedge B \not\leq_W A$$

$$A \equiv_W B \quad \text{iff} \quad A \leq_W B \leq_W A$$

When restricted to Kleene pointclasses, this ordering becomes a quasi-well-orderingd, i.e. it is well-founded, and has antichains of length at most two. Moreover, if A and B are incomparable, then $A \equiv_W B^{\complement}$. The reason for this is that all these properties derive from Borel Determinacy [8]. Indeed, Wadge defined the relation $A \leq_W B$ in terms of the existence of a winning strategy in a suitable game: the Wadge game.

Definition 7 (The Wadge game). Let $A \subseteq \Sigma_A^\omega$, $B \subseteq \Sigma_B^\omega$, $\mathbf{W}(A,B)$ is an infinite two-player game where players (I, and II) take turn playing letters in Σ_A for I, and in Σ_B for II. As opposed to I, player II is allowed to skip provided he plays infinitely many letters.

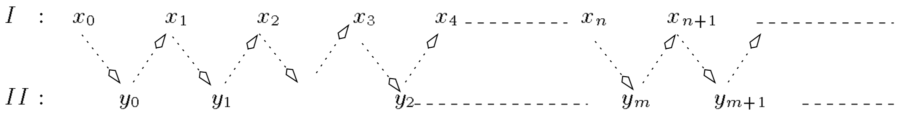

So that at the end of a run (in ω moves), I has produced an ω-word $x \in \Sigma_A^\omega$ and II has produced $y \in \Sigma_B^\omega$. The winning conditions are:

$$II \text{ wins } \mathbf{W}(A,B) \quad \text{iff} \quad (x \in A \Leftrightarrow y \in B)$$

Wadge designed the rules of the game $\mathbf{W}(A,B)$ so that a strategy for II induces a continuous mapping $x \mapsto y$, and conversely; and the winning condition so that

$$II \text{ has a w.s. in } \mathbf{W}(A,B) \quad \text{iff} \quad A \leq_W B.$$

Let us define the equivalence relation $\sim$ by

$$A \sim B \quad \text{iff} \quad A \equiv_W B \ \text{ or } \ A \equiv_W B^{\complement} \ \text{ or } \ A \equiv_W 0B \cup 1B^{\complement}$$

Quotiented by $\sim$, and using determinacy, the Wadge ordering $\leq_W$ turns into a well-ordering (denoted by $\leq_{/\sim}$) whose minimal elements are all clopen sets. This induces the notion of the Wadge degree defined inductively:

$$d^\circ A = 0 \quad \text{iff} \quad A \text{ is clopen}$$

$$d^\circ A = \sup\{d^\circ B + 1 : \ B <_{/\sim} A\}$$

where $<_{/\sim}$ stands for the strict Wadge ordering $<_W$ quotiented by $\sim$.

7.2 Multiplication by ω_1^{ck}

Now, given a topological class, that is a class closed under pre-image by recursive functions (such as Σ_1^0, Σ_n^∂), a set A is complete for the class if it reduces all sets in it. As usual, a complete set is a set of maximal complexity, therefore of maximal Wadge degree. In other words, the Wadge degree of a complete set of a given class is a measure of the topological complexity of this class.

If we look at the sequence of Wadge degrees of complete sets for respectively Σ_1^0, $\Sigma_2^0 = \Sigma_1^\partial$, Σ_2^∂, Σ_3^∂,... We find $1, \omega_1^{ck}, {\omega_1^{ck}}^2, {\omega_1^{ck}}^3, \ldots$. Surprisingly, the progression is precisely multiplication by ω_1^{ck}. More surprisingly indeed, is that multiplication of a Wadge degree by ω_1^{ck} ($\alpha \longmapsto \alpha \cdot \omega_1^{ck}$) corresponds to a simple set theoretical operation ($A \longmapsto A \bullet \omega_1^{ck}$). Namely,

$$A \bullet \omega_1^{ck} = (\Sigma \cup \{a_+, a_-\})^* a_+ A \ \cup \ (\Sigma \cup \{a_+, a_-\})^* a_- A^{\complement}$$

for a_1, a_- two different letters not in Σ. For a better understanding, a player (either I or II) in charge of $A \bullet \omega_1^{ck}$ in a Wadge game is exactly like the same player being in charge of A with the extra possiblity to erase all his moves and decide to start all over again being in charge of $A^{\complement}$ instead of A, and erase everything again and switch from A to $A^{\complement}$, and so on. Playing a_+ or a_- takes care of both the initialization of the play, and the choice between A and $A^{\complement}$. A word containing infinitely many a_+ or a_- being not in $A \bullet \omega_1^{ck}$.

This operation preserves the Wadge ordering

$$A \leq_W B \Rightarrow A \bullet \omega_1^{ck} \leq_W B \bullet \omega_1^{ck}$$

and satisfies the required property:

$$d^\circ(A \bullet \omega_1^{ck}) = d^\circ(A) \ \cdot \omega_1^{ck}$$

7.3 Division by ω_1^{ck}

For inductive proofs on the degree of sets, an inverse operation ω_1^{ck} is needed. It must be a set theoretical counterpart of division by ω_1^{ck}: given any set A the set $\frac{A}{\omega_1^{ck}}$ must verify:

1. $A \leq_W B \Rightarrow \frac{A}{\omega_1^{ck}} \leq_W \frac{A}{\omega_1^{ck}}$
2. $d^\circ A = \alpha \cdot \omega_1^{ck} \ \Rightarrow \ d^\circ \frac{A}{\omega_1^{ck}} = \alpha$

Unfortunately, we do not know how to obtain condition 2 directly. So we weaken our expectations and ask for the following instead:

$$d^\circ A = \alpha \cdot \omega_1^{ck} \ \Rightarrow \ d^\circ \frac{A}{\omega_1^{ck}} = \alpha + 1$$

How to get precisely condition 2 can be deduced from the whole artillery Duparc developed in [3]. The idea to define $\frac{A}{\omega_1^{ck}}$ is that a player in a Wadge game in charge of $\frac{A}{\omega_1^{ck}}$ is like this player being in charge of $A \subseteq \Sigma^\omega$ but having his opponent asking him questions whether or not the infinite word x he is actually constructing step by step, will remain in a tree $T_i \subseteq \Sigma^*$. The opponent is allowed to ask questions about as many trees as he wants as long as the player answers no. Once the player answers yes, there is no more questioning allowed. The precise definition is as follows. Given $x \in \Sigma^\omega$, we write x_{even} for the word $x(0)x(2)x(4)\ldots$.

Definition 8. Given an alphabet Σ, a $\mathscr{T}$ree $\mathscr{T}$ on Σ is some *non empty pruned tree* that satisfies for any $u \in \mathscr{T}$ and any integer $n < \mathrm{lh}(u)$:

if n is even: then $u(n) \in \Sigma$ (these are the nodes that correspond to the main run), and

if n is odd: then $u(n)$ is an auxiliary move with three different options:

Option $\langle \texttt{no} \rangle$: in this case $u(n) = \langle \texttt{no}, v\rangle$ for some $v \in \Sigma^*$ and $u_{even} \subsetneq v$. So v is some position in Σ^* that extends the position u_{even}. But then, we demand that any position w in $\mathscr{T}$ that extends u must verify $w_{even} \subsetneq v$ or $v \subsetneq w_{even}$. Moreover, we also require that $\mathscr{T}$ verifies the following condition: if $(u \restriction n) \langle \texttt{no}, v'\rangle$ belongs to $\mathscr{T}$ with $v' \neq v$, then both $v' \subsetneq v$ and $v \subsetneq v'$ must fail. Or, to say it differently, T must satisfy the condition:

$$\Big(u \in \mathscr{T} \ \wedge \ u \langle \texttt{no}, v\rangle \in \mathscr{T} \wedge \ u \langle \texttt{no}, v'\rangle \in \mathscr{T}\Big) \Rightarrow v = v' \ \vee \ v \perp v'.$$

Option $\langle \texttt{yes} \rangle$: in this case $u(n) = \langle \texttt{yes} \rangle$. This must be regarded as the option to avoid all other positions of the form $(u \restriction n) \langle \texttt{no}, v\rangle$. Formally, this means that any w in $\mathscr{T}$ that extends $(u \restriction n) \langle \texttt{yes} \rangle$ must satisfy

$$v \subsetneq w_{even} \textbf{ fails for any } v \text{ such that } (u \restriction n) \langle \texttt{no}, v\rangle \in \mathscr{T}.$$

Option $\langle - \rangle$: this case should be regarded as no question asked at all. We require $\mathscr{T}$ to satisfy:

if $(u \restriction n) \langle - \rangle \in \mathscr{T}$
then $(u \restriction n) \langle \texttt{yes} \rangle \notin \mathscr{T}$ and for any $v \in \Sigma^*$ $(u \restriction n) \langle \texttt{no}, v\rangle \notin \mathscr{T}$.

Finally, $\mathscr{T}$ must verify:

$$\forall k \ \forall n > k \ \ x(2k+1) = \langle \texttt{yes} \rangle \ \Rightarrow \ x(2n+1) = \langle - \rangle$$

We remark that given any infinite word $x \in \Sigma^\omega$, there exists a unique infinite branch $y \in [\mathscr{T}]$ such that $y_{even} = x$.

Definition 9. Let $A \subseteq \Sigma^\omega$, and $\mathscr{T}$ be a $\mathscr{T}$ree on Σ,

$$A^{\mathscr{T}} = \{x \in [\mathscr{T}] : \ x_{even} \in A\}$$

$$\frac{A}{\omega_1^{ck}} = \text{ a } \leq_{/\sim}\text{-minimal element in } \{A^{\mathscr{T}} : \ \mathscr{T} \text{ a } \mathscr{T}\text{ree on } \Sigma\}$$

Remark 10.

1. If $A \leq_W B$ then for any $\mathscr{T}$ree $\mathscr{T}_B$, one can easily design a $\mathscr{T}$ree $\mathscr{T}_A$ such that $A^{\mathscr{T}_A} \leq_W B^{\mathscr{T}_B}$. By minimality, $\frac{A}{\omega_1^{ck}} \leq_{/\sim} \frac{B}{\omega_1^{ck}}$.
2. If $d^\circ A = \alpha \cdot \omega_1^{ck}$, then take any B with $d^\circ B = \alpha$, consider $B \bullet \omega_1^{ck} \subseteq (\Sigma \cup \{b_+, b_-\})^\omega$. Let $\mathscr{T}$ be the $\mathscr{T}$ree that asks questions about the trees $(\Sigma^* \{b_+, b_-\})^n \Sigma^*$ for each n as long as the opponent does not agree on restricting his moves to such a tree. Clearly

$$\frac{A}{\omega_1^{ck}} \leq_{/\sim} \left(B \bullet \omega_1^{ck}\right)^{\mathscr{T}} \leq_W \bigcup_{n \in \mathbb{N}} 0^{2n} B \cup 0^{2n+1} B^{\complement} \leq_{/\sim} B$$

This shows

$$d^\circ A = \alpha \cdot \omega_1^{ck} \;\Rightarrow\; d^\circ \frac{A}{\omega_1^{ck}} \leq \alpha + 1.$$

The other inequality is by induction on $d^\circ B$ and requires the complete knowledge of the Wadge hierarchy [3].

8 Outlook

The transfinite mu-arithmetic gives us a class of parity games with infinite labellings, but still they are well behaving, there are positional winning strategies. On the other hand, there are examples of automata which induce parity games without positional winning strategies, even requiring an infinite memory for a winning strategy [4]. It is interesting to draw the line between parity games with or without positional strategies sharper. Clearly, if there is only one path in the game tree which is a counter-example of max-closedness, then there is still a positional winning strategy, since this one path can always be left. Thus, the question arises how many ill behaving paths a game tree can allow and still having a positional winning strategy. Which is the right notion of smallness of the set of all ill-behaving paths, maybe meagreness?

9 Acknowledgement

Quickert is supported by the EU Research and Training Network GAMES (Games and Automata for Synthesis and Validation).

References

1. J. Bradfield. Fixpoints, Games and the Difference Hierarchy. *Theoretical Informatics and Applications*, 37:1–15, 2003.
2. 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 293–329. Elsevier Science B.V., 2001.
3. J. Duparc. Wadge Hierarchy and Veblen Hierarchy, Part I: Borel sets of finite rank. *Journal of Symbolic Logic*, 66:56–86, 2001.
4. E. Grädel and I. Walukiewicz. Positional Determinacy of Games with Infinitely Many Priorities. *Manuscript.*
5. A. Kechris. *Classical descriptive set theory.* Springer Verlag, 1995.
6. A. Louveau. Some results in the Wadge hierarchy of Borel sets. In A. Kechris, D. Martin, and Y. Moschovakis, editors, *Cabal Seminar 79-81*, Lecture notes in Mathematics, pages 28–55. Springer, 1983.
7. R. S. Lubarski. μ-definable sets of integers. *Journal of Symbolic Logic*, 58:291–313, 1993.
8. D. Martin. Borel determinacy. *Annals of Mathematics*, 102:336–371, 1975.
9. D. Niwiński. Fixed point characterisation of infinite behavior of finite state systems. *Theoretical Computer Science*, 189:1–69, 1997.