Some remarks on transfinite fixpoint alternation

J. C. Bradfield

LFCS,[2] Division of Informatics, University of Edinburgh, Edinburgh, United Kingdom, EH9 3JZ (email: jcb@dcs.ed.ac.uk)

1 The setting

The connexion between Rabin automata, parity games and fixpoint alternation in modal logic is fundamental to understanding all three, and an essential tool in using and analysing any of the three. The connexion between parity games and fixpoint alternation in particular is an instance in the modal world of a more general relationship in arithmetic between nested (co-)inductive definitions and certain rather simple Gale–Stewart games. The first two levels of this connexion were established long ago, but the extension to all finite levels was, to my knowledge, first remarked only in 1999 in my [Bra99], although several infinitely more complex such results are found in effective descriptive set theory. Of course, in arithmetic, once one has an ω hierarchy, it is natural to try to extend into the transfinite; in [Bra99] I conjectured that this could be done in an "obvious" way. I recently took up this conjecture, and discovered that it is not so easy – and perhaps not doable. Ideas from workshop participants would be most welcome!

In explaining the setting, I assume some standard notations, all defined in [Bra99].

If T_ζ is any increasing (transfinite) sequence of sets or classes, let $T_{<\zeta} = \bigcup_{\xi<\zeta} T_\xi$.

Mu-arithmetic ($L\mu$) is first-order arithmetic augmented with a fixpoint constructor $\mu(x, X).\,\phi$ (and its dual ν) where X is a set variable, interpreted in the usual way. Σ^μ_n and Π^μ_n denote the levels of the *fixpoint alternation hierarchy*, defined analogously to the Kleene hierarchies by counting alternation of μ and ν in the prenex normal form of a formula defining a set.

In a *game* $P \subseteq {}^\omega\omega$ on ω, two players take turns (where the turns need not alternate, but can depend recursively on the history so far) to choose integers $a_0, a_1, \ldots$, subject to recursive rules restricting their choice. If the resulting sequence $\alpha = a_0 a_1 \ldots$ is in P, then the first player (Eloise) wins the play. Eloise *wins* P if she has a winning strategy for the game P. If $Q(\alpha, x) \subseteq {}^\omega\omega \times \omega$, then $\Game\alpha.Q(\alpha, x)$ is the set $\{\, x \mid \text{Eloise wins the game } Q(\alpha, x) \,\}$.

Given a class $\mathscr{S}$ of sets, the *finite difference hierarchy* over $\mathscr{S}$ is given by $\partial_1(\mathscr{S}) = \mathscr{S}$, $\partial_{n+1}(\mathscr{S}) = \mathscr{S} \wedge \neg\partial_n(\mathscr{S})$. The class $\bigcup_n \partial_n(\mathscr{S})$ is equal to the closure of $\mathscr{S}$ under complementation and intersection, so ∂ is one way of stratifying the boolean closure of $\mathscr{S}$. Let Σ^∂_n denote $\partial_n(\Sigma^0_2)$ (and $\Sigma^\partial_0 = \Sigma^0_1$).

The main theorem of [Bra99] is

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

This theorem has three grounds to be interesting: it is the natural generalization of the relationship between modal fixpoint alternation and index of parity games/automata; it characterizes the power of alternating inductive/co-inductive definitions in a much more transparent way than the previous (technically very complex) characterization of Lubarsky; conversely it characterizes the power of (Boolean closure of Σ^0_2) games without invoking higher-type recursion. The last ground is interesting because the only characterization of $\Game\Sigma^0_3$ appears necessarily to require higher types, and the result pushes the boundary before higher type a bit further towards Σ^0_3.

It is then natural to wonder if this result can be pushed farther, by finding a transfinite extension of it.

2 Transfinite fixpoint alternation?

The first requirement in extending such a result is to formulate a transfinite extension of the fixpoint hierarchy.

The natural way to extend one step beyond the finite fixpoint formulae is to allow recursive unions (or intersections) of fixpoint formulae: let $R: \omega \to L\mu$ be a recursive enumeration of formulae of $L\mu$, and say that for any recursive $f: \omega \to \omega$, $\bigvee_{i<\omega} R(f(i))$ is a formula. An enumeration of these formulae is then given via an enumeration of the f. If we now close under positive first-order operators and least fixpoint, we have an apparently plausible definition of Σ^μ_ω, which is clearly more expressive than $L\mu$.

Consider such a formula $\tau \in \mu(x, X). \bigwedge_i \phi_i$, where, say, each ϕ_i is Π^μ_i, of the form $\tau^0_i \in \nu Z^0_i.\tau^1_i \in \mu Z^1_i \ldots$. The satisfaction game for this formula has the usual rules; what is the complexity of its winning conditions? By extension of the usual arguments, the formula is true if Eloise can win the game where on every play the outermost (least j) variable seen infinitely often is a maximal fixpoint variable; that is, is a $Z^{j_0}_i$ for some even j_0. Thus the winning condition has the form $\exists i_0, j_0.$'$Z^{j_0}_{i_0}$ inf. often' $\wedge\ \forall j < j_0.\forall i.$'$Z^j_i$ fin. often' $\wedge$ 'X fin. often'. *Prima facie* the first clause is Σ^0_3 and the second Π^0_3. However, things are not that bad. Because on any winning run X is unfolded finitely often (third clause), the run eventually gets stuck inside one conjunct; thus the second clause can be replaced by $\forall j < j_0.$'$Z^j_{i_0}$ fin. often', which is Σ^0_2. Moreover, for the same reason, the $\exists i_0$ outside the first clause is unnecessary; there can be only one such i_0, and it suffices to check that 'Z^{j_0} (ignoring i) occurs', which is a recursive property of the game positions, happens infinitely often. Hence the winning condition boils down to 'X fin. often' $\wedge$ 'least j seen inf. often is even'. Considering X as Z^{-1}, this is an infinite parity condition. How is it related to the difference hierarchy?

Transfinite difference hierarchies over the open sets were studied almost a century ago by Hausdorff; in the 60s an abstract study of difference hierarchies and 'alternating chains' was undertaken by Addison, although regrettably most of the results on transfinite hierarchies were in a planned paper that never appeared (to my knowledge). Hinman and Burgess also used difference hierarchies to analyse Kolmogorov's R-sets. Part of my hope in trying to extend Theorem 1 is to find the parts of that work that have somewhat simpler formulations analogous to those in the theory of automata, where transfinite hierarchies have been studied by Büchi, Landweber, Barua, Selivanov, and others.

There are several ways to formulate the definitions. One is: Given a class $\mathscr{S}$ of sets, and a sequence $S_0, S_1, \ldots$ of sets in $\mathscr{S}$, the *difference kernel* $\partial(S_0, \ldots)$ is the set of x such that the least i such that $x \notin S_i$ is odd (extending the sequence by an empty set if necessary). For example, $\partial(X, Y, Z) = X \wedge (\neg Y \vee Z)$.) This extends to transfinite sequences, where limit ordinals are even. Then $\partial_\zeta(\mathscr{S})$ is the class of sets representable as a difference kernel over $\mathscr{S}$ of length ζ. This definition coincides at finite levels with the earlier definition, for reasonable $\mathscr{S}$.

The following is the analogue of the relationship between parity (Mostowski) acceptance conditions for automata, and Rabin chain conditions: Let $S_0, \ldots$ be a sequence. Define $R_\zeta = \bigcap_{\xi \leq \zeta} S_\xi$. Then $\partial(S_0, \ldots) = \bigcup_\zeta (R_{2\zeta} - R_{2\zeta+1})$.

Example: let $S_0, S_1, \ldots, S_\omega$ be a sequence of length $\omega+1$. The difference kernel of the sequence is

$$\big(\bigcup_i (\bigcap_{j\le i} S_{2j} - \bigcap_{j\le i} S_{2j+1}\big) \vee (\bigcap_{i<\omega} S_i \wedge S_\omega).$$

The following 'well-known' result on the transfinite difference hierarchy over Kleene pointclasses is why I hoped to extend Theorem 1:

Fact 2 $\partial_{<\omega_1}(\Sigma^0_n) = \Delta^0_{n+1}$.

and in particular, $\Sigma^{\partial}_{<\omega_1} = \Delta^0_3$. (The bold-face version (from Hausdorff) $\partial_{<\aleph_1}(\boldsymbol{\Sigma}^0_2) = \boldsymbol{\Delta}^0_3$ has been used in the work of Burgess mentioned above.) So if Theorem 1 extended neatly into the transfinite, we would have a characterization of $\Game\Delta^0_3$ in terms of transfinite fixpoint iteration, as close to $\Game\Sigma^0_3$ as we could hope.

Now, returning to the satisfaction game above, take S_{j+1} to be the set of game plays in which Z^j occurs finitely often, and S_0 to be plays where X occurs finitely often; the infinite parity condition is the difference kernel of the ω-sequence $S_0, S_1, \ldots$.

This is somewhat disturbing, because now consider a formula

$$\sigma \in \nu(y,Y).\tau \in \mu(x,X).\bigwedge_i \phi_i,$$

with ϕ_i as before, which is in what we would naturally call $\Pi^\mu_{\omega+1}$.

The satisfaction game for this formula is just as before, except that now Y may occur infinitely often on a winning run; taking S_{-1} to be plays where Y occurs finitely often, this is *also* the difference kernel of an ω-sequence $S_{-1}, S_0, \ldots$. Hence there is no obvious difference between the complexity of the winning conditions for the satisfaction games for these two formulae. This raises the obvious question: does this naive fixpoint alternation collapse at ω? This is not the case: Σ^μ_ω etc. can define their own truth sets. The other question is, what, if any, role do the levels of the difference hierarchy beyond ω play in fixpoint alternation?

3 What use is transfinite fixpoint alternation, anyway?

Does this have anything to do with computer science? It is received wisdom that nobody (in verification) needs more than two fixpoints! A counter-argument is that fixpoint alternation and parity conditions on alternating automata are at least as interesting as on normal automata, where they are agreed to be computer science, and fixpoints in arithmetic have been useful in establishing results in modal fixpoint theory. Moreover, it is possible to conceive at least of unbounded, and possibly of infinite, fixpoint alternation in straight modal logic: imagine an unbounded queue of interacting agents, cooperating with a scheduler which schedules them fairly but with priority to the earlier agents. No sane person would specify this in modal mu-calculus, but we might...

References

[Bra99] J. C. Bradfield, Fixpoint alternation and the game quantifier, *Proc. CSL '99*, (LNCS 1683) 350–361 (1999).

(*q.v.* for other references)

Version Control: new.tex:1.4.